



Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria


(A Module of TSGR, FR-440; FR-SONET-17; and FD-29)

Telcordia Technologies Generic Requirements

GR-253-CORE

Issue 4, December 2005

Comments Requested (See Preface)



Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria

This document replaces: GR-253-CORE, Issue 3.

Where material has been added, changed, or deleted, the location of the change is marked by a vertical bar (|) in the outer margin next to the change.

For more information on this document, please contact:

Tom Bowmaster

Telcordia — GR-253-CORE
One Telcordia Drive
Piscataway, NJ 08854-4156

Phone: + 1.732.699.5489
E-mail: tbowmast@telcordia.com

To obtain copies of this document, contact your company's document coordinator or your Telcordia account manager, or call 1.866.672.6997 (USA) or + 1.732.699.6700 (Worldwide), or visit our Web site at <http://telecom-info.telcordia.com>.

Copyright © 1995-2000, 2005 Telcordia Technologies, Inc. All rights reserved.

Trademark Acknowledgments

Telcordia is a registered trademark of Telcordia Technologies, Inc.
NEBS and CLEI are trademarks of Telcordia Technologies, Inc.
All other brand or product names are trademarks of their respective companies or organizations.

Telcordia Technologies Generic Requirements Notice of Disclaimer

This Generic Requirements document (GR) is published by Telcordia Technologies to inform the industry of the Telcordia view of proposed generic requirements for SONET transport systems. The generic requirements contained herein are subject to review and change, and superseding generic requirements regarding this subject may differ from those in this document. Telcordia reserves the right to revise this document for any reason (consistent with applicable provisions of the Telecommunications Act of 1996 and applicable FCC rules).

TELCORDIA MAKES NO REPRESENTATION OR WARRANTY, EXPRESSED OR IMPLIED, WITH RESPECT TO THE SUFFICIENCY, ACCURACY, OR UTILITY OF ANY INFORMATION OR OPINION CONTAINED HEREIN.

TELCORDIA EXPRESSLY ADVISES THAT ANY USE OF OR RELIANCE UPON SAID INFORMATION OR OPINION IS AT THE RISK OF THE USER AND THAT TELCORDIA SHALL NOT BE LIABLE FOR ANY DAMAGE OR INJURY INCURRED BY ANY PERSON ARISING OUT OF THE SUFFICIENCY, ACCURACY, OR UTILITY OF ANY INFORMATION OR OPINION CONTAINED HEREIN.

LOCAL CONDITIONS MAY GIVE RISE TO A NEED FOR ADDITIONAL PROFESSIONAL INVESTIGATIONS, MODIFICATIONS, OR SAFEGUARDS TO MEET SITE, EQUIPMENT, ENVIRONMENTAL SAFETY OR COMPANY-SPECIFIC REQUIREMENTS. IN NO EVENT IS THIS INFORMATION INTENDED TO REPLACE FEDERAL, STATE, LOCAL, OR OTHER APPLICABLE CODES, LAWS, OR REGULATIONS. SPECIFIC APPLICATIONS WILL CONTAIN VARIABLES UNKNOWN TO OR BEYOND THE CONTROL OF TELCORDIA. AS A RESULT, TELCORDIA CANNOT WARRANT THAT THE APPLICATION OF THIS INFORMATION WILL PRODUCE THE TECHNICAL RESULT OR SAFETY ORIGINALLY INTENDED.

This GR is not to be construed as a suggestion to anyone to modify or change any product or service, nor does this GR represent any commitment by anyone, including but not limited to Telcordia, to purchase, manufacture, or sell any product with the described characteristics.

Readers are specifically advised that any entity may have needs, specifications, or requirements different from the generic descriptions herein. Therefore, anyone wishing to know any entity's needs, specifications, or requirements should communicate directly with that entity.

Nothing contained herein shall be construed as conferring by implication, estoppel, or otherwise any license or right under any patent, whether or not the use of any information herein necessarily employs an invention of any existing or later issued patent.

TELCORDIA DOES NOT HEREBY RECOMMEND, APPROVE, CERTIFY, WARRANT, GUARANTEE, OR ENDORSE ANY PRODUCTS, PROCESSES, OR SERVICES, AND NOTHING CONTAINED HEREIN IS INTENDED OR SHOULD BE UNDERSTOOD AS ANY SUCH RECOMMENDATION, APPROVAL, CERTIFICATION, WARRANTY, GUARANTY, OR ENDORSEMENT TO ANYONE.

For general information about this or any other Telcordia documents, please contact:

Telcordia Customer Service
8 Corporate Place, Room 3B-176
Piscataway, NJ 08854-4156
1.866.672.6997 (USA)
+ 1.732.699.6700 (Worldwide)
+ 1.732.336.2226 (FAX)
e-mail: document-info@telcordia.com
<http://telecom-info.telcordia.com>

Table of Contents

Preface	Preface-1
1 Introduction	
1.1 Requirements Terminology	1-1
1.2 Requirement Labeling Conventions	1-2
1.2.1 Numbering of Requirement and Related Objects	1-2
1.2.2 Requirement, Conditional Requirement, and Objective Identification	1-2
1.3 Revision History	1-3
1.3.1 Changes from Issue 3 to Issue 4	1-3
1.3.2 Changes from Issue 2 to Issue 3	1-7
1.3.3 Requirement Object Absolute Number Assignment	1-10
2 Network Compatibility	
2.1 SONET NE Interfaces	2-2
2.1.1 Digital Signal Cross-Connect Interfaces	2-2
2.1.1.1 Electrical Cable Distance	2-3
2.1.1.2 Maintenance Signal Compatibility	2-3
2.1.2 Interface to Fiber Distributing Frames	2-3
2.1.3 Operations Systems Interface	2-4
2.1.4 Synchronization Interface	2-4
2.1.5 Power	2-4
2.2 End-to-End Performance Criteria	2-4
2.2.1 Availability and Reliability	2-5
2.2.2 Protection Switching Performance	2-5
2.2.3 Error Performance	2-6
2.2.4 Jitter	2-6
2.2.5 Transmission Delay	2-6
3 Rates and Formats	
3.1 Synchronous Hierarchical Rates	3-1
3.2 Transport Format	3-2
3.2.1 Frame Structure of the STS-1	3-3
3.2.1.1 Transport Overhead	3-4
3.2.1.2 STS-1 Envelope Capacity and Synchronous Payload Envelope (SPE)	3-5
3.2.2 Frame Structure of the STS-N	3-7
3.2.3 STS Concatenation	3-8
3.2.3.1 Contiguous STS Concatenation	3-8
3.2.3.2 STS Virtual Concatenation	3-11
3.2.4 Virtual Tributary Structure	3-14
3.2.5 VT Concatenation	3-27
3.3 Layered Overhead and Transport Functions	3-29
3.3.1 SONET Interface Layers	3-30
3.3.1.1 Physical Layer	3-30
3.3.1.2 Section Layer	3-30
3.3.1.3 Line Layer	3-31

- 3.3.1.4 Path Layers 3-31
- 3.3.1.5 Interaction of the Layers 3-31
- 3.3.2 STS-1 Overhead Descriptions 3-33
 - 3.3.2.1 Section Overhead 3-35
 - 3.3.2.2 FEC Sub-Layer 3-37
 - 3.3.2.3 Line Overhead 3-38
 - 3.3.2.4 STS Path Overhead 3-42
- 3.3.3 VT Path Overhead 3-49
- 3.4 Payload Mapping 3-55
 - 3.4.1 Sub-STs-1 Mappings 3-55
 - 3.4.1.1 Mappings into Specific Size VT SPEs 3-56
 - 3.4.1.1.1 Byte-Synchronous Mapping for DS1 3-57
 - 3.4.1.1.1.A DS1 Signaling and Framing Bit Transport 3-60
 - 3.4.1.1.1.B Line Codes and Pulse Density Assurance 3-63
 - 3.4.1.1.2 Asynchronous Mapping for DS1 3-63
 - 3.4.1.1.3 Asynchronous Mapping for DS1C 3-65
 - 3.4.1.1.4 Asynchronous Mapping for DS2 3-67
 - 3.4.1.2 Mappings That Apply to Multiple Sizes of VT SPEs 3-70
 - 3.4.1.2.1 ATM to VT SPE Mapping 3-70
 - 3.4.1.2.2 HDLC-Framed Signal to VT SPE Mapping 3-70
 - 3.4.1.2.3 GFP to VT SPE Mapping 3-71
 - 3.4.2 STS-1, STS-Nc, STS-1-Xv and STS-Nc-Xv Mappings 3-71
 - 3.4.2.1 Mappings into Specific Size STS SPEs 3-72
 - 3.4.2.1.1 Asynchronous Mapping for DS3 into an STS-1 SPE 3-72
 - 3.4.2.1.2 Asynchronous Mapping for DS4NA into an STS-3c SPE 3-73
 - 3.4.2.1.3 Asynchronous Mapping for FDDI into an STS-3c SPE 3-75
 - 3.4.2.1.4 DQDB Metropolitan Area Network (MAN) Mapping into an STS-3c SPE 3-76
 - 3.4.2.2 Mappings That Apply to Multiple Sizes of STS SPEs 3-78
 - 3.4.2.2.1 ATM to STS SPE Mapping 3-78
 - 3.4.2.2.2 HDLC-Framed Signal to STS SPE Mapping 3-79
 - 3.4.2.2.3 GFP to STS SPE Mapping 3-80
- 3.5 Payload Pointers 3-80
 - 3.5.1 STS Payload Pointer 3-80
 - 3.5.1.1 Pointer Value 3-81
 - 3.5.1.2 STS Frequency Justification 3-82
 - 3.5.1.3 New Data Flag (NDF) 3-85
 - 3.5.1.4 Concatenation Indicator 3-86
 - 3.5.1.5 STS Payload Pointer Generation Rules 3-86
 - 3.5.1.6 STS Payload Pointer Interpretation 3-88
 - 3.5.2 VT Payload Pointer 3-88
 - 3.5.2.1 VT Pointer Value 3-89
 - 3.5.2.2 VT Frequency Justification 3-90
 - 3.5.2.3 VT Size Indicator 3-91
 - 3.5.2.4 New Data Flag 3-91
 - 3.5.2.5 VT Payload Pointer Generation Rules 3-92
 - 3.5.2.6 VT Payload Pointer Interpretation Rules 3-93

4 Physical Layer

- 4.1 Physical Layer Classifications 4-1

4.2 Optical Parameter Definitions and Interface Requirements	4-5
4.2.1 General	4-5
4.2.2 Optical Line Coding	4-6
4.2.3 Transmitter	4-6
4.2.3.1 Spectral Characteristics	4-7
4.2.3.2 Coupled Transmit Power	4-9
4.2.3.3 Extinction Ratio	4-9
4.2.3.4 Mask of the Eye Diagram	4-9
4.2.4 Receiver	4-13
4.2.4.1 Receiver Sensitivity and Overload	4-13
4.2.4.2 Receiver Reflectance	4-15
4.2.4.3 Optical Power Penalty	4-15
4.2.5 Optical Path	4-17
4.2.5.1 Reflections and ORL	4-17
4.2.5.2 Attenuation	4-19
4.2.5.3 Chromatic Dispersion	4-20
4.2.5.3.1 Chromatic Dispersion for Most Applications at OC-1 through OC-48	4-21
4.2.5.3.2 Chromatic Dispersion for High Bit Rate/Long Range Applications	4-22
4.2.5.3.3 Chromatic Dispersion Accommodation	4-22
4.2.5.3.3.A Passive Dispersion Compensation	4-22
4.2.5.3.3.B Pre-Chirp	4-24
4.2.5.3.3.C Self Phase Modulation	4-24
4.2.5.4 Polarization Effects	4-25
4.2.5.4.1 Polarization Mode Dispersion	4-25
4.2.5.4.2 Polarization Dependent Loss	4-26
4.2.6 Optical Parameter Tables	4-26
4.3 Engineering of a Single-Mode Fiber Optic Transmission System	4-47
4.3.1 Terminal Equipment Transmission Design Information	4-47
4.3.1.1 General System Information	4-48
4.3.1.2 Transmitter Information	4-48
4.3.1.3 Receiver Information	4-49
4.3.1.3.1 General Information	4-49
4.3.1.3.2 Transmission Properties	4-50
4.3.1.4 Attenuators	4-51
4.3.1.5 Wavelength Division Multiplex Device	4-51
4.3.1.6 Passive Dispersion Compensation Devices	4-51
4.3.1.7 Optical Fiber Amplifiers	4-52
4.3.1.8 Connectors	4-52
4.3.1.9 Station Cable	4-53
4.3.1.10 Safety Margin	4-54
4.3.2 Cable Transmission Design Information	4-54
4.3.2.1 Parameters for a Specific Application	4-54
4.3.2.2 Global Fiber Parameters	4-57
4.3.3 Fiber Optic System Transmission Design and Analysis	4-59
4.3.3.1 Design Approach	4-59
4.3.3.2 Loss Budget Constraint	4-60
4.3.3.3 Dispersion Limited Length	4-61
4.3.3.4 Design and Analysis Methodology	4-62

4.4 Electrical Interface Specifications 4-63

5 Network Element Architectural Features

5.1 Multiplex Procedures 5-1

- 5.1.1 Interleaving 5-1
- 5.1.2 Concatenation 5-5
- 5.1.3 Scrambling 5-6
- 5.1.4 An Example of STS-1 and OC-N Signal Composition 5-7

5.2 Overhead Function Usage 5-10

- 5.2.1 Generating and Processing Overhead 5-10
- 5.2.2 Orderwire 5-14
 - 5.2.2.1 Orderwire Access 5-15
 - 5.2.2.2 Orderwire System Communication 5-15
 - 5.2.2.3 Orderwire Operations 5-16
- 5.2.3 User Channels 5-16
- 5.2.4 Remote Error Indications 5-17
- 5.2.5 Virtual Concatenation 5-18
 - 5.2.5.1 Capacity Adjustment 5-18
 - 5.2.5.2 Overhead Bit and Byte Usage 5-19
 - 5.2.5.3 Differential Delay Accommodation 5-24

5.3 Automatic Protection Switching 5-25

- 5.3.1 Protection Switching Boundaries 5-25
- 5.3.2 Linear APS Architectures 5-25
 - 5.3.2.1 1+1 Architecture 5-26
 - 5.3.2.2 1:n Architecture 5-27
 - 5.3.2.3 1:1 Case of the 1:n Architecture 5-28
- 5.3.3 Switch Initiation and Completion Criteria 5-28
 - 5.3.3.1 Switch Initiation Criteria 5-28
 - 5.3.3.2 Switch Initiation Time 5-29
 - 5.3.3.3 Switch Completion Time 5-31
- 5.3.4 Restoral and Clearing of SD and SF Conditions 5-32
- 5.3.5 APS Channel Protocol 5-35
 - 5.3.5.1 K1 Byte 5-36
 - 5.3.5.1.1 Bit Assignments for the K1 Byte 5-36
 - 5.3.5.1.2 K1 Byte Generation 5-38
 - 5.3.5.2 K2 Byte 5-41
 - 5.3.5.2.1 Bit Assignments for the K2 Byte 5-41
 - 5.3.5.2.2 K2 Byte Generation Rules 5-41
 - 5.3.5.2.3 Mode of Operation 5-42
 - 5.3.5.3 Control of the Bridge 5-43
 - 5.3.5.4 Control of the Selector 5-45
 - 5.3.5.5 Transmission and Acceptance of Bytes K1 and K2 5-46
- 5.3.6 Linear APS Commands 5-49
 - 5.3.6.1 Switch Commands 5-49
 - 5.3.6.2 Control Commands 5-50
- 5.3.7 Switch Operation 5-51
 - 5.3.7.1 1:n Architecture 5-51
 - 5.3.7.1.1 1:n Bidirectional Mode 5-51
 - 5.3.7.1.2 1:n Unidirectional Mode 5-56
 - 5.3.7.2 1+1 Architecture 5-56

5.3.7.2.1	1+1 Bidirectional Mode	5-56
5.3.7.2.2	1+1 Unidirectional Mode	5-56
5.4	Network Synchronization	5-57
5.4.1	SONET NE Clock Applications	5-57
5.4.1.1	Physical Interface to Synchronization Network	5-58
5.4.1.2	TDEV and MTIE Measurements	5-59
5.4.2	Synchronization Status Messages	5-59
5.4.3	SONET Timing Modes	5-63
5.4.3.1	External Timing	5-64
5.4.3.2	Line Timing	5-64
5.4.3.3	Loop Timing	5-66
5.4.3.4	Through Timing	5-66
5.4.4	SONET Internal Clock	5-67
5.4.4.1	Stratum Clocks for SONET Applications	5-67
5.4.4.2	SONET Minimum Clock Applications	5-68
5.4.4.2.1	Free-Run Accuracy for SMCs	5-68
5.4.4.2.2	Holdover Stability for SMCs	5-69
5.4.4.2.3	Pull-in/Hold-in for SMCs	5-69
5.4.4.2.4	Wander Tolerance and Transfer for SMCs and Stratum 3 Clocks	5-70
5.4.4.3	All SONET Stratum Clocks and SMCs	5-73
5.4.4.3.1	Clock Hardware	5-73
5.4.4.3.2	Wander Generation	5-74
5.4.4.3.3	Phase Transients	5-76
5.4.4.3.4	Reference Validation and Transitions From Self-Timing to Normal Mode	5-81
5.4.4.3.5	Jitter and Errors During Synchronization Rearrangement Operations	5-81
5.4.4.3.6	Input Tolerance	5-82
5.4.4.3.7	Phase Build-Out	5-83
5.4.5	Timing Distribution	5-83
5.4.5.1	Timing Distribution on Derived DS1 Signals	5-84
5.4.5.2	Synchronization Status Messages for Derived DS1 Signals	5-90
5.4.5.2.1	Switching	5-90
5.4.5.2.2	Message Translation	5-90
5.4.5.3	Timing Distribution on Traffic-Carrying DS1 Payload Signals	5-92
5.4.6	SONET Timing Reference Switching and Entry Into Holdover	5-93
5.4.6.1	Timing Reference Failure Conditions	5-95
5.4.6.2	Performance During Timing Reference Switching	5-96
5.4.6.3	Revertive and Nonrevertive Timing Reference Switching	5-97
5.4.6.4	Synchronization Status Messages and Timing Reference Switching	5-97
5.4.7	Synchronization Status Message Validation and Generation	5-98
5.4.7.1	Message Validation	5-98
5.4.7.2	Message Reaction	5-99
5.4.7.3	Message Generation	5-100
5.4.7.3.1	Externally Timed NEs	5-101
5.4.7.3.2	Line-Timed NEs	5-104
5.4.7.3.3	Through-Timed NEs	5-104
5.5	Framing For SONET Signals	5-105

- 5.6 Jitter 5-107
 - 5.6.1 Network Interface Jitter Criteria 5-108
 - 5.6.2 SONET NE Jitter Criteria 5-110
 - 5.6.2.1 Jitter Transfer 5-111
 - 5.6.2.1.1 Category I Jitter Transfer 5-112
 - 5.6.2.1.2 Category II Jitter Transfer 5-113
 - 5.6.2.2 Jitter Tolerance 5-114
 - 5.6.2.2.1 Category I Jitter Tolerance 5-114
 - 5.6.2.2.2 Category II Jitter Tolerance 5-115
 - 5.6.2.3 Jitter Generation 5-117
 - 5.6.2.3.1 Category I Mapping Jitter 5-118
 - 5.6.2.3.2 Category I Jitter Generation Due to Pointer Adjustments, General 5-119
 - 5.6.2.3.3 Category I Jitter Generation Due to Single Pointer Adjustments 5-121
 - 5.6.2.3.4 Category I Jitter Generation Due to Bursts of Pointer Adjustments 5-122
 - 5.6.2.3.5 Category I Jitter Generation Due to Periodic Pointer Adjustments 5-124
 - 5.6.2.3.6 Category II Jitter Generation 5-129
 - 5.6.2.3.7 Pointer Adjustment Jitter to Bit Stuffing Jitter Conversion 5-129
 - 5.6.2.4 Jitter Enhancement 5-131
- 5.7 Phase Variations on Payload Signals 5-131
 - 5.7.1 Mapping Phase Variations 5-131
 - 5.7.2 Pointer Adjustment Phase Variations 5-133
 - 5.7.2.1 Single Pointer Adjustments 5-134
 - 5.7.2.2 Pointer Adjustment Bursts 5-135
 - 5.7.2.3 Periodic Pointer Adjustments 5-137

6 SONET Network Element Operations Criteria

- 6.1 Memory Administration 6-1
 - 6.1.1 Memory Administration Data 6-1
 - 6.1.2 Data Manipulation 6-2
 - 6.1.3 Administration of Operations Communications Information 6-4
 - 6.1.4 Regenerators 6-4
 - 6.1.5 Memory Backup and Restoration 6-4
 - 6.1.6 System Administration and Security 6-6
 - 6.1.6.1 NE Security Mechanism 6-7
 - 6.1.6.1.1 Identification and Authentication 6-8
 - 6.1.6.1.2 System Access Control 6-9
 - 6.1.6.1.3 Resource Access Control 6-9
 - 6.1.6.1.4 Audit 6-10
 - 6.1.6.2 DCC Security 6-10
 - 6.1.7 Software Generics 6-12
 - 6.1.8 Self Inventory 6-12
- 6.2 Maintenance 6-13
 - 6.2.1 Alarm Surveillance 6-13
 - 6.2.1.1 Directly Detected Defects and Failures 6-16
 - 6.2.1.1.1 Loss of Signal 6-16
 - 6.2.1.1.2 Loss of Frame 6-18

6.2.1.1.3	Loss of Pointer	6-21
6.2.1.1.4	Equipment Failures	6-25
6.2.1.1.5	Loss of Synchronization	6-26
6.2.1.1.6	APS Troubles	6-27
6.2.1.1.6.A	Protection Switching Byte Failure	6-27
6.2.1.1.6.B	Channel Mismatch Failure	6-28
6.2.1.1.6.C	APS Mode Mismatch Failure	6-29
6.2.1.1.6.D	Invalid APS Mode Failure	6-30
6.2.1.1.6.E	Far-End Protection-Line Failure	6-31
6.2.1.1.6.F	Other APS Conditions	6-32
6.2.1.1.7	DCC Failure	6-33
6.2.1.1.8	Signal Label Mismatch	6-33
6.2.1.1.8.A	STS Payload Label Mismatch	6-34
6.2.1.1.8.B	STS Path Unequipped	6-34
6.2.1.1.8.C	VT Payload Label Mismatch	6-37
6.2.1.1.8.D	VT Path Unequipped	6-38
6.2.1.1.9	Trace Identifier Mismatch	6-40
6.2.1.1.9.A	STS Path Trace Identifier Mismatch	6-43
6.2.1.1.9.B	VT Path Trace Identifier Mismatch	6-45
6.2.1.2	Alarm Indication Signal	6-47
6.2.1.2.1	Line AIS	6-48
6.2.1.2.2	STS Path AIS	6-49
6.2.1.2.3	VT Path AIS	6-51
6.2.1.2.4	DSn AIS	6-53
6.2.1.3	Remote Defect Indication and Remote Failure Indication	6-55
6.2.1.3.1	Line Remote Defect Indication (RDI-L) and Remote Failure Indication (RFI-L)	6-55
6.2.1.3.2	STS Path Remote Defect Indication (RDI-P) and Remote Failure Indications (RFI-P)	6-57
6.2.1.3.3	VT Path Remote Defect Indication (RDI-V) and Remote Failure Indication (RFI-V)	6-60
6.2.1.3.4	DSn RDI and RAI Signals	6-65
6.2.1.4	Payload Defect Indication	6-67
6.2.1.4.1	STS Path Payload Defect Indication	6-68
6.2.1.4.2	VT Path Payload Defect Indication	6-70
6.2.1.5	Maintenance Signals for Other Mappings	6-70
6.2.1.6	Trunk Conditioning	6-70
6.2.1.7	Alarm-Related Events	6-72
6.2.1.8	Control of Alarm Processing	6-86
6.2.1.8.1	Alarm Level Designations	6-86
6.2.1.8.2	Single Failure/Single Message	6-87
6.2.1.8.3	Independent Failures	6-90
6.2.1.8.4	Retrieval of an NE's Condition	6-92
6.2.1.8.5	Provisioning of Alarm Levels	6-92
6.2.1.8.6	Clear Messages	6-93
6.2.1.8.7	Non-Intrusive Detection of Defects and Declaration of Failures	6-93
6.2.2	Performance Monitoring	6-94
6.2.2.1	General Accumulation and Thresholding Criteria	6-100
6.2.2.2	Physical Layer PM	6-107

6.2.2.2.1	Physical Layer Parameters	6-107
6.2.2.2.2	Physical Layer PM Criteria	6-108
6.2.2.3	Section Layer PM	6-110
6.2.2.3.1	Section Layer Parameters	6-110
6.2.2.3.2	Section Layer PM Criteria	6-111
6.2.2.4	Line Layer PM	6-112
6.2.2.4.1	Near-End Line Layer Parameters	6-112
6.2.2.4.2	Far-End Line Layer Parameters	6-114
6.2.2.4.3	Line Layer PM Criteria	6-115
6.2.2.5	STS Path Layer PM	6-115
6.2.2.5.1	Near-End STS Path Layer Parameters	6-115
6.2.2.5.2	Far-End STS Path Layer Parameters	6-117
6.2.2.5.3	STS Path Layer PM Criteria	6-118
6.2.2.6	VT Path Layer PM	6-119
6.2.2.6.1	Near-End VT Path Layer Parameters	6-120
6.2.2.6.2	Far-End VT Path Layer Parameters	6-121
6.2.2.6.3	VT Path Layer PM Criteria	6-122
6.2.2.7	Monitoring at DS _n Interfaces	6-123
6.2.2.8	PM During Troubles	6-124
6.2.2.9	Intermediate-Path PM	6-132
6.2.3	Testing Process	6-138
6.2.3.1	Test Access	6-139
6.2.3.1.1	Fiber Access	6-139
6.2.3.1.2	SONET Signal Test Access	6-139
6.2.3.1.3	Digital Test Access	6-140
6.2.3.2	Diagnostics	6-141
6.2.3.2.1	Physical Layer	6-142
6.2.3.2.2	Section Layer	6-142
6.2.3.2.3	Signal Identification	6-144
6.2.3.2.3.A	Section, STS Path and VT Path Trace	6-144
6.2.3.2.3.B	STS and VT Path Signal Label	6-148
6.2.3.2.4	Error Monitoring	6-148
6.2.3.2.5	Internal Clock	6-148
6.2.3.3	Loopbacks	6-149
6.2.3.3.1	SONET Terminal Loopback	6-149
6.2.3.3.2	SONET Cross-Connect Loopback	6-150
6.2.3.3.3	SONET Facility Loopback	6-151
6.2.3.3.4	DS _n Loopback	6-152
6.2.4	Control Features	6-153

7 Other Generic Criteria

7.1	Physical and Environmental Criteria	7-1
7.1.1	Operational Environment for Equipment	7-1
7.1.2	Electromagnetic Compatibility (EMC)	7-1
7.1.3	Outside Plant Cable	7-2
7.2	Equipment Design	7-2
7.3	Documentation and Training	7-2
7.4	Safety	7-3
7.4.1	Station Equipment Safety	7-3
7.4.2	Fiber Optic Cable Safety	7-4

7.5 Quality and Reliability 7-4
7.5.1 Network Equipment Reliability 7-4
7.5.2 Fiber Optic Cable Quality and Reliability 7-5
7.5.3 Component Reliability Assurance 7-5
7.6 Human Factors 7-6

8 SONET Operations Communications

8.1 SONET Operations Communications Architecture 8-2
8.1.1 Architecture Overview 8-2
8.1.2 Operations Communications Protocol Overview 8-4
8.1.3 Gateway NE Requirements 8-6
8.1.4 Intermediate NE Requirements 8-7
8.1.5 End NE Requirements 8-7
8.1.6 Mediation Device 8-8
8.2 Communications Types 8-8
8.2.1 OS/NE Communications 8-8
8.2.2 MD/NE Communications 8-9
8.2.3 NE/NE Communications 8-9
8.2.4 Craftsperson/NE Communications 8-10
8.3 SONET Operations Communications Interface 8-11
8.3.1 Physical Layer 8-12
8.3.1.1 OS/NE 8-12
8.3.1.2 NE/NE – LAN 8-12
8.3.1.3 NE/NE – DCC 8-13
8.3.2 Data Link Layer 8-14
8.3.2.1 OS/NE 8-14
8.3.2.2 NE/NE – LAN 8-14
8.3.2.3 NE/NE – DCC 8-14
8.3.3 Network Layer 8-15
8.3.3.1 OS/NE 8-16
8.3.3.2 NE/NE – LAN and DCC 8-16
8.3.4 Transport Layer – OS/NE and NE/NE 8-19
8.3.5 Session Layer – OS/NE and NE/NE 8-20
8.3.6 Presentation Layer – OS/NE and NE/NE 8-20
8.3.7 Application Layer – OS/NE and NE/NE 8-21
8.3.7.1 ACSE 8-21
8.3.7.2 ROSE/CMISE 8-21
8.3.7.3 FTAM or FTP 8-21
8.3.7.4 Name/Address Translation Services 8-22
8.3.7.5 TL1 8-22
8.4 Interworking between OSs and SONET NEs 8-23
8.4.1 TL1/X.25 [OS] – TL1/OSI [SONET] 8-24
8.4.1.1 Determine Destination NSAP 8-24
8.4.1.2 Directing Autonomous Messages 8-25
8.4.1.3 Establishing Connections 8-27
8.4.1.4 SONET LAN Interworking 8-27
8.4.1.5 SONET DCC Interworking 8-28
8.4.1.6 SONET LAN and DCC Interworking 8-28
8.4.2 TL1/X.25 [OS] – CMISE [SONET] 8-29
8.4.3 CMISE [OS] – CMISE [SONET] 8-29

- 8.4.4 CMISE or TL1/OSI [SONET] – CMISE or TL1/OSI [SONET] 8–30
- 8.5 SONET Operations Communications Routing 8–30
 - 8.5.1 Routing Overview 8–30
 - 8.5.2 ES-IS Requirements 8–32
 - 8.5.3 IS-IS Requirements 8–33
- 8.6 Craftsperson/NE Interfaces 8–34
 - 8.6.1 Craftsperson/WS Interface 8–34
 - 8.6.2 WS/NE Interface 8–35
- 8.7 TARP 8–35
 - 8.7.1 Network Layer Protocol to Support TARP 8–36
 - 8.7.2 TARP PDU Specification 8–36
 - 8.7.2.1 TARP Lifetime (tar-lif) 8–37
 - 8.7.2.2 TARP Sequence Number (tar-seq) 8–37
 - 8.7.2.3 Protocol Address Type (tar-pro) 8–37
 - 8.7.2.4 URC and TARP Type Code (tar-tcd) 8–37
 - 8.7.2.5 TID Target Length (tar-tln) 8–38
 - 8.7.2.6 TID Originator Length (tar-oln) 8–38
 - 8.7.2.7 Protocol Address Length (tar-pln) 8–38
 - 8.7.2.8 TID of Target (tar-ttg) 8–38
 - 8.7.2.9 TID of Originator (tar-tor) 8–38
 - 8.7.2.10 Protocol Address of Originator (tar-por) 8–39
 - 8.7.3 TARP Data Cache 8–39
 - 8.7.4 NE Applications That Use the TARP Processor 8–39
 - 8.7.4.1 Find NET That Matches TID 8–39
 - 8.7.4.2 Find TID That Matches NET 8–40
 - 8.7.4.3 Send Notification of TID or Protocol Address Change 8–40
 - 8.7.5 TARP PDU Processing 8–40
 - 8.7.5.1 Origination of a TARP Type 1 PDU 8–41
 - 8.7.5.2 Origination of a TARP Type 2 PDU 8–41
 - 8.7.5.3 Origination of a TARP Type 3 PDU 8–42
 - 8.7.5.4 Origination of a TARP Type 4 PDU 8–42
 - 8.7.5.5 Origination of a TARP Type 5 PDU 8–42
 - 8.7.5.6 Receipt of a TARP PDU 8–42
 - 8.7.5.6.1 End Systems 8–42
 - 8.7.5.6.2 Level 1 Intermediate Systems 8–43
 - 8.7.5.6.3 Level 2 Intermediate Systems 8–43
 - 8.7.5.7 Loop Detection Procedure (Performed by ISs) 8–43
 - 8.7.5.8 Propagation Procedure (Performed by ISs) 8–44
 - 8.7.6 Management of the TARP Processor 8–45
 - 8.7.7 TARP Echo Function 8–46
 - 8.7.8 Manual TARP Adjacencies 8–46
 - 8.7.9 TARP Example 8–47
 - 8.7.10 TARP Pseudocode 8–48

Appendix A: Deleted Requirement-Object List

- A.1 Requirement-Objects Deleted Through Issue 3 A–1
- A.2 Requirement-Objects Deleted as of Issue 4 A–4
 - A.2.1 Requirement-Objects Deleted From Section 3 A–4
 - A.2.2 Requirement-Objects Deleted From Section 4 A–4
 - A.2.3 Requirement-Objects Deleted From Section 5 A–5

A.2.4 Requirement-Objects Deleted From Section 6 A-6
A.2.5 Requirement-Objects Deleted From Section 8 A-6

Appendix B: Fiber Optic Transmission System Design Worksheets

Appendix C: SONET Operations Communications Lower Layers Protocol Profile

C.1 Introduction C-1
 C.1.1 Source Documents C-1
 C.1.1.1 Base Standards C-1
 C.1.1.2 International Standardized Profiles C-2
 C.1.1.3 American National Standards for Telecommunications C-2
 C.1.1.4 Telcordia Requirements C-2
 C.1.2 Notations Used in the SONET Lower Layer Profile C-2
 C.1.2.1 Status Symbols C-2
 C.1.2.2 Profile Symbols C-2
 C.1.2.3 Support Symbols C-4
 C.1.2.4 References C-4
C.2 LAPD SONET Protocol Profile C-4
 C.2.1 Notations C-4
 C.2.1.1 Abbreviations C-4
 C.2.1.2 Status Symbols C-5
 C.2.1.3 Additional Symbols C-5
 C.2.2 Protocol Capabilities (PC) C-6
 C.2.3 Frames - Protocol Data Units (FR) C-8
 C.2.4 System Parameters (SP) C-9
C.3 LLC Protocol Profile C-10
 C.3.1 Abbreviations and Special Symbols C-10
 C.3.1.1 Status Symbols C-10
 C.3.1.2 Item References C-10
 C.3.2 Claimed Conformance to ISO/IEC 8802-2:1989/Amd. 1, Amd. 2
 and Amd. 4 C-11
 C.3.3 LLC Type 1 Operation - Unacknowledged Connectionless Mode C-11
 C.3.3.1 LLC Type 1 - Supported PDU Types C-11
 C.3.3.2 LLC Type 1 - Supported Parameters in PDUs on Transmission C-12
 C.3.3.3 LLC Type 1 - Supported Parameters in PDUs on Receipt C-13
 C.3.3.4 LLC Type 1 - Miscellaneous C-13
C.4 ISO 8473 Protocol Profile C-15
 C.4.1 Notations C-15
 C.4.1.1 Status Symbols C-15
 C.4.1.2 Additional Symbols C-15
 C.4.2 Major Capabilities C-16
 C.4.3 End Systems C-16
 C.4.3.1 Applicability C-16
 C.4.3.2 Supported Functions C-17
 C.4.3.3 Supported PDUs C-18
 C.4.3.4 Supported Parameters C-19
 C.4.3.4.1 DT Parameters C-19
 C.4.3.4.2 ER Parameters C-20
 C.4.3.4.3 Inactive Network Layer Protocol PDU Parameters C-20

C.4.3.5	Timers	C-21
C.4.4	Intermediate Systems	C-21
C.4.4.1	Applicability	C-21
C.4.4.2	Supported Functions	C-22
C.4.4.3	Supported PDUs	C-23
C.4.4.4	Supported Parameters	C-24
C.4.4.4.1	DT Parameters	C-24
C.4.4.4.2	ER Parameters	C-25
C.4.4.5	Timer and Parameter Values	C-25
C.5	ISO 9542 SONET Protocol Profile	C-26
C.5.1	Notations	C-26
C.5.1.1	Status Symbols	C-26
C.5.1.2	Other Symbols	C-26
C.5.2	PICS Proforma: ISO 9542(1988) – End System	C-26
C.5.2.1	Supported Functions	C-27
C.5.2.2	Supported PDUs	C-27
C.5.2.3	Supported Parameters	C-28
C.5.2.4	Supported Parameter Ranges	C-29
C.5.3	PICS Proforma: ISO 9542(1988) – Intermediate System	C-29
C.5.3.1	Supported Functions	C-29
C.5.3.2	Supported PDUs	C-30
C.5.3.3	Supported Parameters	C-30
C.5.4	Supported Parameter Ranges	C-31
C.6	ISO/IEC 10589 Protocol Profile	C-31
C.6.1	Notations (Status Symbols)	C-31
C.6.2	Protocol Summary: ISO/IEC 10589 General	C-32
C.6.2.1	System Environment: General	C-33
C.6.2.2	Subnetwork Dependent Functions: General	C-33
C.6.2.3	Update Process: General	C-34
C.6.2.4	Decision Process: General	C-35
C.6.2.5	Forward/Receive Process: General	C-35
C.6.3	Protocol Summary: ISO/IEC 10589 Level 1 Specific Functions	C-36
C.6.3.1	Level 1 Subnetwork Dependent Functions	C-36
C.6.3.2	Level 1 Update Process	C-36
C.6.3.3	Level 1 Decision Process	C-36
C.6.4	Protocol Summary: ISO/IEC 10589 Level 2 Specific Functions	C-37
C.6.4.1	Level 2 Subnetwork Dependent Functions	C-38
C.6.4.2	Level 2 Update Process	C-38
C.6.4.3	Level 2 Decision Process	C-39
C.6.4.4	Level 2 Forward/Receive Process	C-39
C.7	ISO/IEC 8073 Protocol Profile	C-40
C.7.1	Notations	C-40
C.7.1.1	Status Symbols	C-40
C.7.2	Protocol Implementation for TP4/CLNS (C4L:.)	C-40
C.7.2.1	Annex B – NCMS	C-40
C.7.2.2	Classes Implemented	C-40
C.7.3	Initiator/Responder Capability for Protocol Classes 0-4	C-41
C.7.4	Supported Functions	C-41
C.7.4.1	Supported Functions for Class 4 (C4L:.)	C-41
C.7.5	Supported TPDUs	C-43

C.7.6 Supported Parameters of Issued TPDU	C-44
C.7.6.1 Parameter Values for CR TPDU (C4L::)	C-44
C.7.6.2 Supported Parameters for Class 4 TPDU (C4L::)	C-44
C.7.7 Supported Parameters for Received TPDU	C-46
C.7.8 User Data in Issued TPDU	C-46
C.7.8.1 Class 4 (C4L::)	C-47
C.7.9 User Data in Received TPDU	C-47
C.7.10 Negotiation	C-48
C.7.10.1 Class Negotiation - Initiator	C-48
C.7.10.2 TPDU Size Negotiation	C-48
C.7.10.3 Use of Extended Format	C-49
C.7.10.4 Expedited Data Transport Service	C-49
C.7.10.5 Non-Use of Checksum (C4L AND T4F29::)	C-49
C.7.10.6 Use of Selective Acknowledgment (See note 20)	C-50
C.7.10.7 Use of Request of Acknowledgment (See note 21)	C-50
C.7.11 Error Handling	C-50
C.7.11.1 Action on Receipt of a Protocol Error	C-50
C.7.11.2 Actions on Receipt of an Invalid or Undefined Parameter in a CR TPDU	C-51
C.7.11.3 Actions on Receipt of an Invalid or Undefined Parameter in a TPDU other than a CR TPDU	C-52
C.7.12 Timers and Protocol Parameters	C-52
C.8 TARP Protocol Implementation Conformance Statement	C-53
C.8.1 Major Function	C-53
C.8.2 Supported PDUs	C-54
C.8.3 Protocol Specifications	C-54
C.8.3.1 TARP PDU CLNP Specifications	C-54
C.8.3.2 TARP PDU Specifications	C-55
C.8.3.3 Protocol Timer Specifications	C-55
C.8.4 Major Capabilities	C-56
C.8.5 TARP Processor Management	C-57
C.8.5.1 LDB Entry Timer Parameters	C-57
C.8.5.2 LDB Flush Timer Parameters	C-58
C.8.5.3 Provisionable TARP PDU Fields	C-58

Appendix D: SONET Operations Communications Upper Layers Protocol Profile

D.1 Introduction	D-1
D.2 Source Documents	D-1
D.2.1 Base Standards	D-1
D.2.2 PICS Proforma	D-1
D.2.3 International Standardized Profiles	D-2
D.2.4 Telcordia Requirements	D-2
D.3 Goals of SONET Upper Layers Profile	D-2
D.4 Structure of SONET Upper Layers Profile	D-2
D.5 Notations Used in the SONET Upper Layers Profile	D-3
D.5.1 Abbreviations	D-3
D.5.2 Status Column	D-3
D.5.3 Profile Column	D-3
D.5.4 Support Column	D-4

D.5.5 PICS Numbers	D-4
D.6 SONET Upper Layers Profile: ACSE	D-4
D.6.1 Additions Beyond Existing ISP Requirements	D-4
D.6.2 Profile Tables	D-5
D.7 SONET Upper Layers Profile: Presentation Layer	D-9
D.7.1 Additions Beyond Existing ISP Requirements	D-9
D.7.2 Profile Tables	D-10
D.8 SONET Upper Layers Profile: Session Layer	D-16
D.8.1 Additions Beyond Existing ISP Requirements	D-16
D.8.2 Profile Tables	D-17

References

Glossary

Requirement-Object Index

List of Figures

Figure 2-1	Simplified Diagram Depicting SONET Section, Line and Path Definitions	2-1
Figure 2-2	Diagram Illustrating SONET Section, Line and Path Definitions	2-2
Figure 3-1	STS-1 Frame	3-4
Figure 3-2	Bit Position Numbering	3-4
Figure 3-3	STS-1 Synchronous Payload Envelope	3-6
Figure 3-4	STS-1 SPE with STS-1 POH and STS-1 Payload Capacity Illustrated	3-6
Figure 3-5	STS-1 SPE in Interior of STS-1 Frames	3-7
Figure 3-6	STS-N Frame	3-8
Figure 3-7	STS-Nc SPE	3-10
Figure 3-8	Transport Overhead Assignment, OC-3 Carrying an STS-3c SPE	3-10
Figure 3-9	Mapping of the STS-1-Xv Payload Capacity Into X STS-1 SPEs	3-12
Figure 3-10	Mapping of the STS-3c-Xv Payload Capacity Into X STS-3c SPEs	3-13
Figure 3-11	VT Sizes	3-15
Figure 3-12	VT Structured STS-1 SPE: All VT1.5s	3-16
Figure 3-13	VT1.5 Locations	3-17
Figure 3-14	VT Structured STS-1 SPE: All VT2s	3-18
Figure 3-15	VT2 Locations	3-19
Figure 3-16	VT Structured STS-1 SPE: All VT3s	3-20
Figure 3-17	VT3 Locations	3-21
Figure 3-18	VT Structured STS-1 SPE: All VT6s	3-22
Figure 3-19	VT6 Locations	3-23
Figure 3-20	Example of VT Structured STS-1 SPE	3-24
Figure 3-21	Correspondence Between Labels and Numbers for the Example in Figure 3-20	3-25
Figure 3-22	VT Superframe and Envelope Capacity	3-26
Figure 3-23	VT SPE and Payload Capacity	3-27
Figure 3-24	Mapping of the VTn-Xv Payload Capacity Into X VTn SPEs	3-29
Figure 3-25	Optical Interface Layers	3-33
Figure 3-26	Transport and STS Path Overhead Byte Designations	3-34
Figure 3-27	STS Path Status Byte (G1)	3-48
Figure 3-28	VT Path Overhead Byte (V5)	3-49
Figure 3-29	Z7 Bit 1 Multiframe Definition	3-53
Figure 3-30	H4 Byte Coding Sequence for VT-Structured STS-1 SPEs	3-56
Figure 3-31	Byte-Synchronous Mapping for DS1 Payload	3-58
Figure 3-32	Byte-Synchronous DS1 Signaling and Framing Bit Assignments	3-59
Figure 3-33	Asynchronous Mapping for DS1 Payload	3-65
Figure 3-34	Asynchronous Mapping for DS1C Payload	3-67
Figure 3-35	Asynchronous Mapping for DS2 Payload	3-69
Figure 3-36	Asynchronous Mapping for DS3 Payload	3-73
Figure 3-37	Asynchronous Mapping for DS4NA Payload	3-74
Figure 3-38	Asynchronous Mapping for FDDI	3-76
Figure 3-39	DQDB Mapping into an STS-3c SPE	3-78
Figure 3-40	Bit Allocation for H4 Byte in DQDB Mapping	3-78
Figure 3-41	STS Payload Pointer (H1, H2) Coding	3-81
Figure 3-42	STS Pointer Offset Numbering	3-82
Figure 3-43	Positive STS Pointer Adjustment Operation (Increment)	3-84

Figure 3-44	Negative STS Pointer Adjustment Operation (Decrement)	3–85
Figure 3-45	VT Payload Pointer (V1, V2) Coding	3–89
Figure 3-46	VT Pointer Offsets	3–90
Figure 4-1	Optical System Interfaces (Points S and R)	4–5
Figure 4-2	SONET Eye Diagram Mask (OC-1 through OC-12)	4–10
Figure 4-3	SONET Eye Diagram Mask (OC-48 through OC-768)	4–11
Figure 4-4	Cable Configuration for Cabled Fiber Cutoff Wavelength Measurement	4–55
Figure 4-5	Wavelength Dependent Attenuation Characteristics	4–56
Figure 4-6	Example Global Fiber Attenuation Characteristics	4–58
Figure 4-7	Example Global Fiber Dispersion Characteristics	4–59
Figure 4-8	Engineering of a Fiber Optic Regenerator Section	4–60
Figure 4-9	Regenerator Section, Sources of Loss	4–61
Figure 4-10	STS-1 Electrical Interface Pulse Mask	4–66
Figure 4-11	STS-1 Electrical Interface Eye Diagram Mask	4–67
Figure 4-12	STS-3 Transmitter Pulse Mask Corresponding to a Binary Zero	4–69
Figure 4-13	STS-3 Transmitter Pulse Mask Corresponding to a Binary One	4–70
Figure 4-14	STS-3 Eye Diagram Mask	4–71
Figure 5-1	Example of Byte-Interleaving Sequence, STS-12	5–3
Figure 5-2	Byte-Interleaving Example, Multiple Level Inputs	5–5
Figure 5-3	Frame Synchronous Scrambler (Functional Diagram)	5–7
Figure 5-4	Example of STS-1 and OC-N (N < 768) Signal Composition	5–9
Figure 5-5	Switch Initiation Time Criteria	5–31
Figure 5-6	Linear APS Switch – 1:n Architecture (in released position)	5–44
Figure 5-7	Linear APS Switch – 1+1 Architecture (in released position)	5–45
Figure 5-8	1:n Linear APS Architecture Example	5–52
Figure 5-9	OC-N and OC-M Example	5–57
Figure 5-10	External-Timing Mode Example	5–64
Figure 5-11	Line-Timing Mode Example	5–65
Figure 5-12	Loop-Timing Mode	5–66
Figure 5-13	Through-Timing Mode Example	5–67
Figure 5-14	OC-N Input and Output Wander Time Deviation	5–72
Figure 5-15	Time Deviation of Filtered Network Input to SONET NEs	5–73
Figure 5-16	MTIE for SONET Clocks	5–75
Figure 5-17	Time Deviation for SONET Clocks	5–75
Figure 5-18	MTIE for Phase Transients from SONET Clocks	5–78
Figure 5-19	Phase-Transient for Stratum 3 Clock and SMC Entry into Holdover	5–79
Figure 5-20	MTIE Masks for Input/Output Phase Transients	5–80
Figure 5-21	DS1 Timing References Derived from a Single OC-N “Interface” Example	5–84
Figure 5-22	DS1 Timing References Derived from Different OC-N “Interfaces” Example	5–85
Figure 5-23	Time Deviation for Derived DS1 Signals	5–89
Figure 5-24	Example Implementation of R5-234 and R5-235	5–104
Figure 5-25	Example Implementation of R5-241 and R5-242	5–104
Figure 5-26	Example Implementation of R5-243	5–105
Figure 5-27	Category I DS1 and DS3 Jitter Transfer Mask	5–113
Figure 5-28	Category II Jitter Transfer Mask	5–114
Figure 5-29	SONET Category II Jitter Tolerance Mask	5–116

Figure 5-30	Single Pointer Adjustment Test Sequence	5-122
Figure 5-31	Maximum Rate Pointer Burst Test Sequence	5-123
Figure 5-32	Phase Transient Pointer Burst Test Sequence	5-123
Figure 5-33	Periodic VT1.5 Pointer Adjustment Test Sequence (26-1 Pattern)	5-126
Figure 5-34	Periodic STS Pointer Adjustment Test Sequence (87-3 Pattern)	5-127
Figure 5-35	Periodic Pointer Adjustment Test Sequence (Continuous Pattern)	5-128
Figure 5-36	DS1 Mapping Phase Variation Limits	5-132
Figure 5-37	DS3 Mapping Phase Variation Limits	5-133
Figure 5-38	Single VT Pointer Adjustment Phase Variation Limits	5-134
Figure 5-39	Single STS-1 Pointer Adjustment Phase Variation Limits	5-135
Figure 5-40	Maximum Rate Pointer Burst Phase Variation Limits	5-136
Figure 5-41	Phase Transient Pointer Burst Phase Variation Limits	5-137
Figure 5-42	Periodic VT Pointer Adjustment Phase Variation Limits	5-138
Figure 5-43	Periodic STS-1 Pointer Adjustment Phase Variation Limits	5-139
Figure 6-1	General Defect Detection and Failure Declaration Model	6-15
Figure 6-2	Maintenance Signals for SONET Layers	6-48
Figure 6-3	STE Maintenance Signals	6-72
Figure 6-4	LTE Maintenance Signals	6-73
Figure 6-5	STS PTE (Asynchronous Mapping for DS3) Maintenance Signals	6-74
Figure 6-6	STS PTE (VT-Structured STS-1 SPE) Maintenance Signals	6-75
Figure 6-7	VT PTE (Asynchronous Mapping for DS _n into VT _n) Maintenance Signals, DS _n Interface	6-76
Figure 6-8	VT PTE (Asynchronous Mapping for DS1 into VT1.5) Maintenance Signals, DS3 Interface with Embedded M13 Multiplex	6-77
Figure 6-9	VT PTE (Byte-Synchronous Mapping for DS1 into a Single VT1.5) Maintenance Signals, DS1 Interface	6-78
Figure 6-10	VT PTE with DS0 Rearrangement Functions (Byte-Synchronous Mapping for DS1 into VT1.5) Maintenance Signals and Trunk Conditioning, OC-N and DS1 Interfaces	6-79
Figure 6-11	VT PTE with DS0 Rearrangement Functions (Byte-Synchronous Mapping for DS1 into VT1.5) Maintenance Signals and Trunk Conditioning, OC-N and DS1 Interfaces (Continued)	6-80
Figure 6-12	DS0 PTE (Byte-Synchronous Mapping for DS1 into VT1.5) Maintenance Signals and Trunk Conditioning	6-81
Figure 6-13	SONET Maintenance Signals for Internal Equipment Failures	6-82
Figure 6-14	Alarm Timing Requirements for Directly Detected Defects and Failures	6-83
Figure 6-15	AIS Timing Requirements	6-84
Figure 6-16	Derived RFI Timing Requirements	6-85
Figure 6-17	RFI-V Signal Timing Requirements (for Byte-Synchronous DS1 Mapping)	6-85
Figure 6-18	SONET PM Accumulation and Thresholding Model	6-97
Figure 6-19	SONET PM Accumulation and Thresholding Model (Continued)	6-98
Figure 6-20	SONET PM Accumulation and Thresholding Model (Continued)	6-99
Figure 6-21	Intermediate-Path PM for Non-PJ-Related Parameters	6-133
Figure 6-22	Examples of Intermediate-Path PM for STS PJ-Related Parameters	6-134
Figure 6-23	Section Layer Diagnostic	6-143
Figure 6-24	SONET Terminal Loopback Example	6-150

Figure 6-25	SONET Cross-Connect Loopback	6-150
Figure 6-26	SONET Facility Loopback	6-151
Figure 6-27	DSn Terminal Loopback	6-152
Figure 6-28	DSn Facility Loopback	6-153
Figure 8-1	SONET Operations Communications: Example NE and Interface Types	8-1
Figure 8-2	Example SONET Operations Communications Architecture	8-2
Figure 8-3	Example Intra-site LAN and Point-to-Point DCC	8-3
Figure 8-4	Example Operations Communications Network for a Survivable Ring	8-3
Figure 8-5	Operations Communications Functions	8-4
Figure 8-6	Interactive CMISE/OSI Protocol Stacks for SONET Operations Communications	8-5
Figure 8-7	SONET NSAP Format	8-18
Figure 8-8	IDP Encoding	8-18
Figure 8-9	TL1/X.25 – LAN Interworking	8-28
Figure 8-10	TL1/X.25-LAN-DCC Interworking	8-29
Figure 8-11	Example Routing Domain	8-32
Figure 8-12	Craftsperson/NE Communications Network	8-34
Figure 8-13	Manual TARP Adjacencies	8-47
Figure 8-14	TARP Example	8-47

List of Tables

Table 3-1	Line Rates for Standard SONET Interface Signals (through N = 768)	3-2
Table 3-2	STS Path Signal Label Assignments	3-45
Table 3-3	STS Path Signal Label Assignments for Signals with Payload Defects	3-46
Table 3-4	VT Path Signal Label Assignments (Non-Extended)	3-52
Table 3-5	Extended VT Path Signal Label Assignments	3-54
Table 3-6	VTn and VTn-Xv SPE Attributes	3-55
Table 3-7	STS SPE Attributes	3-71
Table 3-8	VT Size Indicator	3-91
Table 4-1	Application Categories	4-3
Table 4-2	Attenuation and Group Delay Distortion as a Function of Frequency	4-12
Table 4-3	Tolerance Values of the Attenuation of the Optical Reference Receiver	4-13
Table 4-4	Maximum Attenuation and Range Values	4-20
Table 4-5	Maxwellian DGD Probabilities	4-26
Table 4-6	SR-1 OC-1 Through OC-48 Optical Parameters	4-27
Table 4-7	SR OC-192 Optical Parameters	4-28
Table 4-8	SR OC-768 Optical Parameters	4-29
Table 4-9	IR OC-1 Optical Parameters	4-30
Table 4-10	IR OC-3 Optical Parameters	4-31
Table 4-11	IR OC-12 Optical Parameters	4-32
Table 4-12	IR OC-48 Optical Parameters	4-33
Table 4-13	IR OC-192 Optical Parameters	4-34
Table 4-14	IR OC-768 Optical Parameters	4-35
Table 4-15	LR OC-1 Optical Parameters	4-36
Table 4-16	LR OC-3 Optical Parameters	4-37
Table 4-17	LR OC-12 Optical Parameters	4-38
Table 4-18	LR OC-48 Optical Parameters	4-39
Table 4-19	LR OC-192 Optical Parameters	4-40
Table 4-20	LR OC-768 Optical Parameters	4-41
Table 4-21	VR OC-12 Optical Parameters	4-42
Table 4-22	VR OC-48 Optical Parameters	4-43
Table 4-23	VR OC-192 Optical Parameters	4-44
Table 4-24	UR OC-12 Optical Parameters	4-45
Table 4-25	UR OC-48 Optical Parameters	4-46
Table 4-26	STS-1 Electrical Interface Criteria	4-65
Table 4-27	STS-3 Electrical Interface Criteria	4-68
Table 5-1	Possible Starting Positions for an STS-Mc SPE in an OC-768 Signal .	5-4
Table 5-2	SONET Overhead Generating and Processing Criteria	5-12
Table 5-3	H4 Byte Usage in STS Virtual Concatenation Applications	5-21
Table 5-4	Z7 Bits 1 and 2 Usage in VT Virtual Concatenation	5-23
Table 5-5	Clearing Time Criteria for BER-based SF and SD Conditions	5-33
Table 5-6	K1 Byte, Bits 1 through 4: Type of Request	5-37
Table 5-7	Channel Number Code Assignments, K1 Bits 5 to 8 (and K2 Bits 1 to 4)	5-38
Table 5-8	1:n Bidirectional Switching Example	5-52

Table 5-9	Synchronization Status Message Definitions	5-60
Table 5-10	Example Reference Selection Using Synchronization Status Messages	5-98
Table 5-11	Parameters for Network Interface Jitter Requirements	5-110
Table 5-12	Category I Mapping Jitter Limits	5-119
Table 5-13	Pointer Test Sequence Parameters	5-120
Table 5-14	Jitter Due to Single Pointer Adjustments	5-122
Table 5-15	Jitter Generation Limits for Periodic Pointer Adjustment Sequences	5-125
Table 5-16	SONET Jitter Generation	5-129
Table 6-1	STS-1 Numbers in OC-N Signals	6-3
Table 6-2	STS Signal Label Mismatch Defect Conditions	6-36
Table 6-3	VT Signal Label Mismatch Defect Conditions	6-39
Table 6-4	RDI-P Bit Settings and Interpretation	6-60
Table 6-5	RDI-V Bit Settings and Interpretation	6-64
Table 6-6	Hierarchy of Near-End Failures	6-89
Table 6-7	Hierarchy of Far-End Failures	6-90
Table 6-8	PM Register Sizes and Default Thresholds	6-104
Table 6-9	Section BIP Errors to Trigger a Section SES	6-111
Table 6-10	Line BIP Errors to Trigger a Line SES (Near-End and Far-End) . .	6-113
Table 6-11	STS Path BIP Errors to Trigger an STS Path SES (Near-End and Far-End)	6-116
Table 6-12	VT Path BIP Errors to Trigger a VT Path SES (Near-End and Far-End)	6-120
Table 6-13	Section Layer PM Accumulation During Defects	6-126
Table 6-14	Near-End Line Layer PM Accumulation During Defects	6-127
Table 6-15	Near-End STS Path Layer PM Accumulation During Defects	6-128
Table 6-16	Near-End VT Path Layer PM Accumulation During Defects	6-129
Table 6-17	Far-End Line Layer PM Accumulation During Defects	6-130
Table 6-18	Far-End STS Path Layer PM Accumulation During Defects	6-131
Table 6-19	Far-End VT Path Layer PM Accumulation During Defects	6-132
Table 8-1	TARP PDU Fields	8-37
Table 8-2	TARP Types	8-38
Table 8-3	TARP Timers	8-40
Table 8-4	Provisionable TARP PDU Fields	8-45

Preface

This Preface contains important information about the Telcordia GR process in general, as well as important information about this GR.

The Telcordia Technologies GR Process

Generic Requirements documents (GRs) provide the Telcordia view of proposed generic criteria for telecommunications equipment, systems or services, and involve a wide variety of factors, including interoperability, network integrity, participant expressed needs and other input.

The Telcordia GR process implements Telecommunications Act of 1996 directives relative to the development of industry-wide generic requirements relating to telecommunications equipment, including integral software and customer premises equipment. Pursuant to that Act, Telcordia invites members of the industry to participate in the development process for such GRs. Invitations to participate are issued monthly in the *Telcordia Digest of Technical Information*, and posted on the Telcordia web site at: <http://www.telcordia.com/resources/genericreq/digest>.

At the conclusion of the GR development process, Telcordia publishes the GR, which is available by for license. The license fee entitles the licensee to receive that issue of the GR (GR-CORE) along with any Issues List Report (GR-ILR) and Revisions, if any are released under that GR project. ILRs contain any technical issues that arise during GR development that Telcordia would like further industry interaction on. The ILR may present issues for discussion, with or without proposed resolutions, and may describe proposed resolutions that lead to changes to the GR. Significant changes or additional material may be released as a Revision to the GR-CORE.

Telcordia may also solicit general industry nonproprietary input regarding such GR material at the time of its publication, or through a special Industry Interaction Notice appearing in the *Telcordia Digest of Technical Information*. While unsolicited comments are welcome, any subsequent work by Telcordia regarding such comments will depend on participation in such GR work. Telcordia will acknowledge receipt of comments and will provide a status to the submitting company.

About GR-253-CORE

Participants in GR-253-CORE Issue 4 Development

Telcordia Technologies

Relative Maturity Level

This is a mature technology and the requirements reflect a maintenance mode. Throughout the current document (i.e., GR-253-CORE, Issue 4), the criteria are considered stable except as indicated in the text.

GR-253 Plans

This GR will be periodically evaluated for possible reissue beginning in 2007, as will the need for a replacement for GR-253-ILR, Issue 3A (i.e., an Issue 4A).

To Submit Comments

When submitting comments, please include the GR number, and cite any pertinent section and requirement number. In responding to an ILR, please identify the pertinent Issue ID number. Please provide the name and address of the contact person in your company for further discussion.

Comments should be submitted by January 1, 2007.

Send comments to:

Telcordia — GR-253-CORE
Tom Bowmaster
One Telcordia Drive
Piscataway, NJ 08854-4182
Phone: +1.732.699.5489
E-Mail: tbowmast@telcordia.com

1 Introduction

GR-253-CORE, Issue 4 contains the Telcordia view of Synchronous Optical Network (SONET) generic criteria. In general, GR-253-CORE, Issue 4 expands on, and in some cases changes the criteria contained in GR-253-CORE, *Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria*, Issue 3.

The criteria contained herein are intended to advise the telecommunications industry of the Telcordia view of proposed generic requirements. The criteria reflect recent changes in requirements advocated by the various committees and subcommittees of standards organizations such as the International Telecommunication Union – Telecommunication Standardization Sector (ITU-T), American National Standards Institute (ANSI) accredited Committee T1 and the Electronic Industries Association/Telecommunications Industries Association (EIA/TIA), and agreements reached in the SONET Interoperability Forum (SIF) and the Network and Services Integration Forum (NSIF).

1.1 Requirements Terminology

The following requirements terminology is used in GRs:

- Requirement – Feature or function that, in the view of Telcordia, is *necessary* to satisfy the needs of a typical service provider. Failure to meet a requirement may cause application restrictions, result in improper functioning of the product or hinder operations. A Requirement contains the words *shall* or *must* and is flagged by the letter “**R**.”
- Conditional Requirement – Feature or function that, in the view of Telcordia, is *necessary in specific applications*. If a service provider identifies a Conditional Requirement as necessary, it shall be treated as a requirement for the application(s). Conditions that may cause the Conditional Requirement to apply include, but are not limited to, certain service provider application environments, elements or other requirements, etc. A Conditional Requirement is flagged by the letters “**CR**.”
- Objective – Feature or function that, in the view of Telcordia, is *desirable* and may be required by a service provider. An Objective represents a goal to be achieved. An Objective may be reclassified as a Requirement at a specified date. An objective is flagged by the letter “**O**” and includes the words *should*, *it is desirable* or *it is an objective*.
- Conditional Objective – Feature or function that, in the view of Telcordia, is *desirable in specific applications* and may be required by a service provider. It represents a goal to be achieved in the specified Condition(s). If a service provider identifies a Conditional Objective as necessary, it shall be treated as a requirement for the application(s). A Conditional Objective is flagged by the letters “**CO**.”
- Condition – The circumstances that, in the view of Telcordia, will cause a Conditional Requirement or Conditional Objective to apply. A Condition is flagged by the letters “**Cn**.”

This GR helps establish a foundation for interoperability between different implementations of the functions described herein. It is important to note, however, that a number of optional features are included, both in this GR and the standards documents that are referenced. Criteria in this GR relating to such features are denoted as conditional requirements. In selecting a consistent set of optional features, the following situations have to be taken into account:

- The domain of a single network provider when using a single supplier
- The domain of a single network provider when using multiple suppliers
- Interoperability between domains of different network providers.

1.2 Requirement Labeling Conventions

As part of the Telcordia GR Process, proposed requirements and objectives are labeled using the conventions explained in the following two sections.

1.2.1 Numbering of Requirement and Related Objects

Each Requirement, Objective, Condition, Conditional Requirement, and Conditional Objective object is identified by both a local and an absolute number. The local number consists of the object's document section number and its sequence number in the section (e.g., Requirement **R3-1** is the first Requirement object in Section 3). The local number appears in the margin to the left of the Requirement. A Requirement object's local number may change in subsequent issues of a document if other Requirements are added to the section or deleted.

The absolute number is a permanently assigned number that will remain for the life of the Requirement; it will not change with new issues of the document. The absolute number is presented in brackets (e.g., **[2]**) at the beginning of the requirement text.

Neither the local nor the absolute number of a Conditional Requirement or Conditional Objective depends on the number of the related Condition(s). If there is any ambiguity about which Conditions apply, the specific Condition(s) will be referred to by number in the text of the Conditional Requirement or Conditional Objective.

References to Requirements, Objectives, or Conditions published in other GRs will include both the document number and the Requirement object's absolute number. For example, **R2345-[12]** refers to Requirement **[12]** in GR-2345-CORE.

1.2.2 Requirement, Conditional Requirement, and Objective Identification

A Requirement object may have numerous elements (paragraphs, lists, tables, equations, etc.). To aid the reader in identifying each part of the requirement, rules are used above and below the requirement content:

Introductory information.

Content of Requirement object(s).

Explanatory information.

1.3 Revision History

This section gives a high-level view of the major changes in GR-253-CORE, Issue 4 with respect to GR-253-CORE, Issue 3 and, for historical purposes, the major changes in GR-253-CORE, Issue 3 with respect to GR-253-CORE, Issue 2 (including Revisions 1 and 2). Many of these changes were related to issues that were discussed in GR-253-ILR, Issues 2C and 3A, or were based on industry comments and agreements reached in standards bodies after the previous issue or revision of the document was released. In addition, changes were made in a number of sections to further clarify or better organize the criteria. In general, all of the substantive changes that were made in Issue 4 (with respect to Issue 3) have been marked with change bars in the outside margins of the document.

1.3.1 Changes from Issue 3 to Issue 4

As itemized below, GR-253-CORE, Issue 4 contains technical changes from GR-253-CORE, Issue 3 in most sections of the document.

- Section 2, *Network Compatibility*, changes include:
 - Revisions in [Section 2.1.1.1](#) to reflect the maximum cable distance criteria that were added to Issue 3 of GR-499-CORE
 - The addition of references to the standards that contain transmission delay specifications applicable to SONET NEs that support the FEC function (see [Section 2.2.5](#)).
- Section 3, *Rates and Formats*, changes include:
 - Additions and revisions in multiple subsections to reflect the addition of OC-768 applications and the VT and STS path virtual concatenation functions to this document
 - The revision of criteria and text to reflect that contiguous concatenation of STS-1s (into an STS-Nc) is now defined only for N = 3, 12, 48, 192 and 768 (see [Section 3.2.3.1](#))
 - The addition of criteria and text related to the need for a SONET NE to maintain the bit integrity of the signals that it supports (see [Sections 3.3.1.5](#) and [3.4.1.1.1.A](#))
 - The addition of criteria and text related to the possible use of the J0 byte for a 16-byte Section trace function (see [Section 3.3.2.1](#))
 - Expansion of the information provided (in multiple subsections) that is related to the FEC function defined in ANSI T1.105.08 and ITU-T Rec. G.707

- Changing the name of the byte used for the REI-L function in signals at the OC-1 rate from M0 to M1, and the addition of information related to the interpretation of “invalid” REI-L codes (see [Section 3.3.2.3](#))
- The addition of criteria and text related to STS and VT path supervisory unequipped signals and the signal labels associated with several new payload mappings (see [Sections 3.3.2.4](#) and [3.3.3](#))
- The addition of information related to the generation of REI-P and REI-V codes in the presence of various types of defects detected on the incoming path signal (see [Sections 3.3.2.4](#) and [3.3.3](#))
- Definition of the extended VT signal label and VT path trace functions (see [Section 3.3.3](#))
- The addition of criteria and text related to the patterns an NE that supports the byte-synchronous DS1 mapping and DS0 rearrangement capabilities should insert in unassigned DS0 data bytes (see [Section 3.4.1.1.1](#))
- The definition of mappings for ATM, HDLC-framed, and GFP payloads into VT SPEs (see [Section 3.4.1.2](#))
- The definition of a mapping for GFP payloads into STS SPEs (see [Section 3.4.2.2.3](#)).
- Section 4, *Physical Layer*, changes include:
 - Additions and revisions in multiple subsections to reflect the addition of OC-768 applications to this document
 - Removal of the former Section 4.2.3, *Reflections*, and reassignment of the criteria and text from that section to the appropriate subsections in the current [Sections 4.2.3](#), *Transmitter*, [4.2.4](#), *Receiver*, and [4.2.5](#) *Optical Path*
 - Clarification of conditions under which an OC-N transmitter is expected to generate a signal that meets the applicable criteria, and reorganization of that criteria for consistency with other parts of the document (see [Section 4.2.3](#))
 - Clarification of the intent of the transmitter criteria in cases where two possible central wavelength ranges and corresponding maximum spectral width and D_{SRmax} values are given (see [Section 4.2.3.1](#))
 - Removal of Electrical and Optical Signal to Noise Ratio (ESNR and OSNR) as parameters to be used in the specification of OC-192 transmitters
 - Clarification of conditions under which an OC-N receiver is expected to meet its sensitivity requirement, and the addition of information related to the margins that can be expected in various situations (see [Section 4.2.4.1](#))
 - Reorganization and expansion of the criteria and text related to optical power penalties (see [Sections 4.2.4.3](#) and [4.2.5.1](#))
 - The addition of target distances for each application category listed in [Table 4-4](#)
 - Various clarifications and additions to [Section 4.2.5.3](#), *Chromatic Dispersion*, including the definition of a maximum chromatic dispersion deviation parameter

- The addition, to the optical parameter tables in [Section 4.2.6](#), of information regarding equivalent standards codes and the (few) discrepancies between the specifications in the standards and the criteria in this document
- The addition of new OC-12 VR and UR, OC-48 VR and UR, and OC-192 LR-2 and VR-2b application categories, the removal of the OC-192 VR-1 category, and minor changes to the parameter values for several other categories in order to correct errors or align them with the corresponding standards (see [Section 4.2.6](#)).
- Section 5, *Network Element Architectural Features*, changes include:
 - Additions and revisions in multiple subsections to reflect the addition of OC-768 applications, and the VT and STS path virtual concatenation functions
 - The revision of criteria and text to reflect that contiguous concatenation of STS-1s (into an STS-Nc) is now defined only for N = 3, 12, 48, 192 and 768 (see [Section 5.1.1](#))
 - The addition of criteria related to the generation of REI codes (see [Section 5.2.4](#))
 - Revision of the criteria related to the relative priorities of Forced Switch and SF requests at NEs supporting linear APS (see [Table 5-6](#))
 - The revision of criteria and text related to the mode of operation to be used by an NE the supports linear APS and is not receiving the expected codes from the far-end NE (see [Section 5.3.5.2](#))
 - The revision of criteria and text related to the transmission of invalid codes by an NE the supports linear APS (see [Section 5.3.5.5](#))
 - The addition of criteria to allow an NE that supports the synchronization status messages defined in this document to interoperate with NEs that support an earlier version of those messages (see [Section 5.4.2](#))
 - The revision of criteria and text in multiple subsections of [Sections 5.4.4](#), [5.4.5](#) through [5.4.6](#) to reflect changes that were made in Issue 3 of GR-1244-CORE and align this document with the current versions of various related standards
 - Full definition of the functionality of the Forced Reference Switch command (see [Section 5.4.6](#))
 - Reorganization and various modifications of the criteria and text related to synchronization status messages (see [Section 5.4.7](#))
 - The revision of criteria and text in multiple subsections of [Section 5.6](#) to reflect changes that were made in Issue 3 of GR-499-CORE, better align this document with the current versions of various related standards, and provide explanations for the (few) differences between the specifications in the standards and the criteria in this document.

- Section 6, *SONET Network Element Operations Criteria*, changes include:
 - Additions and revisions in several subsections to reflect the addition of OC-768 applications and the definition of TIM-V as a defect that may be detected by a SONET NE
 - The addition of criteria and text covering cases where equipment that is supposed to receive an STS-Nc path signal instead receives some combination of STS-1 and STS-Mc path signals, or equipment that is supposed to receive some combination of STS-1 and STS-Mc path signals instead receives an STS-Nc path signal (see [Section 6.2.1.1.3](#))
 - Clarification of the criteria in [Section 6.2.1.1.5](#) related to Loss of Synchronization failure declaration and clearing
 - Definition of “Invalid APS Mode” defects and failures to provide indications of provisioning and connection problems that can occur in linear APS systems (see [Section 6.2.1.1.6.D](#))
 - Modification of the PLM-V defect detection and termination criteria to reflect the possibility that the incoming and/or received signal label may be an extended VT signal label carried in the Z7 byte (see [Section 6.2.1.1.8.C](#))
 - The expansion and reorganization of the criteria and text related to TIM defects and failures to account for the possibility of STS path trace messages that meet the specifications in ANSI T1.269, and the use of VT path trace messages for TIM-V detection purposes (see [Section 6.2.1.1.9](#))
 - The addition of information related to the setting of an invalid data flag when an NE’s time of day clock is changed (see [Section 6.2.2.1](#))
 - The addition of criteria and text related to a failure-based TCA-suppression feature similar to that described in ANSI T1.231 (see [Section 6.2.2.1](#)).
 - The addition of criteria and text related to possible nominal values of the OPR Physical layer PM parameter (see [Sections 6.2.2.2.1](#) and [6.2.2.2.2](#))
 - The value of “K” used in the definition of the OC-192 SES-S PM parameter was corrected (see [Table 6-9](#))
 - The revision of criteria and text to indicate that the accumulation of STS and VT PJ-related PM parameters is not supposed to be inhibited as a result of entry into unavailable time for the corresponding path signal (see [Section 6.2.2.8](#))
 - Expansion of [Section 6.2.3.2.3.A](#) on STS path trace diagnostics to also cover Section and VT path trace diagnostics
 - The addition of a reference to the internal clock diagnostics criteria in GR-1244-CORE (see [Section 6.2.3.2.5](#))
 - The definition of a SONET cross-connect loopback function (see [Section 6.2.3.3.2](#))
- Section 8, *SONET Operations Communications*, changes include:
 - Numerous editorial changes (which have generally not been marked with change bars)

- Additions and revisions in multiple subsections to reflect the possible use of the Line or extended Line DCC, FTP, and protocol stacks that include TCP, IP and Ethernet
- Additions and revisions to reflect the possible long-term use of TL1 as an application-layer protocol
- Clarifications regarding the particular implementations in which various criteria apply, and which implementations are not currently addressed via detailed criteria.
- The placeholder that was inserted into Issue 3 when the contents of the previous Appendix A were removed has been replaced by a new appendix entitled “*Deleted Requirement-Object List*”.
- Appendix B, *Fiber Optic Transmission System Design Worksheets*, has been updated to reflect the addition of criteria related to OC-768 systems to the document.
- Several definitions and the complete acronym list in the *Glossary* section were updated, and a number of definitions that were redundant with those provided in the remainder of the document were deleted.
- The *References* section was updated.
- The *Requirement-Objects Index* section was updated.

1.3.2 Changes from Issue 2 to Issue 3

As itemized below, GR-253-CORE, Issue 3 contained technical changes from GR-253-CORE, Issue 2 (including Revisions 1 and 2) in essentially all sections of the document.

- Section 2, *Network Compatibility*, changes included:
 - The addition of text and criteria related to digital signal cross-connect interfaces for systems that support OC-192 signals (see [Section 2.1.1](#))
 - Revisions to update and clarify a number of areas, including cable lengths ([Section 2.1.1.1](#)) and synchronization ([Section 2.1.4](#)).
- Section 3, *Rates and Formats*, changes included:
 - Revisions in multiple sub-sections to account for the addition of OC-192 to the scope of the document
 - The addition of references to recent or current standards work in such areas as virtual concatenation ([Section 3.2.3](#)), forward error correction ([Section 3.3.2.2](#)), an extended REI-L function ([Section 3.3.2.3](#)) and extended VT path signal labels ([Section 3.3.3](#)).
 - The removal of text and criteria related to ZBTSI (see [Section 3.4.1.1.1.B](#))
 - The addition of text and criteria related to ATM and HDLC-over-SONET mappings into STS-48c SPEs ([Section 3.4.3.3](#) in Issue 3, now a part of [Section 3.4.2.2](#)).

- Section 4, *Physical Layer*, changes included:
 - Revisions in multiple sub-sections to account for the addition of OC-192 to the scope of the document (e.g., the addition of information and criteria related to the Very-Long Reach application category, booster amplifiers and pre-amplifiers, source frequency chirp factors, modulation methods, dispersion compensation methods, non-linear fiber effects, polarization mode dispersion)
 - The addition of text and criteria related to the possibility of damage to optical receivers due to high input power levels (see GR-253-ILR Issue ID 253-40 and [Section 4.2.4](#)).
 - Modification of a number of the optical parameter values that previously appeared in GR-1377-CORE or Issue 2 of this document to align them with the current ITU-T standards or draft standards (see GR-253-ILR Issue ID 253-44 and [Section 4.2.6](#)).
- Section 5, *Network Element Architectural Features*, changes included:
 - Revisions in multiple sub-sections to account for the addition of OC-192 to the scope of the document
 - The addition of text related to the possible use of the linear APS Clear command to terminate a WTR period (see [Section 5.3.6.1](#))
 - The addition of information related to SONET versus non-SONET synchronization distribution networks (see [Sections 5.4.1](#) and [5.4.4.2.4](#))
 - The removal of criteria conditionally requiring the support of stratum 3E and 2 clocks by SONET NEs (see [Section 5.4.4.1](#))
 - The addition of an explicit wander tolerance requirement for SMC and stratum 3 clocks (see GR-253-ILR, Issue ID 253-136, and [Section 5.4.4.2.4](#))
 - Clarification of the relationship between the jitter defined to be present on the wander generation input test signal and the level of jitter required to be tolerated by a SONET NE on an OC-N timing reference signal (see [Section 5.4.4.3.2](#))
 - Clarification of the relationship between the input phase transient defined for use in (output) phase transient tests and the type of phase transient required to be tolerated by a SONET NE on an OC-N timing reference signal (see [Section 5.4.4.3.3](#))
 - The addition of criteria concerning manual switch commands and revertive/nonrevertive switching for timing-related switches between working and protection lines (see GR-253-ILR, Issue IDs 253-110 and 253-114, and [Sections 5.4.5.1](#), [5.4.6](#) and [5.4.6.3](#))
 - Solidification of the criteria related to hold-off times for reference switches based on synchronization status messages and the removal of the DUS message (see GR-253-ILR Issue ID 253-60, and [Sections 5.4.6.4](#) and [5.4.7.3.1](#))
 - Clarification of the requirements related to the generation of synchronization status messages upon recovery from holdover (see GR-253-ILR Issue ID 253-138, and [Section 5.4.7.2](#))

- The addition of criteria and text related to false framing (see GR-253-ILR, Issue ID 253-63 and [Section 5.5](#))
- Revision of the text and criteria related to “phase-smoothing” to make it clear that those criteria also apply to the case where a DS_n payload is demultiplexed from a SONET SPE and multiplexed into a higher bit-rate DS_n signal within a single NE (see [Sections 5.6.2.1.1](#) and [5.6.2.3.7](#))
- The addition of an objective for SONET NEs to tolerate higher levels of low-frequency jitter at their OC-12, OC-48 and OC-192 interfaces (see [Sections 5.6.2.2.2](#))
- Section 6, *SONET Network Element Operations Criteria*, changes included:
 - Revisions in several sub-sections to account for the addition of OC-192 to the scope of the document
 - The addition of references to GR-2915-CORE for criteria related to the management of bulk memory restoration processes (see [Section 6.1.5](#))
 - The addition of information related to the emerging Public Key Infrastructure technology and references to NSIF-038-2000 for requirements related to centralized security servers (see GR-253-ILR Issue IDs 253-35, 253-36 and 253-130, and [Sections 6.1.6](#) and [6.1.6.1.1](#))
 - The addition of text and criteria related to the ability of an NE to report adjacency information (see GR-253-ILR Issue ID 253-131, and [Section 6.1.8](#))
 - Revision of the text and criteria related to LOS defect detection and termination, and LOS and LOF failure declaration to address several GR-253-ILR issues and make this document more compatible with the corresponding ANSI SONET standards (see GR-253-ILR Issue ID 253-68 and 253-69, and [Sections 6.2.1.1.1](#), [6.2.1.1.2](#) and [6.2.1.8.2](#))
 - Revisions to make it clear that the criteria in Section 8 of GR-1244-CORE related to alarms, reports and control commands apply to SONET NEs (see [Section 6.2.1.1.5](#))
 - The addition of references to ANSI T1.269 concerning the format of STS path trace messages (see [Sections 6.2.1.1.9](#) and [6.2.3.2.3.A](#))
 - The addition of text related to the possible impact of TIM-P defect detection on VT-accessed STS-1 paths in a BLSR (see [Section 6.2.1.1.9](#))
 - Clarification of text and criteria related to the support and generation of ERDI signals, and the detection of ERDI defects (see GR-253-ILR Issue IDs 253-120, 253-127 and 253-140, and [Sections 6.2.1.3.2](#), [6.2.1.3.3](#), [6.2.2.5.1](#) and [6.2.2.6.1](#))
 - Revision of the text and criteria related to the LBC PM parameter and diagnostic to indicate that they may not be appropriate for use with some optical transmitter technologies (see [Sections 6.2.2.2.1](#), [6.2.2.2.2](#) and [6.2.3.2.1](#))
 - The addition of text and criteria to account for the possibility that an NE that supports Intermediate STS path PM may also be able to be provisioned to monitor for an expected STS path trace message on that path (see [Section 6.2.2.9](#))

- The addition of text and criteria related to DS1 test access (see GR-253-ILR Issue ID 253-128, and [Section 6.2.3.1.3](#))
- Section 7, *Other Generic Criteria*, changes included:
 - Revision of much of the text and criteria related to operational environments for SONET NEs so that they accurately reflect the information and criteria that currently appears in other documents (see [Sections 7.1.1](#) and [7.1.2](#))
 - The addition of a reference to several of the documentation requirements in Section 9 of GR-1244-CORE to make it clear that they apply to SONET NEs (see [Section 7.3](#))
 - Revision of the sub-sections related to quality and reliability to account for the addition of OC-192 to the scope of the document (see [Section 7.5](#)).
- Section 8, *SONET Operations Communications*, changes included:
 - Revisions throughout the section to update references to relevant standards and Telcordia GRs
 - The addition of criteria to support TCP/IP-based WANs for OS/NE communications and FTP for file-oriented applications (see GR-253-ILR Issue ID 253-142, and [Sections 8.2.1](#) and [8.3](#)).
- Appendix A, *Requirement-Object List*, was removed and a placeholder inserted so that the letters assigned to the remaining appendices would not need to be changed.
- Appendix B, *Fiber Optic Transmission System Design Worksheets*, was updated to reflect the addition of criteria related to OC-192 systems to the document.
- The standards and profile lists in Appendix C, *SONET Operations Communications Lower Layers Protocol Profile*, and Appendix D, *SONET Operations Communications Upper Layers Protocol Profile*, were updated.
- The Acronym list in the *Glossary* section was updated.
- The *References* section was updated.
- The *Requirement-Objects Index* section was updated.

1.3.3 Requirement Object Absolute Number Assignment

In general, the absolute number assigned to any particular requirement object can be used (along with the following list) to determine the Issue or Revision of this document in which that object first appeared.

- [1] to [903] appeared in Issue 1
- [904] to [1012] first appeared in Issue 2
- [1013] to [1043] first appeared in Issue 2, Revision 1
- [1044] to [1099] first appeared in Issue 2, Revision 2
- [1100] to [1138] first appeared in Issue 3
- [1139] to [1234] first appeared in the current issue (Issue 4).

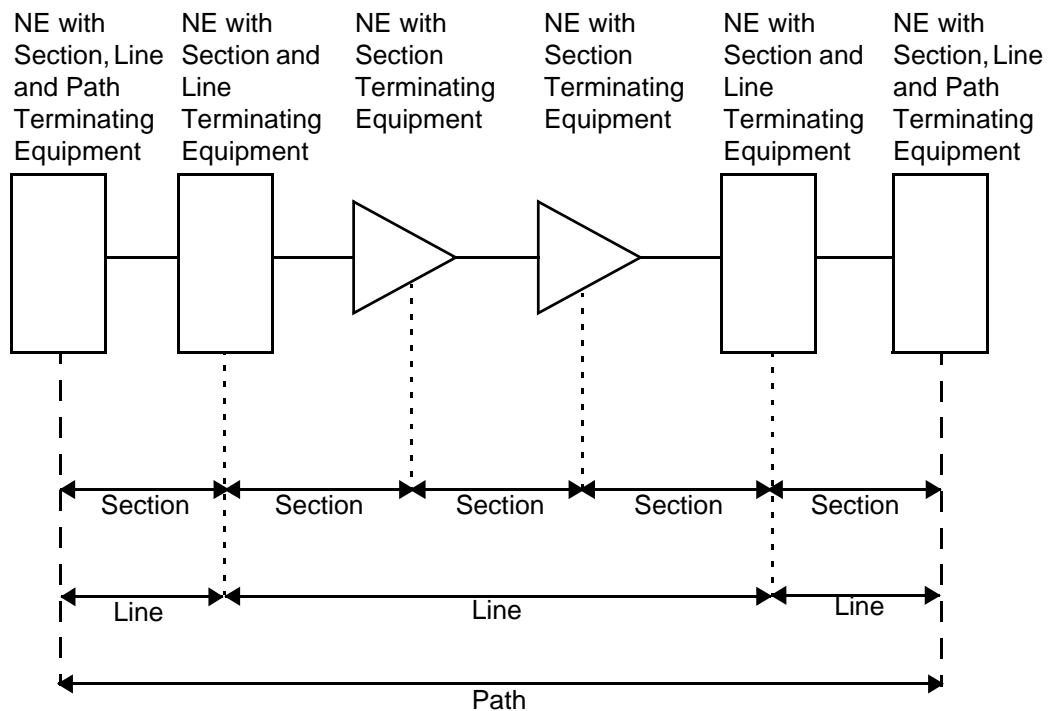
In addition, version numbers (e.g., **[96v2]**, **[418v3]**) are used to mark criteria that have undergone substantive changes in one or more issues of the document. Finally, the absolute numbers assigned to criteria that appeared in a previous issue, but that have now been removed, are not reused.

2 Network Compatibility

This section presents criteria that are intended to help ensure SONET equipment compatibility with the existing network. These criteria deal with interfaces to the existing network and system performance, and are, for the most part, specific to SONET. Generic criteria applicable to SONET and other transmission systems (e.g., asynchronous fiber optic systems or digital radio systems) are found in GR-499-CORE, *Transport Systems Generic Requirements (TSGR): Common Requirements*.

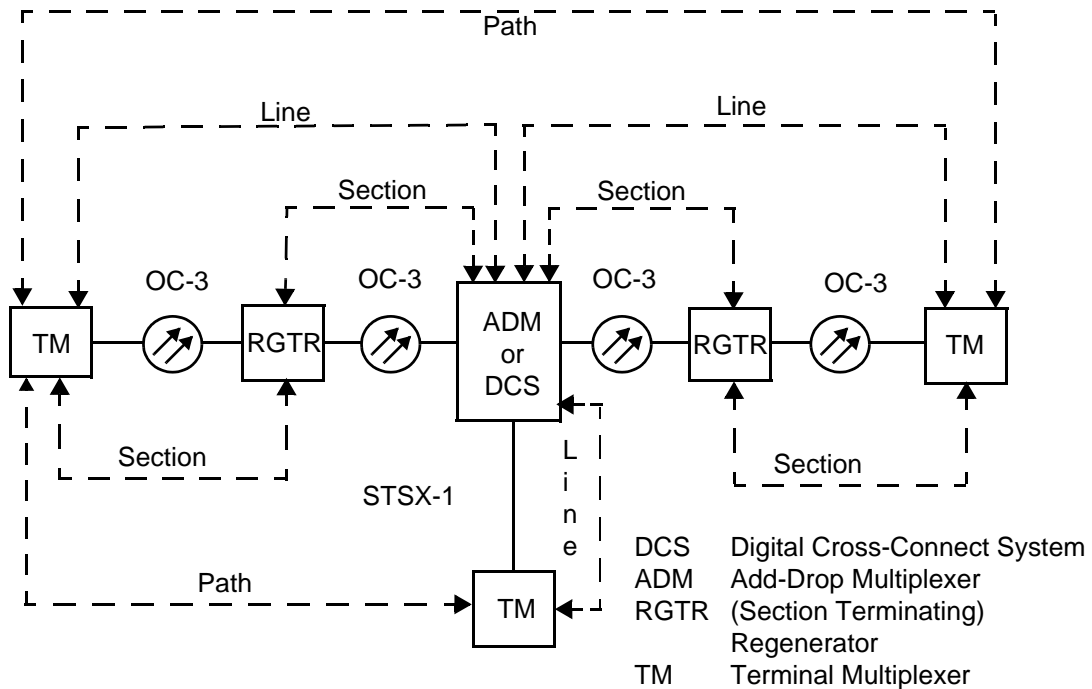
This document uses the terms Section, Line and Path¹ to delineate various segments of the transmission network that interconnect the SONET Network Elements (NEs). Figures 2-1 and 2-2 illustrate these transmission segments between various SONET NEs defined in the Glossary. Specific requirements for SONET NEs such as the Add/Drop Multiplexer (ADM), Terminal Multiplexer (TM), Digital Cross-connect System (DCS) and Regenerator (RGTR), are contained in referenced material.

Figure 2-1 Simplified Diagram Depicting SONET Section, Line and Path Definitions



1. Note that for payloads carried using Virtual Tributaries (VTs, see Section 3.2.4), the path layer discussed in this section actually consists of the Synchronous Transport Signal level 1 (STS-1) path layer and the VT path layer.

Figure 2-2 Diagram Illustrating SONET Section, Line and Path Definitions



2.1 SONET NE Interfaces

SONET NEs interface with existing telecommunications equipment such as other transmission systems, outside-plant fiber, operations systems and power systems. This section discusses these interfaces.

2.1.1 Digital Signal Cross-Connect Interfaces

This document considers SONET equipment that interfaces with other transmission equipment at (among other possibilities) electrical signal cross-connects. Interface requirements for interconnecting at the standard hierarchical Digital Signal Cross-connects (DSX-1, DSX-1C, DSX-2, DSX-3 and DSX-4NA) are found in GR-499-CORE, and are based on ANSI T1.102, *Digital Hierarchy – Electrical Interfaces*, and ANSI T1.107, *Digital Hierarchy – Formats Specifications*. SONET NEs may also provide SONET electrical interfaces, and interconnect with other SONET NEs at SONET STS electrical Cross-connects (i.e., STSX-1s and STSX-3s) as described in [Section 4.4](#).

In most cases, there are no criteria that specify that a SONET NE should be able to support particular types or numbers of interfaces. That is generally considered to be an issue for the SONET NE supplier to determine, based on such factors as the type of NE (e.g., ADM, DCS) and the intended applications. On the other hand, for SONET NEs that support both OC-192 interfaces and DS3/STS-1 electrical interfaces, the following requirement applies.

-
- R2-1 [1100]** A SONET NE that supports one or more OC-192 interfaces, and that also supports DS3 and/or STS-1 electrical interfaces, shall be capable of supporting a minimum of 48 such (electrical) interfaces.
-

In addition, the following criteria apply to all SONET NEs that support DS_n and/or STS-N electrical interfaces.

-
- CR2-2 [1]** SONET NEs may be required to provide electrical signal cross-connect facilities with patching and monitoring jacks for restoration and rearrangements.
- CR2-3 [2]** A bridging repeater or equivalent may be required to allow rearrangement from the monitor jack.
-

2.1.1.1 Electrical Cable Distance

For any DS1, DS1C, DS2, DS3 or DS4NA interfaces provided by a SONET NE, the interface requirements in Section 9 of GR-499-CORE apply. Those criteria define the acceptable characteristics of a signal appearing at the interface, and also the maximum (reference) cable distances between the transmitter and the interface and between the interface and the receiver. In most cases, the interface is expected to be an electrical signal cross-connect frame, and the maximum cable distances are a function of the interface signal's hierarchical level.

2.1.1.2 Maintenance Signal Compatibility

To prevent unwanted propagation of alarms and to help sectionalize failures in a manner consistent with the existing networks, SONET NEs provide and interact with maintenance signals at the standard hierarchical rates. See [Section 6.2.1](#).

2.1.2 Interface to Fiber Distributing Frames

Fiber distributing frames (or other types of optical cross-connect systems) allow optical fibers to be connected and disconnected as necessary without the need to assemble and disassemble splices. Information on fiber distributing frames is contained in GR-449-CORE, *Generic Requirements and Design Considerations for Fiber Distributing Frames*.

-
- R2-4 [3v2]** Suppliers shall provide a description of their fiber distributing frames. This description shall include the number of terminations, storage provisions for excess fiber, and the type of fiber cable connectors provided (see [Section 4](#)).
-

2.1.3 Operations Systems Interface

The protocols, languages and content of the generic networking requirements for interfaces with network providers' operations systems are described in [Section 8](#). The procuring network provider determines the interface arrangement.

2.1.4 Synchronization Interface

In general, SONET uses the synchronization network as described in GR-436-CORE, *Digital Network Synchronization Plan* and ANSI T1.101, *Synchronization Interface Standard*. The goal is to create a fully synchronous optical hierarchy by ensuring that each SONET NE receives timing traceable to a Primary Reference Source (PRS). A PRS is equipment that provides a timing signal whose long-term accuracy is maintained at $\pm 1 \times 10^{-11}$ or better with verification to Universal Time Coordinated (UTC), and whose timing signal is used as the basis of reference for the control of other clocks within a network. Note that while the timing signals used by each SONET NE in a network should be traceable to a PRS, all of the NEs in the network do not need to use signals traceable to the same PRS. That is, a SONET network may use more than one PRS.

The clocks used to provide external timing signals to SONET NEs are stratum 3 (or better quality) clocks as described in ANSI T1.101. Thus, the timing for SONET signals is normally traceable to a primary reference source as GR-436-CORE and ANSI T1.101 describe.

Requirements for the short-term stability of SONET signals are described in [Section 5.4](#), which provides detailed criteria on the SONET NE synchronization interface, timing and clocks.

2.1.5 Power

In most applications, SONET NEs are expected to receive power in the form of a dc voltage with a nominal value of $-48 V_{dc}$. Requirements on SONET equipment pertaining to voltage limits, electrical noise and current drain are contained in GR-499-CORE.

2.2 End-to-End Performance Criteria

This section addresses issues related to SONET system performance from the standpoint of compatibility with the existing network. These criteria ensure that a SONET transport system performs as well as other systems that transport the standard asynchronous hierarchical signals (e.g., DS1, DS3).

The subjects covered in the following subsections include availability (i.e., maximum downtime), error performance, protection switching performance, jitter and transmission delay.

2.2.1 Availability and Reliability

Service availability and system reliability are critical issues for the service providers as well as their customers. Equipment suppliers should consider reliability issues during all phases of system design and manufacture. In terms of design goals, the equipment supplier needs to understand the impact of individual equipment reliability on end-to-end service availability.

The underlying foundation for many system reliability criteria is an end-to-end two-way availability objective of 99.98% for interoffice applications (0.02% unavailability or 105 minutes/year maximum downtime). The corresponding objective for loop transport between the central office and the customer's premises is 99.99%. In interoffice transport, the objective refers to a two-way broadband channel (e.g., DS3, STS-N or OC-N) over a 250-mile path; in loop applications, the objective applies to a two-way narrowband channel (e.g., DS0 or equivalent). In either case, the objective is meant as a long-term average over a large area for regular (non-special) service offerings, such as "plain old telephone service" for loop customers.²

Starting with these end-to-end objectives, a top-down approach allocates maximum downtime to various parts of a generic model network, described in terms of a "hypothetical reference circuit". SONET systems present a challenging problem because a service provider can "mix and match" equipment from different suppliers. To ensure that the total downtime does not exceed the end-to-end objective, this situation for SONET forces allocations down to the level of individual NEs. Special consideration is needed for rings and for metropolitan applications.

A comprehensive discussion of SONET availability, including specific criteria and associated rationale, is presented in GR-418-CORE, *Generic Reliability Assurance Requirements for Fiber Optic Transport Systems*. Specific availability criteria for SONET NEs are provided in Section 2 of that GR.

2.2.2 Protection Switching Performance

Automatic Protection Switching (APS) systems increase system availability by automatically substituting a protection equipment (e.g., an optical transmitter, path and receiver) for failed equipment. A digital channel bank or other digital NE produces false signaling states and eventually a Carrier Group Alarm (CGA) after receiving a signal with sufficiently high Bit Error Ratio (BER) or incoming signal failure for a specified time. Thus, a protection switching system must recognize both situations and respond in a sufficiently short time. Section 5 of GR-499-CORE contains generic criteria on protection switching for fiber optic systems (at both the line and circuit pack levels).

The physical performance of line protection switching is characterized by the time to detect certain switching triggers and thresholds (based on BER) and the time to physically complete the switch. [Section 5.3](#) of this document presents criteria specific to SONET protection switching.

2. Different levels of performance could be requested by a customer or assured by a service provider, based on matters outside of the general discussion here.

2.2.3 Error Performance

Error performance criteria are concerned with such parameters as BER and Errored Seconds observed at standard hierarchical interface rates (e.g., DS1, DS3). Section 4 of GR-499-CORE gives definitions and requirements for these performance parameters.

2.2.4 Jitter

Timing jitter may arise from a number of sources within the digital network. Timing jitter is defined as the short-term variations of the significant instants of a digital signal from their ideal positions in time, where short-term implies phase oscillations of frequencies greater than 10 Hz. [Section 5.6](#) contains the jitter criteria applicable to SONET NEs and systems.

2.2.5 Transmission Delay

Guidelines for transmission delay in telecommunications networks are given in ANSI T1.506-1997 (R2001), *Network Performance – Switched Exchange Access Network Transmission Specifications* and ANSI T1.508-1998, *Loss Plan for Evolving Digital Networks*. In addition, specific SONET NEs have specific transmission delay criteria as described in the appropriate NE-specific GRs and Technical Reference documents (TRs), and [Section 3.3.2.2](#) of this document references ANSI T1.105.08, *Synchronous Optical Network (SONET) – In-band Forward Error Correction Code Specification*, and ITU-T Rec. G.707, *Network node interface for the synchronous digital hierarchy (SDH)*, both of which contain supplemental transmission delay requirements for NEs that support the Forward Error Correction (FEC) function. When incorporating SONET NEs into a transmission path, the processing and propagation delays contributed by the NEs and the propagation delay contributed by the interconnecting media must be taken into consideration in the context of the guidelines in the ANSI standards referenced above.

3 Rates and Formats

This section defines the rates and formats for SONET signals. A primary goal in defining these signals is to articulate a synchronous hierarchy that has sufficient flexibility to carry many different capacity signals. This is realized by defining a basic module with a bit rate of 51.840 Mb/s, and a byte-interleaved multiplex scheme that results in a family of signals with rates of N times 51.840 Mb/s, where N is an integer (Section 3.1).

The basic module can be divided into a portion assigned to overhead and a portion that carries the payload (Section 3.2). This payload portion can be used to transport DS3 signals or a variety of sub-DS3 signals. Because some signals requiring transport have rates greater than the basic rate [e.g., some Broadband Integrated Services Digital Network (B-ISDN) applications], two techniques of linking several basic modules together to build a transport signal of increased capacity are described (see Sections 3.2.3 and 3.2.5). To maintain a consistent payload structure while providing for the transport of a variety of lower rate payloads (e.g., DS1, DS1C and DS2 signals), a structure called a VT is defined (Section 3.2.4). Payloads below the DS3 rate are transported within a VT structure.

Different types of overhead are defined for functions that include maintenance, protection switching, frequency justification, orderwire, identification and user channels (Section 3.3). Also, growth channels are identified to allow for future uses not defined or conceived of at this time. A layered approach to overhead is established, whereby overhead bandwidth is allocated to a layer based on the function addressed by that particular layer. This layered approach allows the creation of equipment that is not required to access all layers of overhead, thereby allowing equipment implementations to meet different needs.

Section 3.4 considers the mapping of various payloads into payload envelopes. Section 3.5 describes payload pointers, which are a mechanism that allows the payload envelopes to slide relative to the overhead, thus accommodating different signal phases and frame rates in multiplexing.

3.1 Synchronous Hierarchical Rates

The STS-1 is the basic module in SONET, and has a bit rate of 51.840 Mb/s. The optical counterpart of the STS-1 is the Optical Carrier - level 1 (OC-1) signal, while the electrical counterpart of the STS-1 is the STS-1 electrical [a.k.a., Electrical Carrier - level 1 (EC-1)] signal defined in Section 4.4.

The definition of the first level also defines the entire hierarchy of SONET signals because higher-level signals are obtained by synchronously multiplexing lower-level modules. When lower-level modules are multiplexed together, the result is denoted as an STS- N (where N is an integer), which can then be converted to an OC- N or STS- N electrical signal. There is an integer multiple relationship between the rates of the basic STS-1 module and the OC- N or STS- N electrical signals (i.e., the rate of an OC- N is equal to N times the rate of an STS-1).

SONET systems support only certain values of N. [Table 3-1](#) lists these values for the standard STS-N electrical and OC-N interface signals up through N equal to 768, along with the corresponding line rates. Values of N greater than 768 may be addressed in future issues of this or other Telcordia documents.

Table 3-1 Line Rates for Standard SONET Interface Signals (through N = 768)

OC-N Level	STS-N Electrical Level	Line Rate (Mb/s)
OC-1	STS-1 electrical	51.84
OC-3	STS-3 electrical	155.52
OC-12	-	622.08
OC-24 ^a	-	1244.16
OC-48	-	2488.32
OC-192	-	9953.28
OC-768	-	39813.12

Note:

- a. This entry is included for consistency with ANSI T1.105, *Synchronous Optical Network (SONET) – Basic Description including Multiplex Structure, Rates and Formats*. However, no criteria specific to the OC-24 level are provided in this document.

3.2 Transport Format

The SONET transport format presented here is based on ANSI T1.105. The format definition in the following sections designates some bits and bytes as undefined. Suppliers are likely to introduce enhanced features by using these bits and bytes in a nonstandard manner. Network providers who wish to deploy these nonstandard features should study the network implications jointly with the supplier, and recognize the potential equipment incompatibilities.

R3-1 [4] A SONET NE shall have the capability to ignore the values contained in all undefined and unused bits and bytes [except for Bit Interleaved Parity (BIP)-8 calculations] to prevent misinterpretation of the received patterns.

Undefined bits and bytes are those for which no standard use has been defined for the transmitting NE. Unused bits and bytes are those for which no standard use has been defined for the NE when it receives the SONET signal. (The criteria on the use of the currently defined overhead bits and bytes are summarized in [Table 5-2](#).)

O3-2 [5] A SONET NE should send all-zeros patterns (before scrambling) in undefined bits and bytes. All-zeros patterns should also be sent in defined bits and bytes if the NE does not support the defined function or if the function has been disabled by the user.

Note that for many bits and bytes, **O3-2 [5]** is overridden by the Alarm Indication Signal (AIS) generation requirements in [Section 6.2.1.2](#) when the NE is transmitting AIS. Also, as discussed in various sections of this document, an NE that transmits particular non-zero values in certain undefined bits and bytes can still meet the objective. For example, although the B1 byte position is not defined for a drop-side STS-1 electrical signal, an NE may transmit either all-zeros or the Section BIP-8 code in that byte (see [Section 3.3.2.1](#)).

-
- R3-3 [6]** If a supplier introduces a nonstandard feature employing SONET overhead, the supplier shall disclose such use of overhead and furnish the network provider with an equipment option to disable the feature (including the transmission of the nonstandard messages).
-

3.2.1 Frame Structure of the STS-1

An STS-1 is a specific sequence of 810 bytes (6480 bits), which includes various overhead bytes and an envelope capacity for transporting payloads. It can be depicted as a 90 column by 9 row structure, as shown in [Figure 3-1](#). With a frame length of 125 μ s (i.e., 8000 frames per second), the STS-1 has a bit rate of 51.840 Mb/s. Using the structure in [Figure 3-1](#), the order of transmission of bytes is row-by-row, from left to right.

-
- R3-4 [7]** The structure of an STS-1 shall be as shown in [Figure 3-1](#).¹
- R3-5 [8]** In each byte of the STS-1, the most significant bit shall be transmitted first, as shown in [Figure 3-2](#).
-

1. What is meant by “the structure of an STS-1 shall be as shown ...” is that the byte sequence shall be such that it can be mapped into the frame structure shown.

Figure 3-1 STS-1 Frame

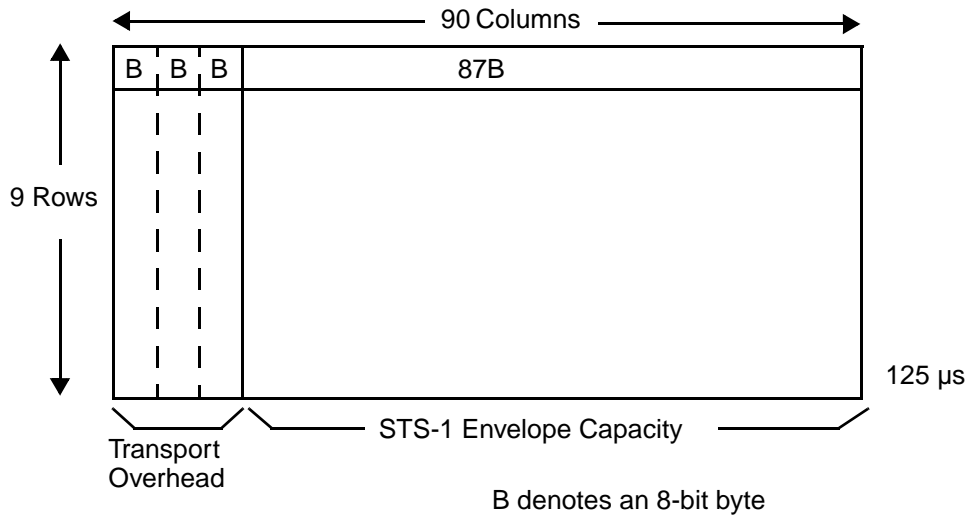



Figure 3-2 Bit Position Numbering



3.2.1.1 Transport Overhead

As [Figure 3-1](#) shows, the first three columns of the STS-1 frame are the transport overhead. These three columns contain 27 bytes, of which nine bytes are overhead for the Section layer (i.e., Section Overhead, or SOH) and 18 bytes are overhead for the Line layer (i.e., Line Overhead, or LOH). [Section 3.3](#) contains the details of these overhead allocations. The remaining 87 columns constitute the STS-1 envelope capacity.

3.2.1.2 STS-1 Envelope Capacity and Synchronous Payload Envelope (SPE)

Figures 3-3, 3-4 and 3-5 depict the STS-1 SPE, which occupies the STS-1 envelope capacity. The STS-1 SPE consists of 783 bytes and can be depicted as an 87 column by 9 row structure. Column 1 contains nine bytes, designated as the STS Path Overhead (POH). Two columns (columns 30 and 59) are not used for payload, but are designated as the “fixed stuff” columns. The 756 bytes in the remaining 84 columns are designated as the STS-1 payload capacity.

R3-6 [9] The structure of an STS-1 SPE shall be as shown in [Figure 3-4](#).

The bytes in the fixed stuff columns are undefined, so the objective in [Section 3.2](#) (to set them to all-zeros) is applicable. However, several possible uses for those bytes have been discussed in the standards bodies, and some suppliers may choose to use them for proprietary purposes. Therefore, for compatibility between the STS-1 path BIP-8 calculation in SONET (which covers all 87 columns of the STS-1 SPE) and the SDH Virtual Container-3 (VC-3) BIP-8 calculation specified in ITU-T Rec. G.707 (which covers all 85 columns of the VC-3), the following requirement applies.

R3-7 [10] The values used to stuff columns 30 and 59 of each STS-1 SPE shall produce even parity in the calculation of the STS-1 path BIP-8 (see [Section 3.3.2.4](#)).

The STS-1 SPE may begin anywhere in the STS-1 envelope capacity. Typically, it begins in one STS-1 frame and ends in the next (although it may be wholly contained in one frame). The STS payload pointer contained in the transport overhead designates the location of the byte where the STS-1 SPE begins. [Section 3.5.1](#) describes STS payload pointers.

STS POH is associated with each payload and is used to communicate various information from the point where a payload is mapped into the STS-1 SPE to where it is delivered. [Section 3.3.2.4](#) contains details on the STS POH.

Figure 3-3 STS-1 Synchronous Payload Envelope

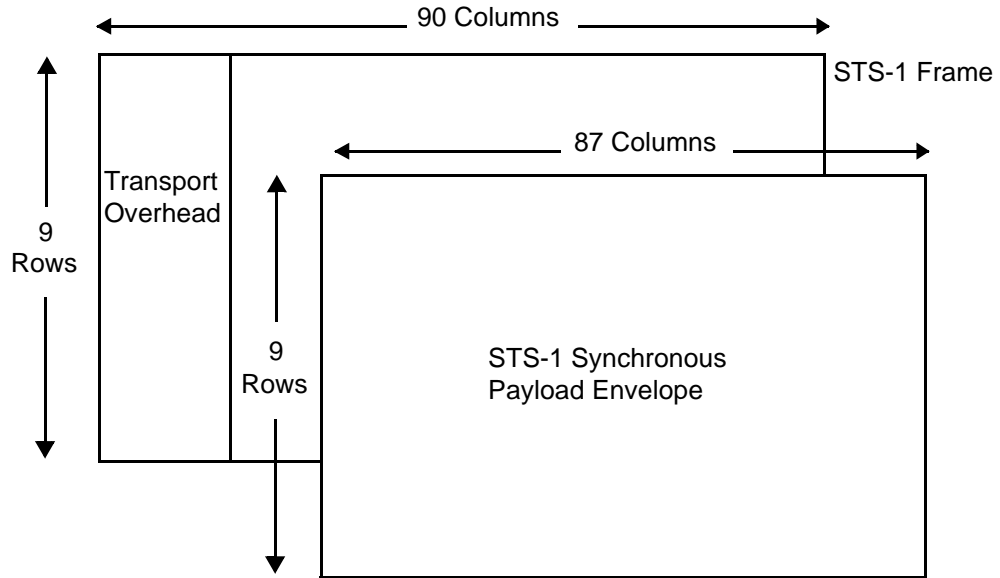


Figure 3-4 STS-1 SPE with STS-1 POH and STS-1 Payload Capacity Illustrated


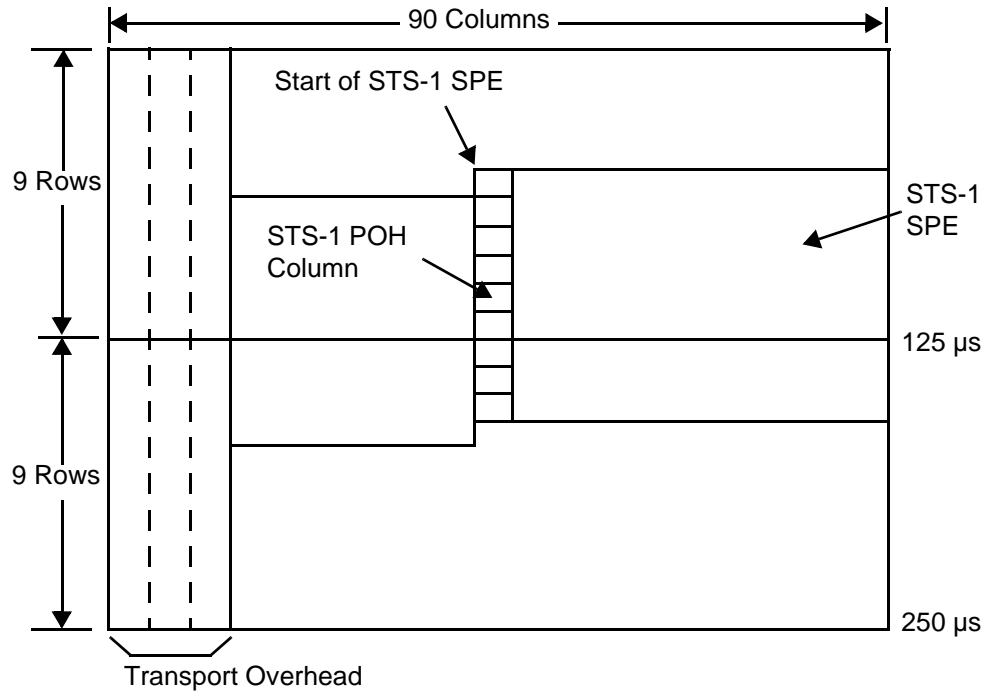
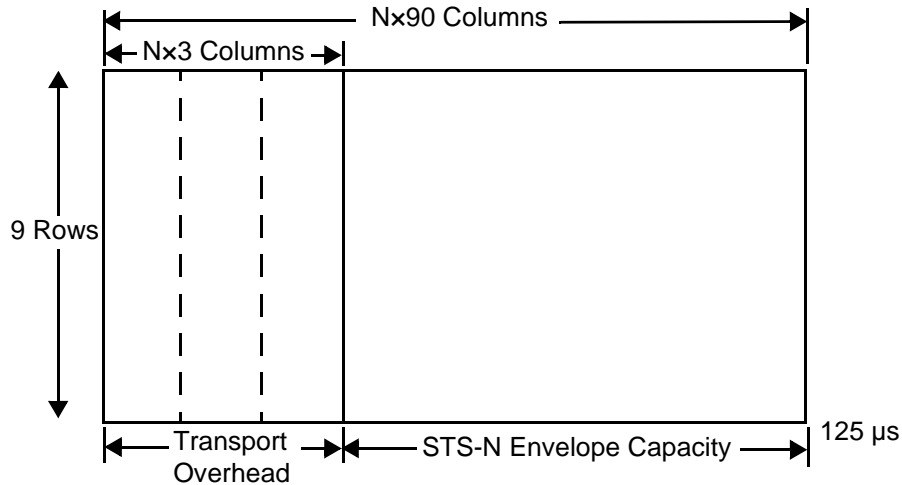


Figure 3-5 STS-1 SPE in Interior of STS-1 Frames



3.2.2 Frame Structure of the STS-N

An STS-N is a specific sequence of $N \times 810$ bytes that can be depicted as the structure shown in Figure 3-6. The STS-N is formed by byte-interleaving STS-1 and STS-M ($3 \leq M < N$) modules. The transport overhead of the individual STS-1 and STS-M modules are frame aligned before interleaving, but the associated STS SPEs are not required to be aligned because each STS-1 has a payload pointer to indicate the location of the SPE (or to indicate concatenation). Section 5.1.1 contains the interleaving requirements.

Figure 3-6 STS-N Frame

3.2.3 STS Concatenation

Multiple STS-1 SPEs are needed to transport super-rate payloads (e.g., some ATM payloads). Two methods of accomplishing this are currently defined, and are discussed in the following sections.

3.2.3.1 Contiguous STS Concatenation

To accommodate a number of the possible super-rate payloads, an “STS-Nc” module is formed by linking N constituent STS-1s together in fixed phase alignment. The super-rate payload is then mapped into the resulting STS-Nc SPE for transport. The STS-Nc SPE can be carried by an OC-N, STS-N electrical, or higher bit-rate signal. This method is referred to as “contiguous concatenation”, and is primarily used to transport payloads that can be mapped efficiently into one of several specific sizes of STS-Nc SPEs.

The need for a SONET NE to be able to generate, multiplex, switch, transport or terminate STS-Nc SPEs depends on the functionality of that NE. Criteria for specific types of NEs are contained in the individual SONET NE GRs and TRs [e.g., GR-496-CORE, *SONET Add-Drop Multiplexer (SONET ADM) Generic Criteria*].

R3-8 [11] If an NE supports the multiplexing, switching or transport of STS-Nc SPEs, then it shall treat each STS-Nc SPE as a single entity.

Concatenation indicators contained in the second through Nth STS payload pointers are used to show that the STS-1s of an STS-Nc are linked together.

An STS-Nc SPE consists of $N \times 783$ bytes, and can be depicted as an $N \times 87$ column by 9 row structure, as shown in [Figure 3-7](#) (which also shows the STS-Nc payload capacity). Only one set of STS POH is required in the STS-Nc SPE. The STS-Nc SPE is carried within the STS-Nc so that the STS POH always appears in the first of the N STS-1s that make up the STS-Nc.

In all of the super-rate payload mappings contained in this document, the first $(N/3)-1$ columns of the STS-Nc SPE following the STS POH are not used for payload, but are designated as fixed stuff columns, (i.e., columns of undefined bytes, see [Section 3.2](#)).² Since the STS-Nc SPE is treated as a single entity, the presence or absence of fixed stuff columns only affects the equipment that generates and terminates the SPE. Therefore, future payload mappings could possibly be defined where these columns are used in the payload mapping.

R3-9 [12] The structure of an STS-Nc SPE shall be as shown in [Figure 3-7](#).

Through Issue 3 of this document, the valid values of “N” for an STS-Nc included 3, 6, 9, ..., 42, 45, 48 and 192. However, only mappings into STS-3c, STS-12c and STS-48c SPEs were defined in that and previous issues of the document, and the new mappings that have been added in this issue do not utilize any of the other possible values between 3 and 48 (see [Section 3.4.2](#)). In addition, most (if not all) SONET products support the transport of only certain size STS-Nc SPEs (e.g., STS-3c, STS-12c, STS-48c and STS-192c SPEs), and the definition of the virtual method of concatenating STS SPEs has essentially eliminated the need for other size STS-Nc SPEs. Therefore, only the values of N noted below (in **R3-10 [1101v2]**) are currently considered defined.

R3-10 [1101v2] Concatenation levels of STS-6c, STS-9c, STS-15c through STS-45c, STS-51c through STS-189c, and STS-195c through STS-765c shall not be used (i.e., contiguous concatenation levels are currently limited to STS-3c, STS-12c, STS-48c, STS-192c and STS-768c).

[Sections 3.5.1.4, 5.1.1 and 5.1.2](#) provide additional details on STS concatenation and the multiplexing of STS-Mc SPEs into STS-N signals. [Section 3.3.2](#) describes the overhead assignments, and [Figure 3-8](#) illustrates the transport overhead assignments for an STS-3 electrical or OC-3 signal carrying an STS-3c SPE. [Section 3.4.2](#) describes various mappings, including mappings of super-rate payloads into STS-Nc SPEs.

2. Note that this implies N must be divisible by 3. It also means that an STS-3c SPE has no fixed stuff columns (although the payload mapping contained in the SPE may contain one or more columns of fixed stuff bytes).

Figure 3-7 STS-Nc SPE

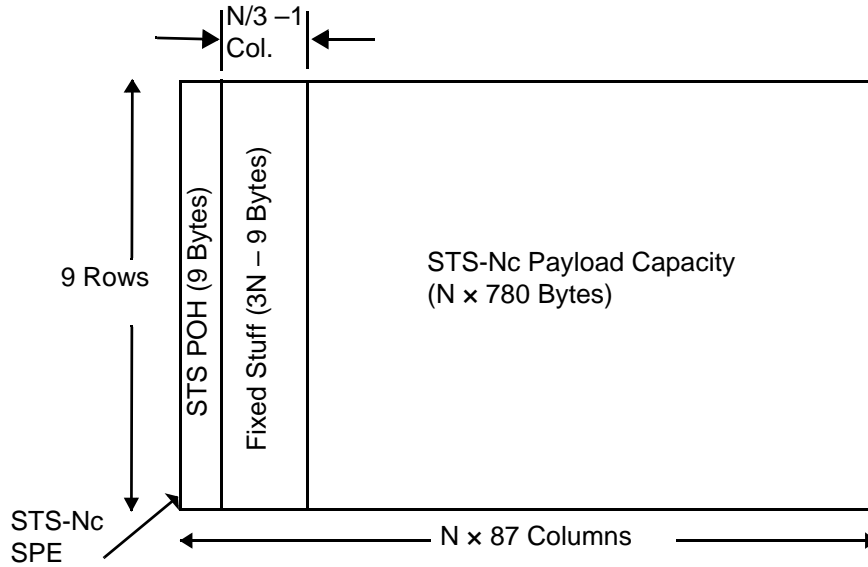
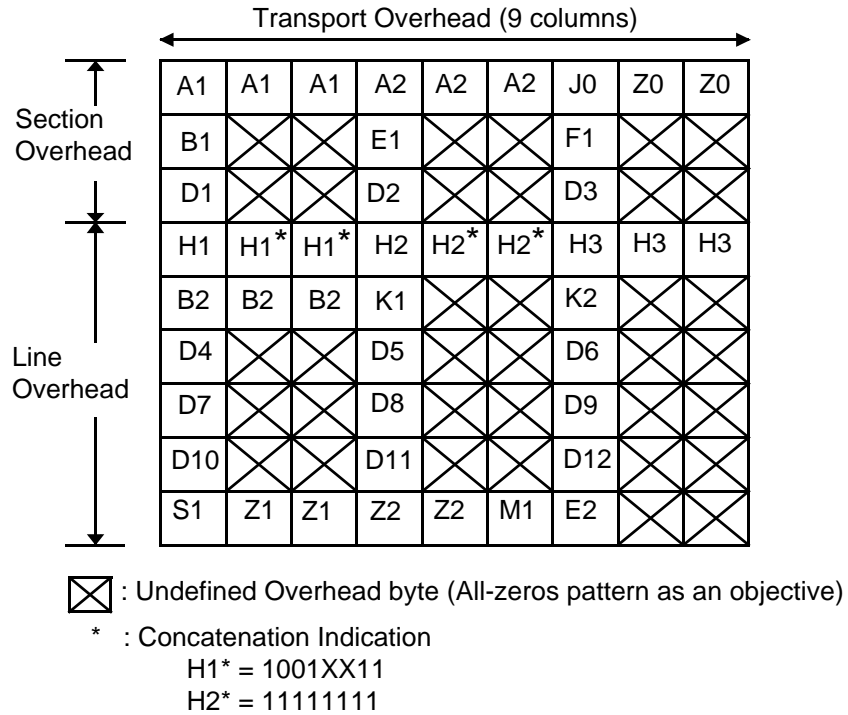


Figure 3-8 Transport Overhead Assignment, OC-3 Carrying an STS-3c SPE



3.2.3.2 STS Virtual Concatenation

In STS virtual concatenation, a single payload is mapped into an STS-1-Xv or STS-3c-Xv ($1 < X \leq 256$) payload capacity, which is then mapped (or demultiplexed) into the payload capacities of X STS-1 or STS-3c SPEs. These constituent SPEs, each of which is identified by a sequence number in the range from 0 to $(X-1)$ and which are referred to collectively as an STS-1-Xv or STS-3c-Xv SPE, are then multiplexed in the desired order into one or more OC-N signals for independent transport across the network. At the far-end, all of the virtually concatenated STS path signals are terminated and the payload is reassembled. Unlike the contiguous concatenation case discussed in [Section 3.2.3.1](#), virtual concatenation need only be supported by the Path Terminating Equipment (PTE) at each end. [That is, intermediate equipment such as Line Terminating Equipment (LTE) and STS path selectors need only be able to multiplex, switch or transport STS-1 or STS-3c SPEs, not some larger entity such as an STS-12c SPE.] In addition, the values of X that can be utilized are not limited like the values of M that can be used in contiguous concatenation, and therefore virtual concatenation can be used to efficiently transport payloads having a wide variety of different bit rates. (Note that the payload capacity of an STS-1-Xv is $X \times 48.384$ Mb/s, while that of an STS-3c-Xv is $X \times 149.760$ Mb/s.)

Since the STS path signals used in the transport of payloads via STS virtual concatenation are STS-1 or STS-3c path signals, no additional criteria related to the formats or frame structures of those entities are needed here. On the other hand, the following requirements apply regarding the mapping of the STS-1-Xv or STS-3c-Xv payload capacity into the X STS-1 or STS-3c SPEs (and the reassembly of the STS-1-Xv or STS-3c-Xv payload capacity at the far end). That mapping is done on a byte-by-byte basis, with the first byte of the STS-1-Xv or STS-3c-Xv payload capacity occupying the first byte of the payload capacity of STS-1 or STS-3c SPE #0 (i.e., the byte immediately following the J1 byte in the first SPE), the second byte of the STS-1-Xv or STS-3c-Xv payload capacity occupying the first byte of the payload capacity of STS-1 or STS-3c SPE #1, etc.

R3-11 [1139] The mapping between an STS-1-Xv payload capacity and the X constituent STS-1 SPEs shall be as shown in [Figure 3-9](#).

R3-12 [1140] The mapping between an STS-3c-Xv payload capacity and the X constituent STS-3c SPEs shall be as shown in [Figure 3-10](#).

In addition, since the constituent STS path signals are transported independently through the network, they may be subject to different amounts of transmission delay [e.g., as a result of traveling in opposite directions between the input and output nodes in a Unidirectional Path Switched Ring (UPSR)]. Thus, it is necessary for the PTE supporting the virtual concatenation function to be able to identify and accommodate such differences. The method by which this done is discussed in [Section 5.2.5](#), which also contains criteria related to a Link Capacity Adjustment Scheme (LCAS) that may be supported by the PTE.

Figure 3-9 Mapping of the STS-1-Xv Payload Capacity Into X STS-1 SPEs





Figure 3-10 Mapping of the STS-3c-Xv Payload Capacity Into X STS-3c SPEs



3.2.4 Virtual Tributary Structure

The VT structure is designed for transport and switching of sub-STS-1 rate payloads. There are four sizes of VTs: VT1.5 (1.728 Mb/s), VT2 (2.304 Mb/s), VT3 (3.456 Mb/s) and VT6 (6.912 Mb/s). These are illustrated in Figure 3-11. In the 87-column by 9-row structure of the STS-1 SPE, these VTs occupy 3, 4, 6 and 12 columns, respectively.

To efficiently accommodate a mix of VT sizes, the VT-structured STS-1 SPE is divided into seven VT groups. Each VT group occupies 12 columns of the 87-column STS-1 SPE, and may contain 4 VT1.5s, 3 VT2s, 2 VT3s or 1 VT6. A VT group can contain only one size of VTs; however, a different VT size is allowed for each VT group in an STS-1 SPE.

Figures 3-12, 3-14, 3-16 and 3-18 each show all seven VT groups in an STS-1 SPE containing one of the four VT sizes. The tables in Figures 3-13, 3-15, 3-17 and 3-19 define the relationship between the VT group number and VT number, and the columns in the STS-1 SPE from Figures 3-12, 3-14, 3-16 and 3-18, respectively. These tables are applicable in all cases, including VT groups in an STS-1 SPE with different VT sizes. Figures 3-20 and 3-21 illustrate an example where the first four VT groups contain VT1.5s, VT2s, VT3s and VT6s.

-
- R3-13** [13] The structure of a VT-structured STS-1 SPE shall be consistent with the structures shown in Figures 3-11 through 3-21.
-

In addition to the division of VTs into VT groups, a 500- μ s structure called a VT superframe is defined for each VT. The VT superframe contains the V1 and V2 bytes (the VT payload pointer), the V3 byte (the VT pointer action byte), the V4 byte (an undefined byte), and the VT envelope capacity, which in turn contains the VT SPE. The VT envelope capacity, and therefore the size of the VT SPE, is different for each VT size, as shown in Figure 3-22. V1 is the first byte in the VT superframe, while V2 through V4 appear as the first bytes in the following frames of the VT superframe, regardless of the VT size.

-
- R3-14** [14] Four consecutive 125- μ s frames of the VT-structured STS-1 SPE shall be organized into a 500- μ s VT superframe, the phase of which is indicated by the H4 (Indicator) byte in the STS POH (see Section 3.4.1).
-

The VT payload pointer provides for flexible and dynamic alignment of the VT SPE within the VT envelope capacity, independent of other VT SPEs. Section 3.5.2 further describes the VT payload pointers. Figure 3-23 illustrates the VT SPEs corresponding to the four VT sizes. Each VT SPE contains four bytes of VT POH (V5, J2, Z6 and Z7), and the remaining bytes constitute the VT payload capacity, which is different for each VT size. Section 3.4.1 describes the mappings for various payloads (e.g., DS1, DS1C, DS2) into VT SPEs.

-
- R3-15** [15] The structure of a VT SPE shall be as shown in Figure 3-23.
-

Figure 3-11 VT Sizes




Figure 3-12 VT Structured STS-1 SPE: All VT1.5s




Figure 3-13 VT1.5 Locations

<u>VT Group #, VT #</u>	<u>Columns</u>	
1, 1	2, 31, 60	
2, 1	3, 32, 61	
3, 1	4, 33, 62	Column 1 = STS-1 POH
4, 1	5, 34, 63	30 = Fixed Stuff
5, 1	6, 35, 64	59 = Fixed Stuff
6, 1	7, 36, 65	
7, 1	8, 37, 66	
1, 2	9, 38, 67	
2, 2	10, 39, 68	
3, 2	11, 40, 69	
4, 2	12, 41, 70	
5, 2	13, 42, 71	
6, 2	14, 43, 72	
7, 2	15, 44, 73	
1, 3	16, 45, 74	
2, 3	17, 46, 75	
3, 3	18, 47, 76	
4, 3	19, 48, 77	
5, 3	20, 49, 78	
6, 3	21, 50, 79	
7, 3	22, 51, 80	
1, 4	23, 52, 81	
2, 4	24, 53, 82	
3, 4	25, 54, 83	
4, 4	26, 55, 84	
5, 4	27, 56, 85	
6, 4	28, 57, 86	
7, 4	29, 58, 87	

Note that the corresponding table in the 2003 version of ITU Rec. G.707 (i.e., Table VI/G.707/Y.1322, *Relationship between TU-11 address and location of columns within a VC-3*) is incorrect. On the other hand, the equation provided in clause 7.3.13 of that document, *Numbering of TU-11s in a VC-3*, provides results that are consistent with the information provided here (once the fixed stuff columns, which are present in an STS-1 SPE but not in a VC-3, are taken into account).

Figure 3-14 VT Structured STS-1 SPE: All VT2s




Figure 3-15 VT2 Locations

<u>VT Group #, VT #</u>	<u>Columns</u>	
1, 1	2, 23, 45, 67	
2, 1	3, 24, 46, 68	
3, 1	4, 25, 47, 69	
4, 1	5, 26, 48, 70	
5, 1	6, 27, 49, 71	Column 1 = STS-1 POH
6, 1	7, 28, 50, 72	30 = Fixed Stuff
7, 1	8, 29, 51, 73	59 = Fixed Stuff
1, 2	9, 31, 52, 74	
2, 2	10, 32, 53, 75	
3, 2	11, 33, 54, 76	
4, 2	12, 34, 55, 77	
5, 2	13, 35, 56, 78	
6, 2	14, 36, 57, 79	
7, 2	15, 37, 58, 80	
1, 3	16, 38, 60, 81	
2, 3	17, 39, 61, 82	
3, 3	18, 40, 62, 83	
4, 3	19, 41, 63, 84	
5, 3	20, 42, 64, 85	
6, 3	21, 43, 65, 86	
7, 3	22, 44, 66, 87	

Figure 3-16 VT Structured STS-1 SPE: All VT3s




Figure 3-17 VT3 Locations

<u>VT Group #, VT #</u>	<u>Columns</u>	
1, 1	2, 16, 31, 45, 60, 74	
2, 1	3, 17, 32, 46, 61, 75	
3, 1	4, 18, 33, 47, 62, 76	
4, 1	5, 19, 34, 48, 63, 77	Column 1 = STS-1 POH
5, 1	6, 20, 35, 49, 64, 78	30 = Fixed Stuff
6, 1	7, 21, 36, 50, 65, 79	59 = Fixed Stuff
7, 1	8, 22, 37, 51, 66, 80	
1, 2	9, 23, 38, 52, 67, 81	
2, 2	10, 24, 39, 53, 68, 82	
3, 2	11, 25, 40, 54, 69, 83	
4, 2	12, 26, 41, 55, 70, 84	
5, 2	13, 27, 42, 56, 71, 85	
6, 2	14, 28, 43, 57, 72, 86	
7, 2	15, 29, 44, 58, 73, 87	

Figure 3-18 VT Structured STS-1 SPE: All VT6s




Figure 3-19 VT6 Locations

<u>VT Group #, VT #</u>	<u>Columns</u>
1, 1	2, 9, 16, 23, 31, 38, 45, 52, 60, 67, 74, 81
2, 1	3, 10, 17, 24, 32, 39, 46, 53, 61, 68, 75, 82
3, 1	4, 11, 18, 25, 33, 40, 47, 54, 62, 69, 76, 83
4, 1	5, 12, 19, 26, 34, 41, 48, 55, 63, 70, 77, 84
5, 1	6, 13, 20, 27, 35, 42, 49, 56, 64, 71, 78, 85
6, 1	7, 14, 21, 28, 36, 43, 50, 57, 65, 72, 79, 86
7, 1	8, 15, 22, 29, 37, 44, 51, 58, 66, 73, 80, 87

Column 1 = STS-1 POH
30 = Fixed Stuff
59 = Fixed Stuff

Figure 3-20 Example of VT Structured STS-1 SPE

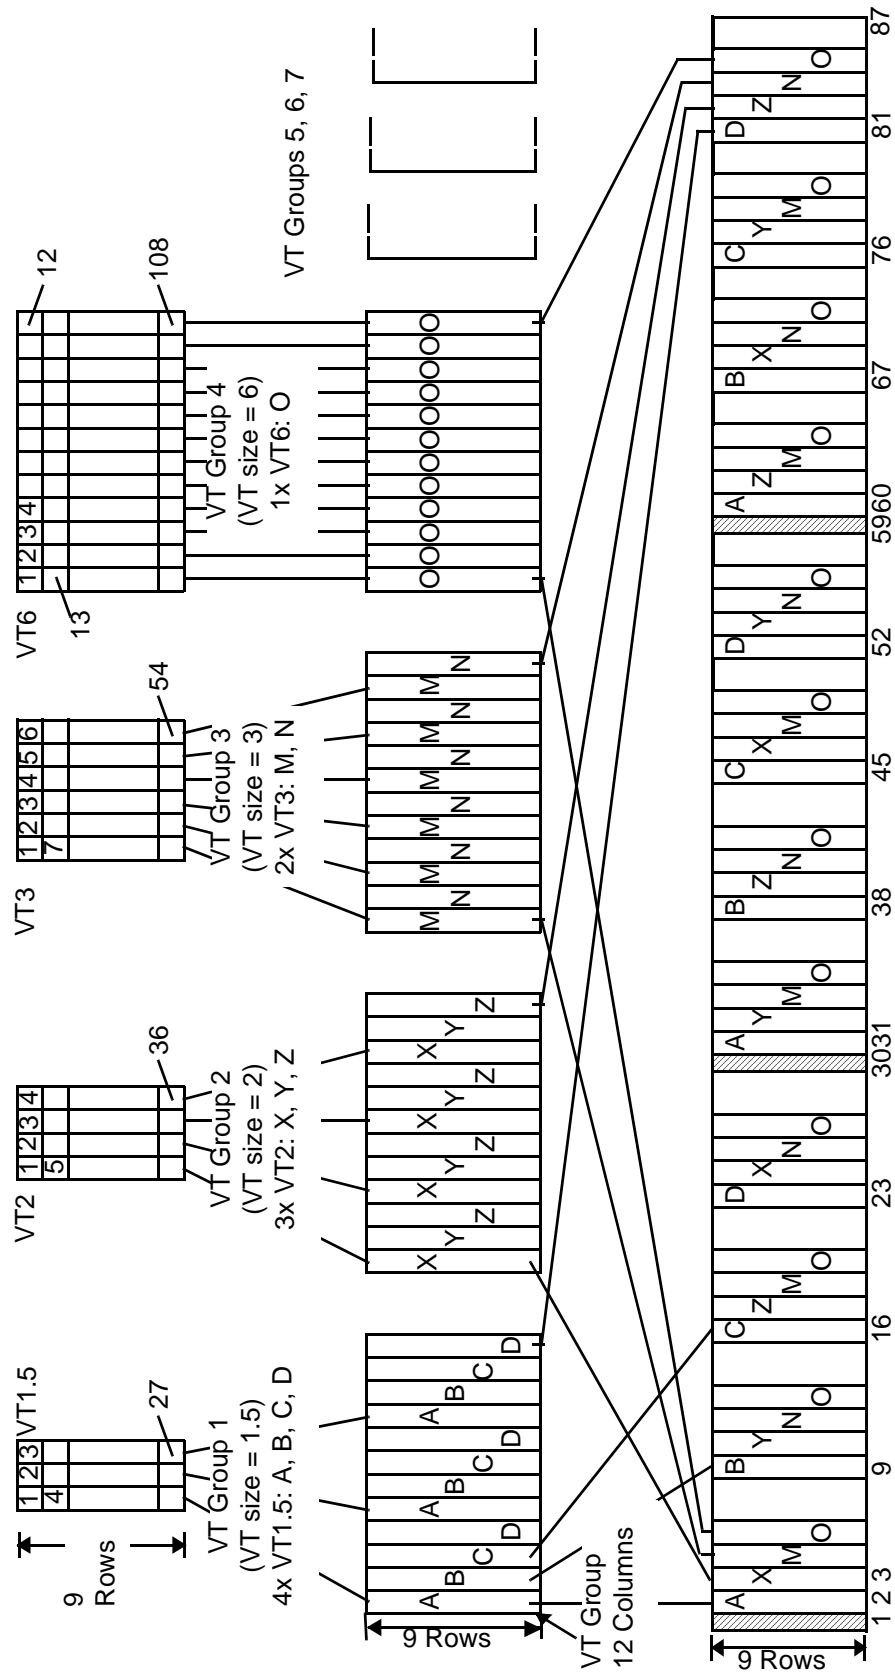


Figure 3-21 Correspondence Between Labels and Numbers for the Example in [Figure 3-20](#)

<u>Label</u>	<u>VT Group #, VT #</u>
A	1, 1
B	1, 2
C	1, 3
D	1, 4
X	2, 1
Y	2, 2
Z	2, 3
M	3, 1
N	3, 2
O	4, 1

Figure 3-22 VT Superframe and Envelope Capacity

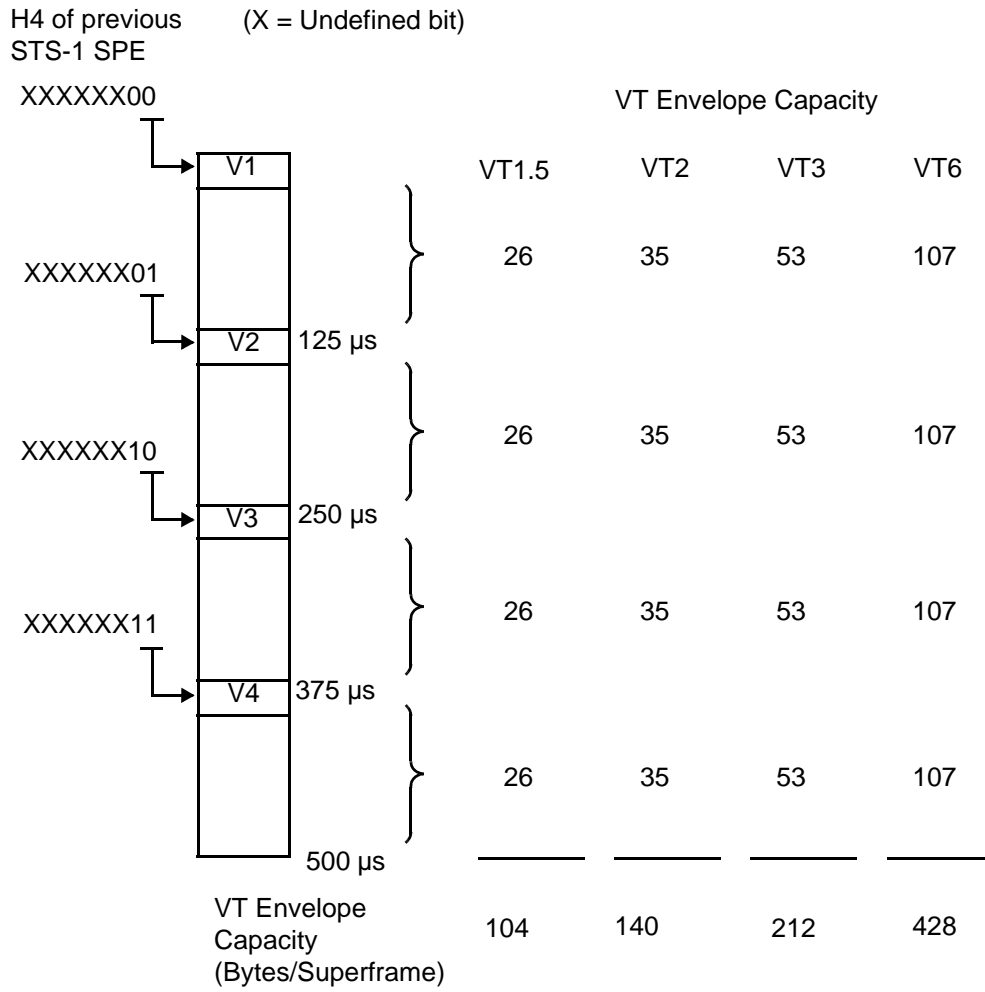
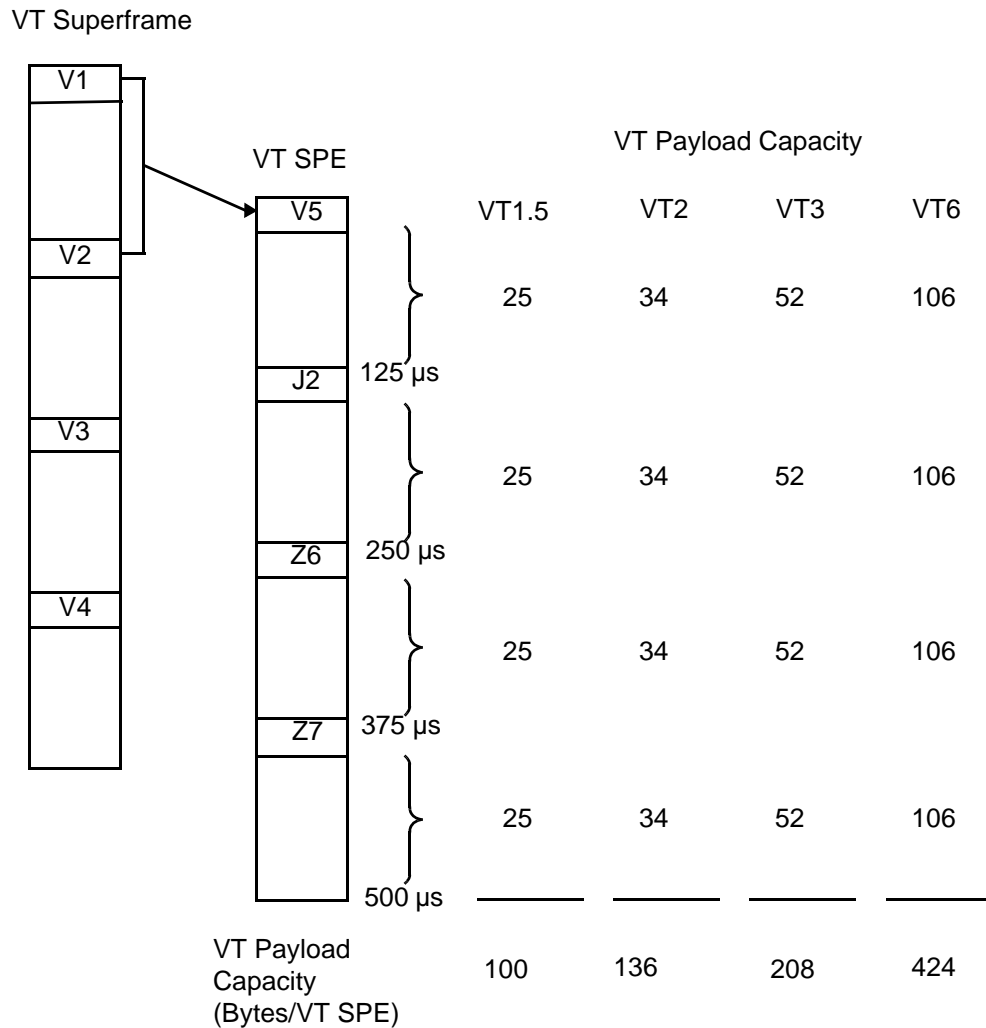


Figure 3-23 VT SPE and Payload Capacity



3.2.5 VT Concatenation

Although contiguous concatenation of VC-2s within a VC-3 is defined for use in SDH networks in ITU-T Rec. G.707, virtual concatenation is currently the only type of concatenation that is defined for VTs (i.e., in SONET). Similar to the STS case discussed in Section 3.2.3.2, in VT virtual concatenation a single payload is mapped into a VT n -X v ($n = 1.5$, or possibly 2, 3 or 6, and $1 < X \leq 64$) payload capacity, which is then mapped (or demultiplexed) into the payload capacities of X VT n SPEs.³ These constituent SPEs, each of which is identified by a sequence number in the range from 0 to ($X-1$) and which are referred to collectively as a VT n -X v SPE, are then multiplexed in the desired order into one or more STS-1 (or possibly an STS-3c, see footnote 3) SPEs for independent transport across the network. At the far-end, all of the virtually concatenated VT n path signals are terminated and the payload is reassembled. Also as in the STS virtual concatenation case, VT virtual concatenation need only be supported by the PTE at each end. (That is, intermediate equipment

need only be able to multiplex, switch or transport VTn SPEs.) In addition, the values of X that can be utilized are limited only by the number of possible sequence numbers, and therefore virtual concatenation can be used to efficiently transport payloads having a wide variety of different bit rates. (Note that the payload capacity of a VT1.5- X_v is $X \times 1.6$ Mb/s.)

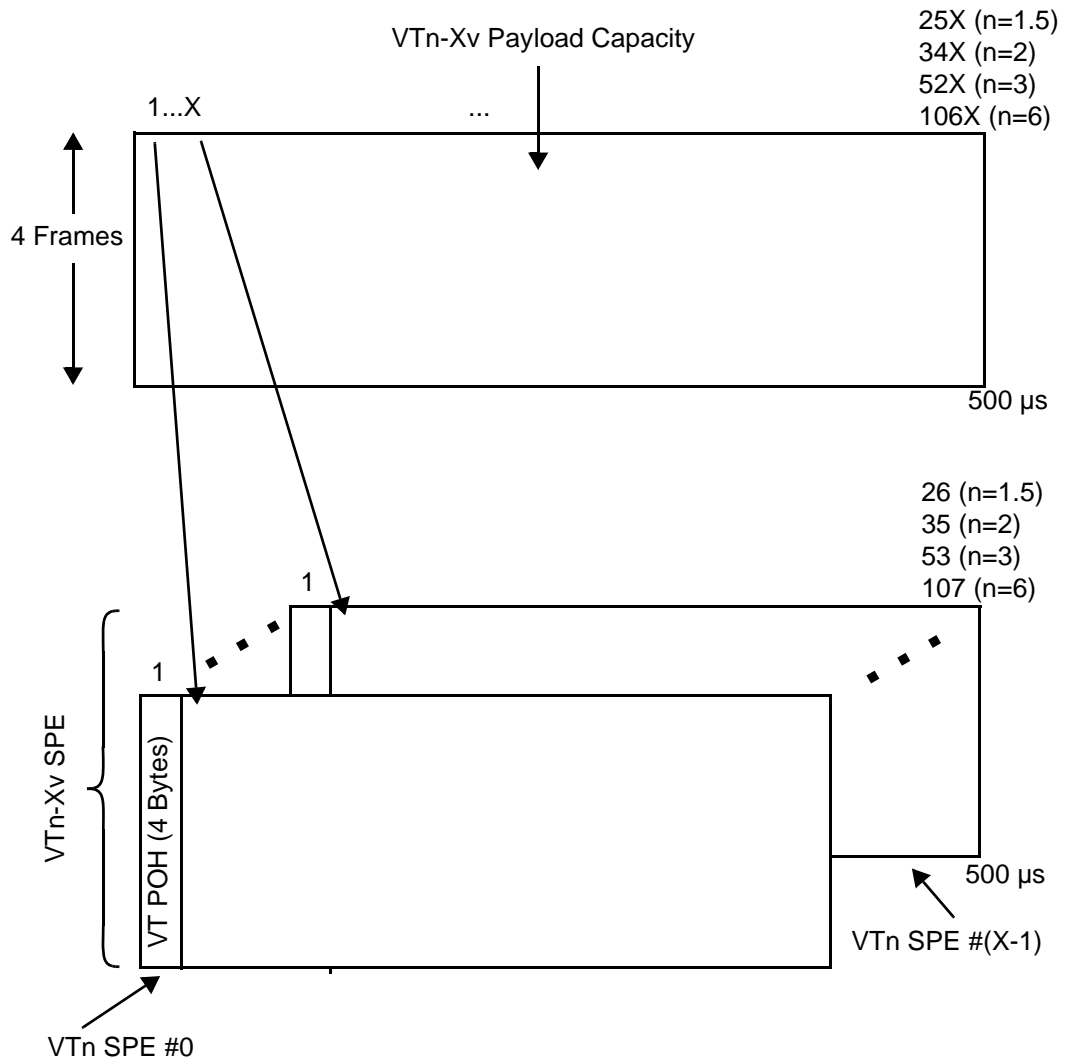
Since the VTn path signals used in the transport of payloads via VT virtual concatenation are VT1.5, VT2, VT3 or VT6 path signals, no additional criteria related to the formats or frame structures of those entities are needed here. On the other hand, the following requirement applies regarding the mapping of the VTn- X_v payload capacity into the X VTn SPEs (and the reassembly of the VTn- X_v payload capacity at the far end). That mapping is done on a byte-by-byte basis, with the first byte of the VTn- X_v payload capacity occupying the first byte of the payload capacity of VTn SPE #0 (i.e., the byte immediately following the V5 byte in the first SPE), the second byte of the VTn- X_v payload capacity occupying the first byte of the payload capacity of VTn SPE #1, etc.

R3-16 [1141] The mapping between a VTn- X_v payload capacity and the X constituent VTn SPEs shall be as shown in [Figure 3-24](#).

In addition, since the constituent VTn path signals are transported independently through the network, they may be subject to different amounts of transmission delay. Thus, it is necessary for the PTE supporting the virtual concatenation function to be able to identify and accommodate such differences. The method by which this done is discussed in [Section 5.2.5](#), which also contains criteria related to the LCAS function that may be supported by the PTE.

-
3. Due to the minimal support of VT sizes other than VT1.5 by existing SONET NEs, and the lack of criteria and standards related to the mapping of VTs into STS-3c SPEs, it is expected that practically all SONET VT virtual concatenation applications will utilize VT1.5s carried in STS-1 SPEs. However, the use of VT2s, VT3s, VT6s and STS-3c SPEs is also allowed, and therefore those cases are included here for completeness. In addition, it should be noted that in SDH most lower order VC (VT) virtual concatenation applications are expected to utilize VC-2s (the SDH equivalent of VT6s) carried in VC-4s (the equivalent of STS-3c SPEs). Finally, if a SONET application happens to utilize an STS-3c SPE to transport VTs, the mapping of those VTs would be expected to meet the specifications in ITU-T Rec. G.707 for mapping VC-11s, VC-12s or VC-2s into a VC-4 [via Tributary Unit Group-3s (TUG-3s)].

Figure 3-24 Mapping of the VT_n-X_v Payload Capacity Into X VT_n SPEs



3.3 Layered Overhead and Transport Functions

The overhead and transport functions in SONET are broken into layers. These are the Physical, Section, Line, STS path and (in cases where the payloads are carried using VT-structured STS-1 SPEs) VT path layers.⁴ The layers have a hierarchical relationship and are considered from the top down to provide a general introduction to the individual layers and their functionalities.

Each layer requires the services of all lower layers to perform its functions (see Figure 3-25). For example, when two STS path layer processes exchange DS3s, the path layer maps the DS3 signal and the STS POH into an STS-1 SPE, which is then given (as an internal STS path layer signal) to the Line layer. The Line layer

multiplexes several SPEs from the STS path layer (frame and frequency aligning each one) and adds LOH. Finally, the Section layer adds SOH and performs scrambling before transmission by the Physical layer.⁵

3.3.1 SONET Interface Layers

This section describes each layer in detail. Each description includes a broad classification of the layer, followed by a specification of the main functions it provides. Finally, examples of system hardware associated with the layer are given to clarify the role it plays. [Figure 3-25](#) depicts the relationship of the layers to each other.

3.3.1.1 Physical Layer

The Physical layer deals with the transport of bits as optical or electrical pulses across the physical medium. No overhead is associated with the Physical layer.

The main function of this layer is conversion between internal STS-N signals and external optical or electrical SONET signals. Issues dealt with at this layer include pulse shape, power levels and line code. As an example, electro-optical units communicate at this level.

3.3.1.2 Section Layer

The Section layer deals with the transport of an STS-N frame across the physical medium. This layer uses the Physical layer for transport.

Functions of this layer include framing, scrambling, Section error monitoring and Section-level communications overhead [e.g., Local Orderwire (LOW)]. The SOH is interpreted and modified or created by Section Terminating Equipment (STE).

4. In addition to these layers, an STS (or higher order) tandem connection sub-layer has been defined in ANSI T1.105.05, *Synchronous Optical Network (SONET): Tandem Connection Maintenance* and ITU-T Rec. G.707, and a VT (or lower order) tandem connection sub-layer has been defined in Rec. G.707. When invoked, these sub-layers occur between the Line and STS path layers and between the STS path and VT path layers, respectively, and provide specific performance monitoring capabilities. These sub-layers are currently not discussed further in this document. Also, an FEC sub-layer that can reside between the Section and Line layers has been defined for OC-48, OC-192 and OC-768 signals in ANSI T1.105.08 and Rec. G.707 (see [Section 3.3.2.2](#)).
5. Although this description (e.g., that the Line layer “adds” the LOH) could be interpreted to mean that the lower layer overhead bytes are not present in the signals passed from one layer to the next lower layer, that is not the intent. There are no criteria concerning the format of the internal signals used in an NE, and some (or all) of the overhead bytes may be present in the internal signals. If they are present, then it would be the function of the lower layer to overwrite those bytes as necessary to create the appropriate signal to pass to the next layer.

The Section and Physical layers can be used in some equipment [e.g., the STE regenerator described in TR-NWT-000917, *SONET Regenerator (SONET RGTR) Equipment Generic Criteria*] without involving the higher layers.

3.3.1.3 Line Layer

The Line layer deals with the transport of STS path layer payloads across the physical medium. All lower layers exist to provide transport for this layer.

This layer provides synchronization and multiplexing functions for the STS path layer. The overhead associated with these functions includes overhead for maintenance and line protection purposes and is inserted into the LOH channels. The LOH is interpreted and modified or created by LTE. To access the LOH, the SOH must first be terminated. Therefore, an NE that contains LTE will also contain STE.

An example of system equipment that communicates at this level is an OC-M to OC-N multiplex.

3.3.1.4 Path Layers

The STS and VT path layers deal with the transport of various payloads between SONET terminal multiplexing equipment. Examples of such payloads are DS1s and DS3s.

The path layers map the payloads into the format required by the Line layer. In addition, these layers communicate end-to-end via the STS and VT Path Overhead (POH), which is interpreted and modified or created by PTE. To access the STS POH, the SOH and LOH must first be terminated. Therefore, an NE that contains STS PTE will also contain STE and LTE. Similarly, to access the VT POH the Section, Line and STS path overhead must first be terminated, and therefore an NE that contains VT PTE will also contain STE, LTE and STS PTE.

An example of the system equipment that communicates at this level is DS3 to STS-1 mapping circuits.

3.3.1.5 Interaction of the Layers

Figure 3-25 depicts the interaction of the layers for the case of an optical interface. Each layer:

- Communicates horizontally to peer equipment in that layer
- Processes certain information and passes it vertically to the adjacent layers.

The interactions are described in terms of each level's horizontal and vertical transactions.

Figure 3-25 also shows payloads as inputs to the VT and/or STS path layers. These layers transmit payloads and POH horizontally to their peer entities. The VT path layer (if present) maps the payloads and VT POH into VT SPEs that it passes vertically to the STS path layer as internal VT path layer signals. The STS path layer

creates STS SPEs by synchronizing and multiplexing VT SPEs or mapping payloads and inserting STS POH, and then passes those STS SPEs vertically to the Line layer as internal STS path layer signals.

The Line layer transmits the STS SPEs and LOH to its peer entities. It maps the STS SPEs and LOH into internal Line layer signals. The STS SPEs are synchronized and multiplexed at this time, and then the internal Line layer signal is passed to the Section layer.

The Section layer transmits STS-N signals to its peer entities. It maps the internal Line layer signals and the SOH into an internal STS-N signal that is handed to the Physical layer, which transmits optical or electrical pulses to its peer entities.

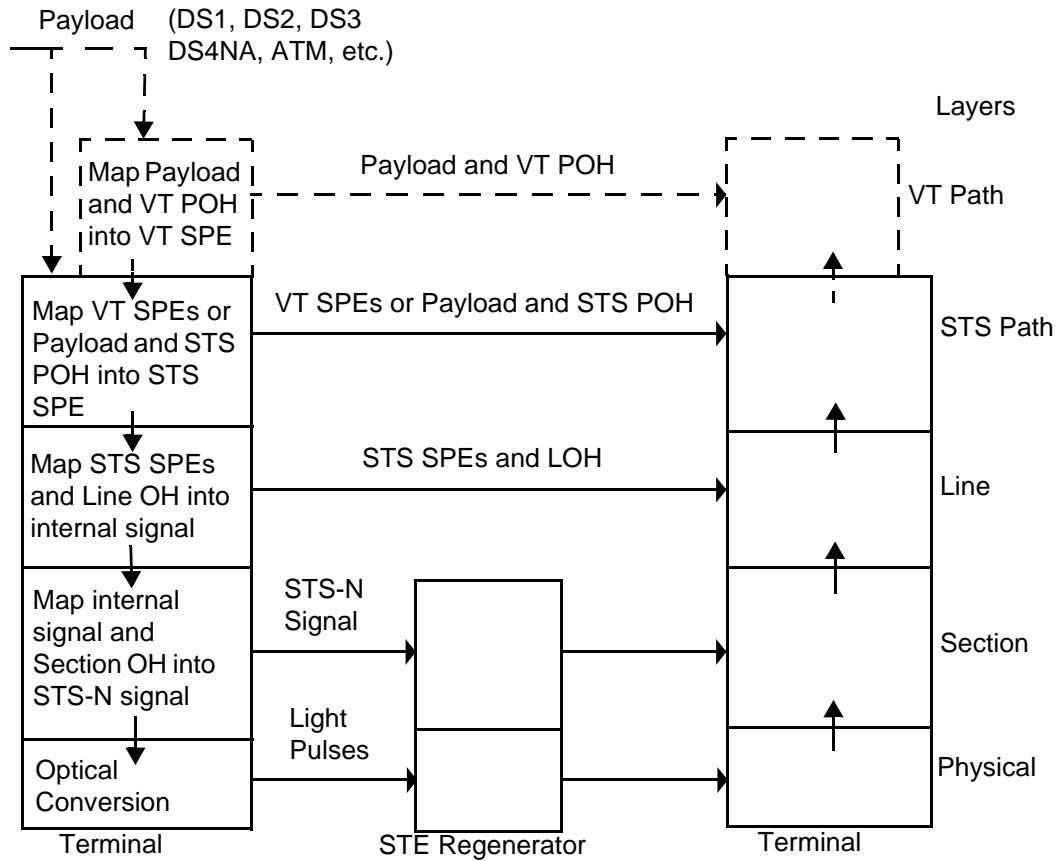
Access to all of the layers is not required of every SONET NE. For example, an STE regenerator would use only the first two layers (Physical and Section). Similarly, an NE that merely routes STS SPEs and does not accept any new inputs from the STS path layer uses only the first three layers (Physical, Section and Line). Note however, that NEs may monitor (or in some cases, may be required to monitor) the overhead of layers that they do not terminate.

R3-17 [1142] Except in cases where other criteria apply [e.g., criteria related to DS1 Robbed Bit Signaling (RBS) removal and insertion, the generation of AIS, or payload bit-rate retiming], bit integrity shall be maintained by a SONET NE that maps or demaps a payload signal (see [Section 3.4](#)).

R3-18 [1143] Except in cases where other criteria apply (e.g., criteria related to the generation of AIS, or the use or compensation of certain overhead bytes by a sub-layer that supports a non-terminated SONET layer), bit integrity shall be maintained by a SONET NE for any SONET layers that it does not terminate.

For the purpose of these requirements, to maintain bit integrity means that the logical content of the payload or non-terminated SONET signal is unchanged except for bit errors caused by internal processes, and that the BER resulting from such errors is significantly less than the BER caused by optical or electrical transmission issues (e.g., the internal processing BER is much less than 1×10^{-12}). Also, unless access is provided to the internal signals used by a SONET NE (e.g., in a product design and development environment), it will generally not be possible to separate the internal processes' contribution to the overall BER from the receiver's contribution. Therefore, to test a SONET NE against these requirements it will generally be necessary to minimize the BER at the receiver. For example, this may be done by providing some type of supplier-specified "optimal" signal (with respect to the power level, pulse shape or eye diagram, etc.) to the receiver. Finally, in order to meet **R3-17 [1142]**, a SONET desynchronizer that supports an asynchronous payload mapping must recover the payload clock and use it for timing the output signal. Thus, a SONET NE is transparent to payload timing when it uses one of the asynchronous mappings defined in [Section 3.4](#).

Figure 3-25 Optical Interface Layers



3.3.2 STS-1 Overhead Descriptions

Figure 3-5 illustrated the location of the overhead bytes in the STS-1 frame. The functions assigned to the Section and Line layers have been combined into a structure of 27 bytes called the transport overhead, which occupies the first 3 columns of the frame. The functions of the STS path layer have been assigned nine bytes in the first column of the STS SPE.

This section defines each of the overhead bytes. Each byte is assigned to a layer and a position in the overhead columns as shown in Figure 3-26. The overhead associated with a given layer is modified or created by the equipment terminating that layer before insertion on the outgoing signal. [Note that where a sub-layer is defined (see footnote 4 on page 3-30), equipment supporting that sub-layer function will modify certain overhead bytes assigned to one or both of the adjacent layers. Also note that there is no overhead associated with the Physical layer.]

In addition, two types of interface signals are defined. These are line-side signals, which have full transport overhead functionality, and drop-side signals, which have reduced transport overhead functionality. Line-side signals are suitable for inter-office connections, although they could also be used for intra-office

connections. Drop-side signals are suitable for intra-office connections. The overhead functionality criteria applicable to line-side and drop-side signals are summarized in [Table 5-2](#).

Figure 3-26 Transport and STS Path Overhead Byte Designations

	Transport Overhead			STS Path Overhead
Section	Framing	Framing	Trace/Growth	Trace
	A1/Reserved ^a	A2/Reserved ^a	J0/Z0 ^b	J1
	BIP-8	Orderwire	User	BIP-8
	B1/Undefined ^{b, c}	E1/Undefined ^{b, c}	F1/Undefined ^b	B3
Line	Data Com	Data Com	Data Com	Signal Label
	D1/Undefined ^{b, c}	D2/Undefined ^{b, c}	D3/Undefined ^{b, c}	C2
	Pointer	Pointer	Pointer Action	Path Status
	H1	H2	H3	G1
	BIP-8	APS	APS	User Channel
	B2	K1/Undefined ^{b, c}	K2/Undefined ^{b, c}	F2
	Data Com	Data Com	Data Com	Indicator
	D4/Undefined ^{b, d}	D5/Undefined ^b	D6/Undefined ^{b, c}	H4
	Data Com	Data Com	Data Com	Growth
	D7/Undefined ^{b, d}	D8/Undefined ^b	D9/Undefined ^{b, c}	Z3
Data Com	Data Com	Data Com	Growth	
D10/Undefined ^{b, d}	D11/Undefined ^b	D12/Undefined ^{b, c}	Z4	
Sync Status/Growth	REI-L ^e /Growth	Orderwire	STS Tandem Connection	
S1/Z1/Undefined ^{b, c}	M1 ^f or M1/Z2/Undefined ^b or M0&M1/Z2/Undefined ^g	E2/Undefined ^b	N1	

Notes:

- For signals at rates from OC-1 through OC-192, all of the bytes in these positions are framing bytes (i.e., A1 or A2 bytes). For OC-768 signals, only the last 64 of the 768 A1 byte positions and the first 64 of the 768 A2 byte positions contain framing bytes. The remaining 1408 byte positions are reserved for future standardization.
- For these entries, the first label shown is applicable for one STS-1 in an STS-N electrical or OC-N signal, the second label is applicable for certain (OC-192 and OC-768) or all (OC-3/STS-3 electrical through OC-768) of the remaining STS-1s, and the third label (if present) is applicable to the remaining STS-1s in an OC-192 or OC-768 signal.
- For OC-48, OC-192 and OC-768 signals, certain of these bytes are defined for use in the FEC function (see [Section 3.3.2.2](#)).
- For OC-768 signals, 144 of these bytes are reserved for possible use in an extended Line Data Communications Channel (DCC) function.
- REI-L (Line Remote Error Indication) was originally referred to as Line FEBE.
- This entry, which previously appeared as M0, applies for signals at the OC-1 rate.
- This entry is applicable for OC-768 signals.

3.3.2.1 Section Overhead

This section defines each of the SOH bytes.

Framing (A1 and A2) – For signals at rates from STS-1 through STS-192, two bytes are allocated in each STS-1 for framing. At the STS-768 rate, the A1 byte positions in the last 64 STS-1s (in order of transmission) and the A2 byte positions in the first 64 STS-1s (also in order of transmission) are allocated for framing. The corresponding byte positions in the remaining 704 STS-1s are reserved for future standardization, and must be set to values as indicated in **R3-20 [1144]**.⁶

R3-19 [16v2] The A1 byte shall be set to ‘1111 0110’ and the A2 byte shall be set to ‘0010 1000’ in all STS-1s within an STS-N where N is less than or equal to 192. For N equal to 768, each of the last 64 STS-1s in order of transmission shall contain an A1 byte set to ‘1111 0110’, and each of the first 64 STS-1s in order of transmission shall contain an A2 byte set to ‘0010 1000’.

R3-20 [1144] In an STS-768, the A1 byte positions in the first 704 STS-1s and the A2 byte positions in the last 704 STS-1s (both in order of transmission) shall be set to values that, after scrambling (see [Section 5.1.3](#)):

- Will result in approximately equal numbers of “zeros” and “ones”
- Will not mimic a pattern that might be used for framing (e.g., one or more A1 byte patterns followed by one or more A2 byte patterns)
- Will provide sufficient transitions for proper operation of the clock recovery process at the receiving NE.

[Section 5.5](#) contains the framing criteria for SONET NEs.

Section Trace (J0)/Section Growth (Z0) – One byte in the first STS-1 of an STS-N is allocated for a Section trace function, while the corresponding byte position in each of the remaining N–1 STS-1s is allocated for Section-level future growth.

O3-21 [17v2] STE that supports line-side signals should support a feature that allows the contents of the J0 bytes in those signals to be provisioned (on a per-SONET-signal basis) by the user.

O3-22 [1145] If a SONET NE supports a feature that allows the contents of the J0 byte to be provisioned, then the following apply:

- The feature should allow the user to enter the contents of the Section trace such that it meets the specifications in ANSI T1.269, *Information Interchange - Structure and Representation of Trace Message Formats for the North American Telecommunications System*, related to 16-byte trace messages.

6. More specific criteria regarding the values to be transmitted in these 1408 bytes and the 767 Z0 bytes (see **R3-25 [1146]**) in an STS-768 may be provided in a future issue of this document. In addition, **O3-2 [5]** is not considered applicable to any of these bytes.

- The feature should allow the user to enter the contents of the Section trace such that it consists of any single value from 00 through FF (hex).

The format for a 16-byte trace message is described in ANSI T1.269 [or ITU-T G.831, *Management capabilities of transport networks based on the Synchronous Digital Hierarchy (SDH)*], as a header byte consisting of a frame start marker (i.e., a '1' in bit 1 and the results of a Cyclic Redundancy Check-7 (CRC-7) calculation over the previous frame in bits 2 through 8, followed by 15 ASCII [or ITU-T Rec. T.50, *International Reference Alphabet (IRA) (Formerly International Alphabet No. 5 or IA5) – Information technology – 7-bit coded character set for information interchange*] characters beginning with either the country code as defined in ITU-T E.164, *The international public telecommunication numbering plan*, or the alphabetic character country code as defined in ISO 3166-1:1997, *Codes for the representation of names of countries and their subdivisions – Part 1: Country codes*. Additional criteria related to the use of the Section trace appears in [Section 6.2.3.2.3.A](#). Also, the ability to access the J0 byte or transmit values other than that indicated in **R3-23 [18v3]** is not required for STE that only supports drop-side signals.

For $N \leq 48$, both the J0 and Z0 bytes were originally defined to be STS-1 Identifier (C1) bytes, and were required to be set to indicate their order of appearance in the STS-N frame. To promote compatibility between (older) equipment designed to meet that earlier requirement and current equipment, the following requirement applies. In addition, **R3-24 [1102]** is intended to (among other things) assure sufficient transitions for proper operation of the clock recovery function at equipment receiving an OC-192 signal. Finally, for $N = 768$, the Z0 bytes are scrambled (see [Section 5.1.3](#)), and therefore **R3-25 [1146]** applies.

R3-23 [18v3] Unless it is being used for a defined purpose (e.g., to carry a Section trace message) the J0 byte and, in an STS-48 or lower bit-rate signal, each Z0 byte shall be set to a binary number corresponding to its order of appearance in the STS-N frame (i.e., the J0 byte shall be set to '00000001', the first Z0 byte shall be set to '00000010', the second Z0 byte to '00000011', etc.).

R3-24 [1102] Unless it is being used for a (future) defined purpose, each Z0 byte in an STS-192 bit-rate signal shall be set to the fixed pattern '11001100'.

R3-25 [1146] With the exception of those being used for (future) defined purposes, the Z0 bytes in an STS-768 bit-rate signal shall be set to values that, after scrambling

- Will result in approximately equal numbers of “zeros” and “ones”
- Will not mimic a pattern that might be used for framing (e.g., one or more A1 byte patterns followed by one or more A2 byte patterns)
- Will provide sufficient transitions for proper operation of the clock recovery process at the receiving NE.

The preceding requirements are applicable for both line-side and drop-side signals. Also, since no standard use has been defined for the Z0 bytes (or for the former C1 bytes) received by an NE, these bytes are currently considered unused at the receiving STE (see [Section 3.2](#)).

Section BIP-8 (B1) – The B1 byte is located in the first STS-1 of an STS-N, and is defined for a Section error monitoring function in line-side signals. The value contained in the B1 byte in a drop-side signal is undefined, so the criteria in [Section 3.2](#) are applicable.⁷ Except for those used in the optional FEC function (see [Section 3.3.2.2](#)), the corresponding byte locations in the second through Nth STS-1s of both line-side and drop-side signals are also currently undefined.

R3-26 [19] The B1 byte in a line-side signal shall carry a BIP-8 code, using even parity. The Section BIP-8 shall be calculated over all bits of the previous STS-N frame after scrambling and placed in the B1 byte of the current STS-N frame before scrambling.

Orderwire (E1) – The E1 byte is located in the first STS-1 of an STS-N, and is used for an LOW channel to provide voice communication between regenerators, hubs and remote terminal locations. Except for those used in the optional FEC function (see [Section 3.3.2.2](#)), the corresponding byte locations in the second through Nth STS-1s are currently undefined. [Section 5.2.2](#) contains the orderwire criteria.

Section User Channel (F1) – The F1 byte is located in the first STS-1 of an STS-N, and is available for use by the network provider. The corresponding byte locations in the second through Nth STS-1s are currently undefined. [Section 5.2.3](#) contains the Section user channel criteria.

Section Data Communications Channel (D1, D2 and D3) – The D1, D2 and D3 bytes are located in the first STS-1 of an STS-N, and are used for Section data communications. Except for those used in the optional FEC function (see [Section 3.3.2.2](#)), the corresponding byte locations in the second through Nth STS-1s are currently undefined.

These three bytes are considered as one 192-kb/s, message-based channel for alarms, maintenance, control, monitoring, administration and other communication needs between STE. This channel is used for internally generated, externally generated and supplier-specific messages. [Section 8](#) contains the data communications channel criteria.

3.3.2.2 FEC Sub-Layer

As of the publication of ANSI T1.105.08 in 2001 and ITU-T Rec. G.707 in 2003, a number of OC-768 and previously undefined OC-48 and OC-192 SOH and LOH bytes were reserved for an FEC function. That function is performed at an FEC sub-layer that resides between the Section and Line layers, and uses bytes borrowed from the SOH and LOH to provide error correction on the SONET Line layer signal (i.e., the LOH and the STS-48, STS-192 or STS-768 envelope capacity) plus 1, 4 or 16 bytes of the SONET SOH for OC-48, OC-192 or OC-768 signals, respectively. (Note that for FEC calculation purposes, the remaining SOH bytes are set to all-zeros at both the encoder and the decoder.)

7. It is also acceptable for the B1 byte in a drop-side signal to carry the BIP-8 code as described for line-side signals. However, STE cannot assume the BIP-8 code will be present in the received drop-side signal, and therefore it must be capable of ignoring the value in that byte.

In general, a number of revisions and new detailed information and criteria would be needed in this document in order for it to fully cover the FEC function. While those changes and additions have not been included in this issue of GR-253-CORE, it is possible that they will be made in a future issue. Until they are included, the following criteria apply.

CR3-27 [1103v2] A SONET NE may be required to support an FEC function at its OC-48, OC-192 and/or OC-768 interfaces (if provided).

R3-28 [1104v2] If a SONET NE supports an FEC function, that function shall be supported according to the specifications contained in ANSI T1.105.08 and ITU-T Rec. G.707.

Note that ANSI T1.105.08 and ITU-T Rec. G.707 covers such topics as the particular bytes and error correction codes used, code block definitions, compensation of the B2 bytes, encoder and decoder locations, delay characteristics, possible functionalities of (intermediate) SONET regenerators, FEC operational states, and transitions between those operational states. Also note that at the time this document was being prepared, it appeared that the only difference between the FEC-related specifications in T1.105.08 and Rec. G.707 was that T1.105.08 did not address the OC-768 [or Synchronous Transport Module-256 (STM-256)] bit rate.

3.3.2.3 Line Overhead

This section defines each of the LOH bytes.

STS Payload Pointer (H1 and H2) – Two bytes are allocated to a pointer that indicates the offset in bytes between the pointer and the first byte of the STS SPE. The pointer bytes are used in all STS-1s within an STS-N to align the STS-1 transport overheads in the STS-N, and to perform frequency justification (see [Section 3.5](#)).

These bytes are also used to indicate concatenation, and to detect STS Path AIS (AIS-P).

Pointer Action Byte (H3) – The pointer action byte is allocated for SPE frequency justification purposes. The H3 byte is used in all STS-1s within an STS-N to carry the extra SPE byte in the event of a negative pointer adjustment. The value contained in this byte when it is not used to carry the SPE byte is undefined.

Line BIP-8 (B2) – One byte is allocated in each STS-1 for a Line error monitoring function. The N Line BIP-8 bytes in an STS-N electrical or OC-N signal are intended to form a single error monitoring facility capable of measuring BERs up to 10^{-3} , independent of the value of N.

R3-29 [20] The B2 byte shall be provided in all STS-1s within an STS-N to carry a Line BIP-8 code, using even parity. The Line BIP-8 shall be calculated over all bits of the LOH and the envelope capacity of the previous STS-1 frame before scrambling, and placed in the B2 byte of the current STS-1 frame before scrambling.

APS Channel (K1 and K2) – The K1 and K2 bytes are located in the first STS-1 of an STS-N, and are used on the protection line for APS signaling between LTE that uses Line-level protection switching [e.g., in systems using linear APS, or in Bidirectional Line Switched Rings (BLSRs)]. Except for those used in the optional FEC function (see [Section 3.3.2.2](#)), the corresponding byte locations in the second through Nth STS-1s are currently undefined.

The K2 byte is also used to detect Line AIS (AIS-L) and Line Remote Defect Indication (RDI-L) signals (see [Sections 6.2.1.2.1](#) and [6.2.1.3.1](#)).

Line Data Communications Channel (D4 through D12) – The D4 through D12 bytes are located in the first STS-1 of an STS-N, and are used for Line data communication. Except for those used in the optional FEC function (see [Section 3.3.2.2](#)) or extended Line DCC function (see below), the corresponding byte locations in the second through Nth STS-1s are currently undefined.

These nine bytes are considered as one 576-kb/s, message-based channel for alarms, maintenance, control, monitoring, administration and other communication needs. This channel is available for internally generated, externally generated and supplier-specific messages. [Section 8](#) contains the data communications channel criteria.

Extended Line Data Communications Channel (D13 through D156) – The D13 through D156 bytes are defined only for OC-768 signals and are located in the ninth through fifty-sixth STS-1s (in order of transmission) in the D4 (D13 through D60), D7 (D61 through D108) and D10 (D109 through D156) byte positions illustrated in [Figure 3-26](#).⁸ These 144 bytes are considered as one 9216-kb/s, message-based channel for alarms, maintenance, control, monitoring, administration and other Line data communication needs. This channel is available for internally generated, externally generated and supplier-specific messages. [Section 8](#) contains the data communications channel criteria.

Synchronization Status (S1) – The S1 byte is located in the first STS-1 of an STS-N, and bits 5 through 8 of that byte are allocated to convey the synchronization status of the NE. [Section 5.4.2](#) contains the synchronization status message criteria. Bits 1 through 4 of the S1 byte are currently undefined.

Growth (Z1) – The Z1 bytes are located in the second through Nth STS-1s of an STS-N where $3 \leq N \leq 48$, or STS-1 numbers “1,2” through “16,3” of an STS-192 or STS-768,⁹ and are allocated for future growth. Except for those used in the optional FEC function (see [Section 3.3.2.2](#)), the use of these bytes is currently undefined.

8. Note that this description is consistent with that provided in ITU-T Rec. G.707. On the other hand, the version of ANSI T1.105 that was available when this document was being prepared (i.e., T1.105-2001) indicated that these bytes are located in the second through forty-ninth STS-1s (in order of transmission) rather than the ninth through fifty-sixth STS-1s. If in the future Rec. G.707 is changed to be consistent with T1.105, then the description provided here should be considered preempted by that provided in (both of) the standards.

9. These STS-1s can also be identified as STS-1 numbers 2 through 48 (see [Section 5.1.1](#)). In order of appearance in the byte-interleaved OC-192 signal, these are the 2nd through 16th, 65th through 80th and 129th through 144th STS-1s. In order of appearance in the byte-interleaved OC-768 signal, they are the 2nd through 16th, 257th through 272nd and 513th through 528th STS-1s.

Note that an OC-1 or STS-1 electrical signal does not contain a Z1 byte. Also note that the corresponding byte locations in STS-1 numbers “17,1” through “64,3” or “256,3” (or 49 through 192 or 768) of an STS-192 or STS-768 are currently undefined except for possible use in the FEC function.

Line Remote Error Indication (M1 or M0&M1) – As described below, either one byte (M1) or two bytes (M0 and M1) are allocated for an REI-L function. That function conveys the error count detected by LTE (using the Line BIP code) back to its peer LTE. (Also see [Section 5.2.4](#).)

For OC-1 and STS-1 electrical signals, bits 5 through 8 of the M1 byte (located in the only available STS-1) are used for the REI-L function,¹⁰ while bits 1 through 4 of that byte are currently undefined.

R3-30 [21v2] LTE terminating an OC-1 or STS-1 electrical signal shall set bits 5 through 8 of the M1 byte to indicate (to the upstream LTE) the count of the interleaved-bit block errors that it has detected based on the Line BIP-8 (B2) byte. The error count shall be a binary number from zero (i.e., ‘0000’) to 8 (i.e., ‘1000’). The remaining seven values represented by the four REI-L bits (i.e., ‘1001’ through ‘1111’) shall not be transmitted, and shall be interpreted by receiving LTE as zero errors.

For STS-3 electrical signals and OC-N signals where $3 \leq N \leq 192$, the M1 byte is located in the third STS-1¹¹ in order of appearance in the byte-interleaved STS-N electrical or OC-N signal.

R3-31 [22v2] LTE terminating an OC-N or STS-N electrical signal where $3 \leq N \leq 192$ shall set the M1 byte to indicate (to the upstream LTE) the count of the interleaved-bit block errors that it has detected using the Line BIP-8 (B2) bytes. For values of N less than 48, the error count shall be a binary number from zero to 8N. The remaining

10. Note that through Issue 3 of this document, as well as in the 2001 version of ANSI T1.105, the byte that is used for the REI-L function in OC-1 and STS-1 electrical signals was labeled M0 instead of M1. However with the addition of the two-byte REI-L function at the OC-768 bit rate (in which M0 is used to identify the byte containing the most significant bits of the two-byte code and M1 is used to identify the byte containing the least significant bits of that code), that designation was considered too confusing. In addition, the designations used here are consistent with those that appear in ITU-T Rec. G.707.

11. The “third STS-1” is defined as the third STS-1 in order of their appearance in the byte-interleaved STS-N electrical or OC-N signal. Using the two-level STS numbering scheme discussed in [Figures 5-1](#) and [5-2](#), the third STS-1 in an OC-12 or higher rate signal would be labeled “3,1”.

It is important to recognize that the numbering of the transport overhead bytes as they appear in an STS-N electrical or OC-N signal can be separated from the numbering of the STS-1 and STS-M inputs to the byte-interleaver (see [Section 5.1.1](#)). Except for the pointer bytes, when the NE adds (or overwrites) the LOH to the output of the byte-interleaver, it does so independently of the particular mix of input STS-1s and STS-Ms to the byte-interleaver. For example, if the first input to the byte-interleaver shown in [Figure 5-2](#) was “STS-12 Number 1,1” (instead of “STS-1 Number 1,1”) then there would be no inputs shown with numbers of “1,2”, “1,3”, “2,1”, ..., “4,3” (including “3,1”). However, the M1 byte would still appear in a particular transport overhead byte position, and that byte position can be said to be contained in STS-1 “3,1”.

255 – (8×N) possible values represented by the eight REI-L bits shall not be transmitted, and shall be interpreted by the receiving LTE as zero errors. For N equal to 48 or 192, the count shall be truncated at 255.

For OC-768 signals, the M0 byte is located in the second STS-1 and the M1 byte is located in the third STS-1 (both in order of appearance in the byte-interleaved signal).

R3-32 [1147] LTE terminating an OC-768 signal shall set the M0 and M1 bytes to indicate (to the upstream LTE) the count of the interleaved-bit block errors that it has detected using the Line BIP-8 (B2) bytes. The error count shall be a binary number from zero to 6144, with bit 1 of the M0 byte being the most significant bit and bit 8 of the M1 byte being the least significant. The remaining 65535 – 6144 = 59391 possible values represented by the 16 REI-L bits shall not be transmitted, and shall be interpreted by the receiving LTE as zero errors.

Note that consistent with ANSI T1.105, **R3-30 [21v2]** indicates that LTE receiving an OC-1 or STS-1 electrical signal must consider only bits 5 through 8 of the M1 byte, while **R3-31 [22v2]** and **R3-32 [1147]** indicate that LTE receiving higher bit rate signals must consider the entire M1 and (if applicable) M0 bytes. ITU-T Rec. G.707, on the other hand, specifies that for STM-0, STM-1 and STM-4 signals (i.e., signals at bit rates from OC-1 through OC-12) the receiving NE should consider bits 2 through 8 of the M1 byte. For an STM-0 signal this means (for example) that a received value of '0100 0001' would be interpreted as an invalid code and therefore zeros errors, rather than one error as required by **R3-30 [21v2]**. Conversely, for an STM-1 or STM-4 signal, a received value of '1000 0001' (for example) would be interpreted as a valid code indicating one error, rather than an invalid code and therefore zeros errors as required by **R3-31 [22v2]**.

Also consistent with T1.105, only a single-byte REI-L function is defined above for use at the OC-192 bit rate. ITU-T Rec. G.707, on the other hand, defines a two-byte REI-L function for all STM-N signals with $N \geq 64$ (i.e., signals at rates of OC-192 and above). While a SONET NE may certainly support SDH STM-N interfaces that meet the Rec. G.707 specifications, it needs to meet the requirements shown above at its SONET interfaces.

Growth (Z2) – The Z2 bytes are located in the first and second STS-1s of an STS-3, the first, second and fourth through Nth STS-1s of an STS-N where $12 \leq N \leq 48$, STS-1 numbers "1,1" through "2,3" and "3,2" through "16,3" of an STS-192, and STS-1 numbers "1,1" through "1,3", "2,2", "2,3", and "3,2" through "16,3" of an STS-768.¹² These bytes are allocated for future growth, and their use is currently undefined.

12. In an STS-192, these STS-1s can also be identified as STS-1 numbers 1 through 6 and 8 through 48, while in an STS-768 they can be identified as STS-1 numbers 1, 2, 3, 5, 6 and 8 through 48 (see [Section 5.1.1](#)). In order of appearance in the byte-interleaved OC-192 signal, these are the 1st, 2nd, 4th through 16th, 65th through 80th and 129th through 144th STS-1s. In order of appearance in the byte-interleaved OC-768 signal, they are the 1st, 4th through 16th, 257th through 272nd and 513th through 528th STS-1s. In addition, note that ANSI T1.105 indicates (in separate definitions) that the byte in this position in the second STS-1 (in order of appearance) in a byte-interleaved OC-768 signal is both the M0 byte and a Z2 byte.

Note that an OC-1 or STS-1 electrical signal does not contain a Z2 byte. Also note that the corresponding byte locations in STS-1 numbers “17,1” through “64,3” or “256,3” (or 49 through 192 or 768) of an STS-192 or STS-768 are currently undefined.

Orderwire (E2) – The E2 byte is located in the first STS-1 of an STS-N, and is used for an Express Orderwire (EOW) channel between Line entities. The corresponding byte locations in the second through Nth STS-1s are currently undefined.

[Section 5.2.2](#) contains the orderwire criteria.

3.3.2.4 STS Path Overhead

STS POH is assigned to each STS SPE for functions necessary in transporting its payload. The STS POH supports the following classes of functions:

- A. Payload independent functions with standard format and coding – all payloads require these functions
- B. Mapping dependent functions with standard format and coding that are specific to the type of payload – these functions are needed for one or more types of payload, but not all payloads
- C. Application-specific functions – appropriate GRs, TRs or standards documents (e.g., ANSI T1.105.05 for applications using the STS tandem connection sub-layer) specify the format and coding for these functions, which may share the same overhead capacity.

STS POH capacity not yet assigned to Class A, B or C functions may be defined in the future for supporting any of those classes, with Class A having priority. Also, this classification scheme does not preclude the allocation of other overhead bits or bytes within the STS payload capacity for specific mappings (e.g., the stuff control bits for the asynchronous DS3 mapping).

Class A Functions

STS Path Trace (J1) – This byte is used to transport a repetitive 64-byte message so that the receiving STS PTE can verify its continued connection to the intended transmitting STS PTE.

[Sections 6.2.1.1.9](#) and [6.2.3.2.3.A](#) contain the criteria related to loading and detecting the STS path trace message, along with information about possible message formats.

R3-33 [23] If no message has been loaded by the user for transmission in the J1 byte, then that byte shall be set to all zeros (i.e., to ASCII NULL characters).

In addition, it should be noted that in SDH this byte is used to transport a repetitive 16-byte message with a format as described in [Section 3.3.2.1](#) (for use in the J0 byte). Therefore, SONET NEs that are intended to interoperate with SDH equipment may need to also support the generation and processing of STS path trace messages in that format.

STS Path BIP-8 (B3) – The B3 byte is allocated for an STS path error monitoring function.

- R3-34 [24]** The B3 byte shall carry a BIP-8 code, using even parity. The STS path BIP-8 shall be calculated over all bits (783 bytes for an STS-1 SPE or $N \times 783$ bytes for an STS-Nc SPE, regardless of any pointer adjustments) of the previous STS SPE before scrambling, and placed in the B3 byte of the current STS SPE before scrambling.

STS Path Signal Label (C2) – The C2 byte is allocated to indicate the content of the STS SPE, including the status of the mapped payloads. Of the 256 possible binary values [00 to FF (hex)], only the codes defined in Tables 3-2 and 3-3 have been assigned for the mappings described in Section 3.4 or other documents (as noted). The codes in Table 3-2 are assigned for use by all STS PTE, and the code generated (under normal conditions) by a particular STS PTE depends on its provisioned (or only supported) functionality. The additional codes in Table 3-3 are assigned for use by STS PTE that support the STS Path Payload Defect Indication (PDI-P) feature, and are generated automatically based on the status of the mapped payloads. The remaining codes are reserved to be assigned as needed for future STS payload-specific mappings.

PDI-P is an application-specific feature that uses the STS path signal label to indicate to downstream equipment that there is a defect in one or more directly mapped, embedded payloads in the STS SPE. For VT-structured STS-1 SPEs, 28 codes are defined to indicate that 1 through 28 of the embedded VTn paths have defects. In this scheme a VT1.5 payload defect is assigned the same priority as a VT2, VT3 or VT6 payload defect. For example, an STS-1 SPE with two VT6 payload defects would be assigned the code 'E2,' which is the same code as would be used for an STS-1 SPE with two VT1.5 payload defects. For non-VT-structured STS-1 or STS-Nc SPEs, one code is defined to indicate that the mapped payload has a defect. See Section 6.2.1.4.1 for additional criteria related to PDI-P.

- R3-35 [25v2]** Except for the case covered in **R3-37 [1148]**, STS PTE supporting an STS path connection that is “provisioned” shall generate a valid non-zero STS path signal label. If the content of the STS SPE is one of the specific possibilities listed in Table 3-2 (including the case of a mapping that is under development), then the corresponding code from Table 3-2 or (if PDI-P is supported and one or more payload defects are present) Table 3-3 shall be used. If the content is not specifically listed and does not need to be identified via an STS path signal label, then the code for “Equipped – Nonspecific Payload” shall be used.
- R3-36 [26]** For STS path connections that are not equipped, or that are equipped but not provisioned, the NE (e.g., the NE’s LTE) shall generate all-zeros STS SPEs with “valid” STS payload pointers.

In the above requirements, “provisioned” means that the STS PTE has been configured for a mapping (or only supports one mapping) and has been assigned a timeslot (or is hardwired to a specific timeslot) in a SONET signal. Unless both of these conditions are met, the connection is not considered provisioned. In addition,

a supplier may choose to not consider a connection to be provisioned unless the STS PTE has been assigned one or more signals (e.g., one or more VTs) to map or multiplex into the STS SPE that it is originating.

Note that Issue 1 of this document required that a connection not be considered provisioned unless the STS PTE had been assigned a signal (rather than leaving it as a supplier option). That meant (for example) that STS PTE terminating a VT-structured STS-1 SPE was not considered provisioned unless it had been assigned one or more specific VTs to place into the STS-1 SPE that it was originating. However, in some applications it may be useful to confirm STS PTE to STS PTE connectivity before they have been assigned specific signals to map or multiplex. NEs that provide that capability [e.g., by sending non-zero STS signal labels or STS path trace messages (along with the appropriate STS path BIP-8) before being assigned a signal] should not be considered nonconforming, and therefore the definition of provisioned was revised.

Also note that the “valid” pointers referred to in **R3-36 [26]** may contain any value that can be interpreted correctly by downstream STS pointer processors without causing an STS Path Loss of Pointer (LOP-P, see [Section 6.2.1.1.3](#)), and that an all-zeros STS SPE results in a valid STS path BIP-8 code and a signal label of ‘00’ (hex), indicating “Unequipped”.

R3-37 [1148] STS PTE or supervisory generation equipment supporting an STS path connection that is provisioned to carry an STS Supervisory-unequipped (TEST-P) signal shall generate a valid STS path trace (J1), the Unequipped STS path signal label (C2 = ‘00’), a valid STS path BIP-8 (B3), and a valid STS path status (G1). In addition, it shall set the STS tandem connection byte (N1) to all zeros.

The contents of the remaining STS POH bytes and the STS payload capacity in a TEST-P signal are currently undefined (see [Section 3.2](#)). Also, ANSI T1.105 discusses possible applications of the TEST-P signal, and information and criteria related to such applications may be included in a future issue of this document.

Table 3-2 STS Path Signal Label Assignments

Code (Hex)	Content of the STS SPE	Code (Hex)	Content of the STS SPE
00	Unequipped or Supervisory Unequipped	17	Mapping for SDL with Self-Synchronizing Scrambler ^c
01	Equipped – Nonspecific Payload	18	Mapping for HDLC/LAPS Framed Signals
02	VT-Structured STS-1 SPE ^a		
03	Locked VT Mode ^a	19	Mapping for SDL with Set-Reset Scrambler ^c
04	Asynchronous Mapping for DS3		
05	Mapping under development ^b	20	Asynchronous mapping of ODUk (k=1, 2) into VC-4-Xv (X=17, 68) ^d
12	Asynchronous Mapping for DS4NA		
13	Mapping for ATM	1A	Mapping for 10 Gb/s Ethernet ^e
14	Mapping for DQDB	1B	GFP Mapping
15	Asynchronous Mapping for FDDI	1C	Mapping for 10 Gb/s Fibre Channel ^f
16	Mapping for HDLC/PPP Framed Signals (a.k.a., the HDLC-Over-SONET Mapping)	CF	Reserved ^g
		FE	O.181 Test Signal Mapping ^h
		FF	STS SPE AIS ⁱ

Notes:

- a In early SONET standards and criteria documents, two modes were defined for VT-structured STS-1 SPEs. These were the Floating mode, which was assigned the code ‘02,’ and the Locked mode, which was assigned ‘03.’ The Locked VT Mode was subsequently removed from the SONET standards and criteria. However, for backward compatibility between future mappings and equipment that supports the Locked VT Mode, the signal label that was assigned to that mode remains defined.
- b This code is specified to be used for developmental or experimental activities (i.e., for cases where the payload mapping is not one of the specific mappings for which a different code has been assigned). Note that prior to the definition of this code (in the 2001 version of ANSI T1.105 and Issue 4 of this document) the “Equipped – Nonspecific Payload” code was often used for such activities; however, since ‘01’ is required to be considered “Matched” by all STS PTE (see [Section 6.2.1.1](#)), the use of that code did not provide the necessary isolation between developmental or experimental activities and the rest of the SONET network. Also note that in cases where the developmental or experimental activities result in the specification of a new mapping that is assigned its own payload-specific code, no compatibility will exist between equipment that uses this code and subsequent equipment that uses the new code. If such compatibility is desired, it will be necessary for the equipment that uses this code to be able to be reconfigured to utilize the new code.
- c This mapping is considered proprietary.
- d Specifications related to this mapping appear in ITU-T Rec. G.707, and may be added to a future issue of this document.

- e This code/mapping assignment is specified in ITU-T Rec. G.707. It is expected to be used by equipment that supports the 10 Gb/s Ethernet Wide Area Network (WAN) interface defined in IEEE 802.3ae, *Standard for Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications-Media Access Control (MAC) Parameters, Physical Layer and Management Parameters for 10 Gb/s Operation*. That interface is very similar to an STM-64 with a VC-4-64c (i.e., an OC-192 signal carrying an STS-192c SPE), with the primary differences being certain limitations on the use of overhead bytes and different clock accuracy specifications.
- f This mapping is under study, and the signal label has been provisionally allocated.
- g This code was used for a High-level Data Link Control/Point-to-Point Protocol (HDLC/PPP) payload mapping that is now obsolete.
- h This code/mapping assignment is specified in ITU-T Rec. G.707 for use in SDH networks (to indicate any mapping that is defined in ITU-T Rec. O.181, *Equipment to assess error performance on STM-N interfaces*, and does not correspond to a mapping defined in Rec. G.707).
- i This code assignment is specified in ANSI T1.105 and ITU-T Rec. G.707 to indicate STS SPE AIS (a.k.a., VC-AIS). It is generated by an STS tandem connection maintenance source if no valid incoming signal is available and a replacement signal is being generated.

Table 3-3 STS Path Signal Label Assignments for Signals with Payload Defects

Code (Hex)	Content of the STS SPE	Code (Hex)	Content of the STS SPE
E1	VT-structured STS-1 SPE with 1 VTn Payload Defect (STS-1 w/1 VTn PD)	F0	STS-1 w/16 VTn PDs
		F1	STS-1 w/17 VTn PDs
E2	STS-1 w/2 VTn PDs	F2	STS-1 w/18 VTn PDs
E3	STS-1 w/3 VTn PDs	F3	STS-1 w/19 VTn PDs
E4	STS-1 w/4 VTn PDs	F4	STS-1 w/20 VTn PDs
E5	STS-1 w/5 VTn PDs	F5	STS-1 w/21 VTn PDs
E6	STS-1 w/6 VTn PDs	F6	STS-1 w/22 VTn PDs
E7	STS-1 w/7 VTn PDs	F7	STS-1 w/23 VTn PDs
E8	STS-1 w/8 VTn PDs	F8	STS-1 w/24 VTn PDs
E9	STS-1 w/9 VTn PDs	F9	STS-1 w/25 VTn PDs
EA	STS-1 w/10 VTn PDs	FA	STS-1 w/26 VTn PDs
EB	STS-1 w/11 VTn PDs	FB	STS-1 w/27 VTn PDs
EC	STS-1 w/12 VTn PDs	FC	VT-structured STS-1 SPE with 28 VT1.5 Payload Defects, or a non-VT- structured STS-1 or STS-Nc SPE with a Payload Defect
ED	STS-1 w/13 VTn PDs		
EE	STS-1 w/14 VTn PDs		
EF	STS-1 w/15 VTn PDs		

Path Status (G1) – The G1 byte is allocated to convey the STS path terminating status and performance back to the originating STS PTE. This feature permits the status and performance of the complete duplex path to be monitored at either end, or at any point along that path. [Figure 3-27](#) illustrates the bit assignments in the G1 byte. Bits 1 through 4 are allocated for an STS Path REI function (REI-P, formerly referred to as STS path FEBE). (Also see [Section 5.2.4](#).)


R3-38 [27v2] Except for when it has detected an LOP-P or AIS-P defect, STS PTE shall set bits 1 through 4 of the G1 byte to indicate (to the upstream STS PTE) the count of interleaved-bit block errors that it has detected based on the STS path BIP-8 byte (B3). The error count shall be a binary number from zero (i.e., ‘0000’) to 8 (i.e., ‘1000’). The remaining seven values represented by the four REI-P bits (i.e., ‘1001’ through ‘1111’) shall not be transmitted, and shall be interpreted by receiving STS PTE as zero errors.

Note that when STS PTE has detected an LOP-P or AIS-P defect, it either cannot locate the B3 byte or that byte does not exist.¹³ In either case, an accurate error count cannot be provided in the REI-P bits, and therefore those bits are considered undefined. While the criteria in [Section 3.2](#) indicate that undefined bits should be set to all zeros, it is also considered acceptable for G1 bits 1 through 4 to be set to indicate an “apparent” error count in this situation [e.g., a count based on processing of the STS path signal using the most recent valid pointer value (and thus the last known B3 byte location)]. Also note that while **R3-38 [27v2]** is consistent with the specifications in ANSI T1.105, it differs from those in ITU-T Rec. G.707. In particular, Rec. G.707 indicates that the REI-P bits are also undefined if the STS PTE has detected an STS Path Unequipped or Trace Identifier Mismatch (UNEQ-P or TIM-P) defect.

Bits 5, 6 and 7 of the G1 byte are allocated for an STS Path RDI (RDI-P) signal, while bit 8 of that byte is currently undefined. [Section 6.2.1.3.2](#) contains the criteria related to the generation and detection of RDI-P.

13. The same cannot be said for the case of LTE that is monitoring for Line BIP-8 (B2) errors and detects an AIS-L defect. That is, although they will have been overwritten with all ones, the B2 bytes are still present and can be located based on their fixed offsets from the framing (A1 and A2) bytes. Therefore, the REI-L generation requirements in [Section 3.3.2.3](#) do not contain exception statements equivalent to the statement in **R3-38 [27v2]**.

Figure 3-27 STS Path Status Byte (G1)



REI-P Coding:

0	0	0	0	0	0	Errors
0	0	0	1	0	1	Error
0	0	1	0	0	2	Errors
0	0	1	1	0	3	Errors
0	1	0	0	0	4	Errors
0	1	0	1	0	5	Errors
0	1	1	0	0	6	Errors
0	1	1	1	0	7	Errors
1	0	0	0	0	8	Errors
1	0	0	1	0		
1	0	1	0	0		
1	0	1	1	0		
1	1	0	0	0	}	0 Errors
1	1	0	1	0		
1	1	1	0	0		
1	1	1	1	0		
1	1	1	1	1		

Class B Functions

Indicator (H4) – The H4 byte is allocated for use as a mapping-specific indicator byte. Currently, it is used only in VT-structured STS-1 SPEs, the Distributed Queue Dual Bus (DQDB) mapping, and virtually concatenated STS-1 or STS-3c SPEs. For VT-structured STS-1 SPEs, the byte is used as a multiframe indicator and the applicable criteria appear in [Section 3.4.1](#). In the DQDB mapping it is used to carry the DQDB Link Status Signal, and to indicate the offset to the boundary of the next DQDB slot. See [Section 3.4.2.1.4](#) for details. Finally, in STS virtual concatenation applications the byte is used as a sequence and multiframe indicator, and the applicable criteria are provided in [Section 5.2.5](#).

For mappings where a standard use of this byte has not been defined, it is considered undefined.

Class C Functions

Path User Channel (F2) – The F2 byte is allocated for user communication purposes between STS path terminating NEs. For applications where this byte is not used, it is considered undefined. [Section 5.2.3](#) contains the criteria for the path user channel.

The F2 byte is also used in the DQDB mapping to carry DQDB Layer Management information (a Class B function). See [Section 3.4.2.1.4](#) for details.

Tandem Connection (N1) – The N1 byte (formerly referred to as the Z5 byte) is allocated for functions related to tandem connections, including the STS tandem connection incoming error count in bits 1 through 4, and STS tandem connection maintenance or the STS tandem connection data link in bits 5 through 8. Refer to ANSI T1.105 and T1.105.05, and ITU-T Rec. G.707 for details.

Other


STS Path Growth (Z3, Z4) – The Z3 and Z4 bytes (which are referred to in SDH as the F3 and K3 bytes) are allocated for future growth, and the use of these bytes is currently undefined for most mappings and applications. In the DQDB mapping the Z3 byte is used to carry DQDB Layer Management information (a Class B function). See Section 3.4.2.1.4 for details.

3.3.3 VT Path Overhead

Four bytes (V5, J2, Z6 and Z7) are allocated for VT POH. The first byte of a VT SPE (i.e., the byte in the location pointed to by the VT payload pointer) is the V5 byte, while the J2, Z6 and Z7 bytes occupy the corresponding locations in the subsequent 125-µs frames of the VT superframe (as shown in Figure 3-23).

V5 – The V5 byte provides the same functions for VT paths that the B3, C2 and G1 bytes provide for STS paths; namely error checking, signal label and path status. The bit assignments for the V5 byte are illustrated in Figure 3-28.

Figure 3-28 VT Path Overhead Byte (V5)



REI-V Coding

- 0 0 Errors
- 1 1 or 2 Errors

Bits 1 and 2 of the V5 byte are allocated for error performance monitoring. A BIP-2 scheme is defined as follows:

R3-39 [28] Bit 1 of the V5 byte shall be set so that the parity of all of the odd-numbered bits (i.e., bits 1, 3, 5 and 7) in all bytes in the previous VT SPE is even. Bit 2 shall be set so that the parity of all of the even-numbered bits (2, 4, 6 and 8) in all bytes in the previous VT SPE is even.

Bit 3 of the V5 byte is allocated for a VT Path REI function (REI-V, formerly referred to as VT path FEBE) to convey the VT path terminating performance back to an originating VT PTE. (Also see [Section 5.2.4](#).)

- R3-40 [29v2]** Except for when it has detected a VT Path LOP or AIS (LOP-V or AIS-V) defect, VT PTE shall set bit 3 of the V5 byte to '1' if one or more errors were detected using the BIP-2. It shall set bit 3 to '0' if zero errors were detected.

Similar to the case discussed in [Section 3.3.2.4](#) regarding REI-P, when VT PTE has detected an LOP-V or AIS-V defect, it either cannot locate the V5 byte or that byte does not exist. In either case, an accurate error indication cannot be provided in the REI-V bit, and therefore that bit is considered undefined. While the criteria in [Section 3.2](#) indicate that undefined bits should be set to all zeros, it is also considered acceptable for V5 bit 3 to be set to indicate "apparent" errors in this situation [e.g., to indicate errors that it detects based on processing of the VT path signal using the most recent valid pointer value (and thus the last known V5 byte location)]. As in the REI-P case, this is consistent with the specifications in ANSI T1.105, but it differs from those in ITU-T Rec. G.707 [which indicate that the REI-V bit is also undefined if the VT PTE has detected a VT Path UNEQ or TIM (UNEQ-V or TIM-V) defect].

Bit 4 of the V5 byte is allocated for a VT Path Remote Failure Indication (RFI-V) in the byte-synchronous DS1 mapping. See [Section 6.2.1.3.3](#) for the criteria related to the generation and detection of RFI-V signals.

Bits 5 through 7 of the V5 byte are allocated for a VT path signal label to indicate the content of the VT SPE or the presence of an extended VT path signal label in the Z7 byte (see below). The codes in [Table 3-4](#) are assigned for use by all VT PTE, and the code generated by a particular VT PTE depends on its provisioned (or only supported) functionality.

- R3-41 [30v3]** Except for the case covered in **R3-43 [1149]**, VT PTE supporting a VT path connection that is "provisioned" shall generate a valid, non-zero VT signal label. If the mapping contained in the VT SPE is one of the specific possibilities listed in [Table 3-4](#) or [Table 3-5](#) (including the case of a mapping that is under development), then the corresponding code from the table shall be used. If the mapping is not specifically listed and does not need to be identified via a VT path signal label, then the code for "Equipped – Nonspecific Payload" shall be used.

- R3-42 [31]** For VT path connections that are not equipped, or that are equipped but not provisioned, the NE (e.g., the NE's STS PTE) shall generate all-zeros VT SPEs with "valid" VT payload pointers.

In the above requirements, "provisioned" means that the VT PTE has been configured for a mapping (or only supports one mapping), and has been assigned a timeslot (or is hardwired to a specific timeslot) in an STS-1 SPE. Unless both of these conditions are met, the connection is not considered provisioned. In addition, a supplier may choose not to consider a connection to be provisioned unless the VT PTE has been assigned a signal (e.g., via a cross-connection) to map or multiplex into the VT SPE that it is originating.

In addition, note that the “valid” pointers referred to in **R3-42 [31]** may contain any value that can be interpreted correctly by downstream VT pointer processors without causing an LOP-V (see [Section 6.2.1.1.3](#)), and that an all-zeros VT SPE results in a valid VT path BIP-2 code and a signal label of ‘000’, indicating “Unequipped”.

R3-43 [1149] VT PTE or supervisory generation equipment supporting a VT path connection that is provisioned to carry a VT Supervisory-unequipped (TEST-V) signal shall generate an Unequipped VT path signal label (V5 bits 5 through 7 set to ‘000’), a valid VT path BIP-2 (V5 bits 1 and 2), a valid VT path status (V5 bits 3 and 8, and possibly Z7 bits 5 through 7), and a valid or all-zeros VT path trace (J2). In addition, it shall set the Z6 byte to all-zeros.

The contents of the remaining VT POH bits and bytes and the VT payload capacity in a TEST-V signal are currently undefined (see [Section 3.2](#)). Also, information and criteria related to the use of the TEST-V signal may be included in a future issue of this document.

Bit 8 of the V5 byte (along with bits 5 through 7 of the Z7 byte, see below) is allocated for a VT Path Remote Defect Indication (RDI-V) signal. [Section 6.2.1.3.3](#) contains the criteria related to the generation and detection of RDI-V.

Table 3-4 VT Path Signal Label Assignments (Non-Extended)

Code	Content of VT SPE			
	VT1.5	VT2 ^a	VT3	VT6
000	Unequipped			
001	Equipped – Nonspecific Payload			
010	Asynchronous Mapping for DS1	Asynchronous Mapping for 2.048 Mb/s	Asynchronous Mapping for DS1C	Asynchronous Mapping for DS2
011	Bit-synchronous Mapping for DS1 ^b	Bit-synchronous Mapping for 2.048 Mb/s	–	–
100	Byte-synchronous Mapping for DS1	Byte-synchronous Mapping for 2.048 Mb/s	–	–
101	Extended Signal Label ^c			
110	O.181 Test Signal Mapping ^d			
111	VT SPE AIS ^e			

Notes:

- a Although the VT2 mappings for 2.048 Mb/s signals are not included in this document, the signal label codes for those mappings have been assigned in ANSI T1.105.
- b The DS1 bit-synchronous mapping has been removed from the SONET standards and criteria. However, for backward compatibility between future mappings and equipment that supports the DS1 bit-synchronous mapping, the signal label that was assigned to that mapping will remain defined.
- c This code is used to indicate that an extended VT path signal label code is present in the Z7 byte.
- d This code/mapping assignment is specified in ITU-T Rec. G.707 for use in SDH networks (to indicate any non-virtually concatenated mapping defined in ITU-T Rec. O.181 that does not correspond to a mapping defined in Rec. G.707).
- e This code assignment is specified in ANSI T1.105.02, *Synchronous Optical Network (SONET) – Payload Mappings*, and ITU-T Rec. G.707 to indicate VT SPE AIS (a.k.a., VC-AIS). It is generated by a VT tandem connection maintenance source if no valid incoming signal is available and a replacement signal is being generated.

VT Path Trace (J2) – This byte is used to transport a repetitive 16-byte message [or Access Point Identifier (API)] so that the receiving VT PTE can verify its continued connection to the intended transmitting VT PTE.

Sections 6.2.1.1.9 and 6.2.3.2.3.A contain the criteria related to loading and detecting the VT path trace message, along with information about possible message formats.

R3-44 [1150] If no message has been loaded by the user for transmission in the J2 byte, then that byte shall be set to all zeros (i.e., to ASCII NULL characters).

VT Path Growth (Z6) – The Z6 byte (which is referred to in SDH as the N2 byte) was originally allocated for future growth, and has now been reserved for use in a VT tandem connection function. Refer to ANSI T1.105 and T1.105.05, and ITU-T Rec. G.707 for details.

VT Path Growth (Z7) – Bit 1 of the Z7 byte (which is referred to in SDH as the K4 byte) is allocated for an extended VT path signal label that is used to indicate the content of the VT SPE in cases where additional codes (beyond those that could be carried in V5 bits 5 through 7) are needed. If the signal label in V5 bits 5 through 7 is ‘101’, then this bit carries an extended VT path signal label as described below. Otherwise, this bit is undefined (see [Section 3.2](#)).

In cases where Z7 bit 1 carries an extended VT path signal label, that signal label is contained in VT superframes 12 (bit 1) through 19 (bit 8) of a 32-VT superframe multiframe as defined in [Figure 3-29](#). As shown in that figure, the code ‘0111 1111 110’ is used in VT superframes 1 through 11 as a Multiframe Alignment Signal (MFAS), and therefore any future use of the bits that are currently “reserved for future use” will need to avoid any sequence that includes nine consecutive “ones”. Also as shown in the figure, the bit in VT superframe 20 is set to ‘0’, and the remaining 12 bits (in VT superframes 21 through 32) are reserved for future standardization, and therefore are currently considered undefined (see [Section 3.2](#)).

Figure 3-29 Z7 Bit 1 Multiframe Definition

SFN	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	
Use	MFAS											Signal Label								–	Reserved for Future Use												
NV	0	1	1	1	1	1	1	1	1	1	0	b	b	b	b	b	b	b	b	b	0	0	0	0	0	0	0	0	0	0	0	0	0
												1	2	3	4	5	6	7	8														

SFN VT Superframe Number

NV Normal Value

Of the 256 possible values that can be carried in the extended VT signal label [i.e., 00 to FF (hex)], only the codes defined in [Table 3-5](#) have been assigned for the mappings described in [Section 3.4](#) or other documents (as noted). The codes in [Table 3-5](#) are assigned for use by all VT PTE that supports extended VT signal labels, and the code generated (under normal conditions) by a particular VT PTE depends on its provisioned (or only supported) functionality. The remaining codes are reserved to be assigned as needed for future VT payload-specific mappings.

Table 3-5 Extended VT Path Signal Label Assignments

Code (Hex)	Content of the VT SPE
00 to 07	Reserved ^a
08	Experimental or developmental mapping ^b
09	ATM mapping
0A	Mapping of HDLC/PPP framed signal
0B	Mapping of HDLC/LAPS framed signals
0C	Virtually concatenated O.181 test signal ^c
0D	GFP mapping
D0 to DF	Reserved for proprietary use ^d
FF	Reserved

Notes:

- a These codes are reserved to avoid possible confusion with the non-extended VT path signal label codes listed in [Table 3-4](#).
- b This code is specified to be used for developmental or experimental activities (i.e., for cases where the payload mapping is not one of the specific mappings for which a different code has been assigned). Note that prior to the definition of this code (in the 2001 version of ANSI T1.105 and Issue 4 of this document) the “Equipped – Nonspecific Payload” code (in V5 bits 5 through 7) was often used for such activities; however, since that code is required to be considered “Matched” by all VT PTE (see [Section 6.2.1.1](#)), its use did not provide the necessary isolation between developmental or experimental activities and the rest of the SONET network. Also note that in cases where the developmental or experimental activities result in the specification of a new mapping that is assigned its own payload-specific code, no compatibility will exist between equipment that uses this code and subsequent equipment that uses the new code. If such compatibility is desired, it will be necessary for the equipment that uses this code to be able to be reconfigured to utilize the new code.
- c This code/mapping assignment is specified in ITU-T Rec. G.707 for use in SDH networks (to indicate any virtually concatenated mapping defined in ITU-T Rec. O.181 that does not correspond to a mapping defined in Rec. G.707).
- d These codes are not subject to further standardization. See Annex A of ITU-T Rec. G.806 for information related to their use.

Bit 2 of the Z7 byte is allocated for a VT virtual concatenation string. The bit is multiframed in 32 VT superframes to form a 32-bit string. This function is described in [Section 5.2.5](#).

Bits 3, 4 and 8 of the Z7 byte are allocated for future growth, while bits 5 through 7 of that byte (along with bit 8 of the V5 byte, see above) are allocated for an RDI-V signal. [Section 6.2.1.3.3](#) contains the criteria related to the generation and detection of RDI-V.

3.4 Payload Mapping

This section describes the mapping of various payloads into VT and STS SPEs.

3.4.1 Sub-STS-1 Mappings

This section describes mappings of various payloads into the payload capacities of VT_n or VT_n-X_c SPEs (see Figures 3-23 and 3-24), which are then transported in VT-structured STS-1 SPEs. These are organized into mappings that are specific to particular sizes of VT SPEs (e.g., the asynchronous mapping for DS1 signals into VT1.5 SPEs), and mappings that apply across multiple sizes of VT SPEs (e.g., the ATM mapping). For reference, Table 3-6 lists various attributes of VT SPEs.

Table 3-6 VT_n and VT_n-X_v SPE Attributes

Type	Number of Payload Capacity Bytes per VT Superframe	Nominal Capacity (Mb/s)
VT1.5	100	1.6
VT1.5-X _v	X × 100	X × 1.6
VT2 ^a	136	2.176
VT2-X _v ^a	X × 136	X × 2.176
VT3 ^a	208	3.328
VT3-X _v ^a	X × 208	X × 3.328
VT6 ^a	424	6.784
VT6-X _v ^a	X × 424	X × 6.784

Note:

- a. Few existing SONET NEs support the capability to multiplex, switch or transport VT2, VT3 and VT6 SPEs, and therefore the applications in which mappings into these types of VT SPEs are utilized may be severely limited.

In addition, to the payload-specific criteria in the following subsections, **R3-45 [32]** applies in all cases where VT-structured STS-1 SPEs are utilized.

R3-45 [32] The H4 byte shall be used to indicate the phase of the V1 through V4 bytes in the 500-μs (i.e., 4-frame) VT superframe. The allocation of the bits in the H4 byte, and the correspondence of the H4 code with the V1 through V4 bytes shall be as shown in Figure 3-30.

When the H4 byte is used to indicate the phase of the V1 through V4 bytes, bits 1 through 6 are considered to be undefined. Therefore the criteria in Section 3.2 are applicable, including the objective that undefined bits and bytes be set to zero. However, early versions of ANSI T1.105 and the corresponding Telcordia criteria documents required H4 bits 1 through 6 to be set to non-zero values, while ITU-T Rec. G.707 indicates H4 bits 3 and 4 must be set to '11'. To accommodate equipment

built to these standards or previous requirements, an NE that does either of the following will also be considered to have met the objective (for the undefined bits in H4):

- Sets bits 1 through 6 to all-ones
- Sets bits 1 through 6 to the 48-frame sequence described below:
 - All of the bits (including bits 7 and 8) are zero in frame 0
 - Bits 1 and 2 count from '00' through '11' over 24 frames (i.e., '00' appears in 6 consecutive frames, then '01' for 6 frames, etc.) and then repeat
 - Bits 3 and 4 count from '00' through '10' over 6 frames, and then repeat
 - Bits 5 and 6 count from '00' through '11' over 16 frames, and then repeat.

Figure 3-30 H4 Byte Coding Sequence for VT-Structured STS-1 SPEs

H4 Byte bits								V1 through V4	Time
1	2	3	4	5	6	7	8	byte in the next frame	
X	X	X	X	X	X	0	0	V1	0
X	X	X	X	X	X	0	1	V2	
X	X	X	X	X	X	1	0	V3	
X	X	X	X	X	X	1	1	V4	

500- μ s VT Superframe

Although there are no specific criteria related to the alignment algorithm that STS PTE uses in determining the phase of the received H4 byte sequence, it is important that the algorithm be tolerant of single bit errors in the H4 sequence, and that it gain alignment quickly after the H4 byte is located. It is also important that the algorithm quickly recognize and respond to actual changes in the phase of the H4 bytes. If STS PTE cannot determine the phase of the H4 sequence, then it also cannot determine the phase of the V1 through V4 bytes. This is likely to result in a disruption of traffic and the detection of multiple LOP-V defects (see [Section 6.2.1.1.3](#)).

3.4.1.1 Mappings into Specific Size VT SPEs

Asynchronous mappings are defined for clear-channel transport of (among other possibilities) signals that meet the DSX requirements in GR-499-CORE. At asynchronous interfaces, frame acquisition and generation are not required for mapping purposes, although frame acquisition may be required for monitoring purposes (for example, see [Section 6.2.2.9](#)). Byte-synchronous mappings, which are currently defined (in this document) only for DS1 signals, allow direct identification and access to the DS0 channels that are carried. At byte-synchronous interfaces, frame acquisition and generation capabilities are required.

3.4.1.1.1 Byte-Synchronous Mapping for DS1

A byte-synchronous mapping of a DS1 into the payload capacity of a VT1.5 SPE is defined to allow downstream SONET NEs direct identification and access to the 24 DS0 channels that are carried. In some applications, a DS1 signal that appears at a DS1 interface is byte-synchronously mapped into the VT1.5, while in other applications the mapping is used to transport 24 DS0 channels without a DS1 interface [e.g., in an Integrated Digital Loop Carrier (IDLC) application with SONET-based Remote Digital Terminals (RDTs) and Integrated Digital Terminals (IDTs)]. In addition, some NEs may provide DS0 rearrangement (i.e., cross-connect and grooming) capabilities.

Figure 3-31 shows the byte-synchronous mapping for a DS1 (or 24 DS0s) into a VT1.5. The S_1 , S_2 , S_3 and S_4 bits are allocated to carry signaling for the 24 DS0 channels, the F-bit is allocated to carry the DS1 frame bit, and the P-bits are allocated for indicating the phase of the signaling and the frame bits on a per-VT basis.

R3-46 [33] If the byte-synchronous DS1 mapping is provided, it shall be as shown in Figure 3-31.

CR3-47 [1151] A SONET NE that supports the byte-synchronous DS1 mapping and DS0 rearrangement capabilities may be required to be capable of inserting a pattern other than all zeros (e.g., all ones, '0111 1111') in any unassigned DS0 data bytes within that mapping.

In general, criteria concerned with the patterns that a SONET NE should insert in any unassigned DS0 data bytes within a DS1 signal that it originates are provided in NE-specific or application-specific documents such as GR-303-CORE, *Integrated Digital Loop Carrier System Generic Requirements, Objectives, and Interface* and TR-NWT-000170, *Digital Cross-Connect System (DSC 1/0) Generic Criteria*, and therefore have not been included here. On the other hand, the capability referred to in **CR3-47 [1151]** could be important in (for example) the case where an NE that supports DS0 rearrangements is connected to an NE that does not support DS0 rearrangements. In that case the contents of the outgoing DS1 at the latter NE would be dependent on the patterns inserted by the first NE, and therefore the first NE would need to be able to insert the same patterns as an NE that originates a DS1 signal.

In some applications the S-bits or the F-bit are not used to carry signaling or framing (see Section 3.4.1.1.1.A). In those cases they are considered undefined and the criteria in Section 3.2 are applicable. Similarly, if the P-bits are not needed to indicate the phase of the S-bits or the F-bits in a particular application, then they are also considered undefined.¹⁴

14. It is acceptable for the VT PTE to set the P-bits to either all-zeros or the 24-frame sequence in Figure 3-32 when they are undefined.

R3-48 [34] If the P-bits are being used to indicate the phase of the S-bits or the F-bits, then they shall set to the 24-frame sequence shown in [Figure 3-32](#).

Although there are no specific criteria related to the alignment algorithm that VT PTE uses in determining the phase of the received P-bit sequence (and therefore the phase of the S-bits or the F-bits), it is important that the algorithm be tolerant of single bit errors in the P-bit sequence, and that it gain alignment quickly after the P-bits are located. It is also important that the algorithm quickly recognize and respond to actual changes in the phase of the P-bits.

Figure 3-31 Byte-Synchronous Mapping for DS1 Payload

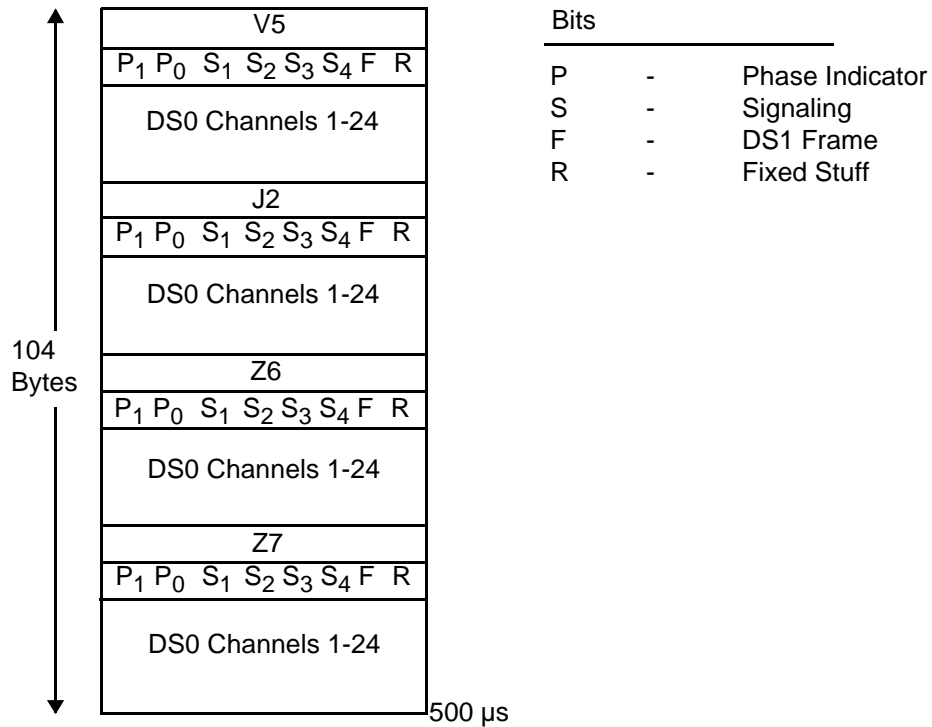


Figure 3-32 Byte-Synchronous DS1 Signaling and Framing Bit Assignments

	Signaling												DS1 Format	
	2-State				4-State				16-State				SF	ESF
P ₁ P ₀	S ₁	S ₂	S ₃	S ₄	S ₁	S ₂	S ₃	S ₄	S ₁	S ₂	S ₃	S ₄	F	F
00	A ₁	A ₂	A ₃	A ₄	A ₁	A ₂	A ₃	A ₄	A ₁	A ₂	A ₃	A ₄	F ₁	M ₁
00	A ₅	A ₆	A ₇	A ₈	A ₅	A ₆	A ₇	A ₈	A ₅	A ₆	A ₇	A ₈	S ₁	C ₁
00	A ₉	A ₁₀	A ₁₁	A ₁₂	A ₉	A ₁₀	A ₁₁	A ₁₂	A ₉	A ₁₀	A ₁₁	A ₁₂	F ₂	M ₂
00	A ₁₃	A ₁₄	A ₁₅	A ₁₆	A ₁₃	A ₁₄	A ₁₅	A ₁₆	A ₁₃	A ₁₄	A ₁₅	A ₁₆	S ₂	F ₁
00	A ₁₇	A ₁₈	A ₁₉	A ₂₀	A ₁₇	A ₁₈	A ₁₉	A ₂₀	A ₁₇	A ₁₈	A ₁₉	A ₂₀	F ₃	M ₃
00	A ₂₁	A ₂₂	A ₂₃	A ₂₄	A ₂₁	A ₂₂	A ₂₃	A ₂₄	A ₂₁	A ₂₂	A ₂₃	A ₂₄	S ₃	C ₂
01	A ₁	A ₂	A ₃	A ₄	B ₁	B ₂	B ₃	B ₄	B ₁	B ₂	B ₃	B ₄	F ₄	M ₄
01	A ₅	A ₆	A ₇	A ₈	B ₅	B ₆	B ₇	B ₈	B ₅	B ₆	B ₇	B ₈	S ₄	F ₂
01	A ₉	A ₁₀	A ₁₁	A ₁₂	B ₉	B ₁₀	B ₁₁	B ₁₂	B ₉	B ₁₀	B ₁₁	B ₁₂	F ₅	M ₅
01	A ₁₃	A ₁₄	A ₁₅	A ₁₆	B ₁₃	B ₁₄	B ₁₅	B ₁₆	B ₁₃	B ₁₄	B ₁₅	B ₁₆	S ₅	C ₃
01	A ₁₇	A ₁₈	A ₁₉	A ₂₀	B ₁₇	B ₁₈	B ₁₉	B ₂₀	B ₁₇	B ₁₈	B ₁₉	B ₂₀	F ₆	M ₆
01	A ₂₁	A ₂₂	A ₂₃	A ₂₄	B ₂₁	B ₂₂	B ₂₃	B ₂₄	B ₂₁	B ₂₂	B ₂₃	B ₂₄	S ₆	F ₃
10	A ₁	A ₂	A ₃	A ₄	A ₁	A ₂	A ₃	A ₄	C ₁	C ₂	C ₃	C ₄	F ₁	M ₇
10	A ₅	A ₆	A ₇	A ₈	A ₅	A ₆	A ₇	A ₈	C ₅	C ₆	C ₇	C ₈	S ₁	C ₄
10	A ₉	A ₁₀	A ₁₁	A ₁₂	A ₉	A ₁₀	A ₁₁	A ₁₂	C ₉	C ₁₀	C ₁₁	C ₁₂	F ₂	M ₈
10	A ₁₃	A ₁₄	A ₁₅	A ₁₆	A ₁₃	A ₁₄	A ₁₅	A ₁₆	C ₁₃	C ₁₄	C ₁₅	C ₁₆	S ₂	F ₄
10	A ₁₇	A ₁₈	A ₁₉	A ₂₀	A ₁₇	A ₁₈	A ₁₉	A ₂₀	C ₁₇	C ₁₈	C ₁₉	C ₂₀	F ₃	M ₉
10	A ₂₁	A ₂₂	A ₂₃	A ₂₄	A ₂₁	A ₂₂	A ₂₃	A ₂₄	C ₂₁	C ₂₂	C ₂₃	C ₂₄	S ₃	C ₅
11	A ₁	A ₂	A ₃	A ₄	B ₁	B ₂	B ₃	B ₄	D ₁	D ₂	D ₃	D ₄	F ₄	M ₁₀
11	A ₅	A ₆	A ₇	A ₈	B ₅	B ₆	B ₇	B ₈	D ₅	D ₆	D ₇	D ₈	S ₄	F ₅
11	A ₉	A ₁₀	A ₁₁	A ₁₂	B ₉	B ₁₀	B ₁₁	B ₁₂	D ₉	D ₁₀	D ₁₁	D ₁₂	F ₅	M ₁₁
11	A ₁₃	A ₁₄	A ₁₅	A ₁₆	B ₁₃	B ₁₄	B ₁₅	B ₁₆	D ₁₃	D ₁₄	D ₁₅	D ₁₆	S ₅	C ₆
11	A ₁₇	A ₁₈	A ₁₉	A ₂₀	B ₁₇	B ₁₈	B ₁₉	B ₂₀	D ₁₇	D ₁₈	D ₁₉	D ₂₀	F ₆	M ₁₂
11	A ₂₁	A ₂₂	A ₂₃	A ₂₄	B ₂₁	B ₂₂	B ₂₃	B ₂₄	D ₂₁	D ₂₂	D ₂₃	D ₂₄	S ₆	F ₆

Notes:

SF & ESF Format: F₁– F₆ = Frame Alignment Bits

SF Format: S₁– S₆ = Signaling Framing Bits

ESF Format: C₁–C₆ = Cyclic Redundancy Check-6 Bits
 M₁–M₁₂ = Data Link Bits

The criteria in the remainder of this section apply to VT PTE with DS1 interfaces that use the byte-synchronous DS1 mapping.

R3-49 [35] VT PTE with byte-synchronous DS1 interfaces shall be capable of accepting DS1 signals using the DS1 Superframe (SF) and Extended Superframe (ESF) formats defined in GR-499-CORE.

R3-50 [36] If the DS1 signal uses the ESF format and the F-bit is not used to transport the framing bits, then the following apply:

- The CRC-6 code in the received DS1 signals shall be monitored, and detected errors subsequently reported in the Data Link performance report message on the outgoing signal
- The correct CRC-6 code shall be calculated and inserted on the outgoing DS1 signals
- The NE shall send a performance report message every second to the sink, and receive a performance report message every second from the source.

If an NE provides DS0 rearrangement capabilities for incoming DS1 signals, then it would need to provide slip buffers so that the DS0s from different DS1s could be aligned for mapping into the VT SPE (which would be synchronized to the NE's clock, see [Section 5.4](#)). However, if the DS0 rearrangement capabilities are not provided, then the use of slip buffers could cause unnecessary DS1 frame slips in some synchronization failure scenarios. Therefore, the following objective applies, and an NE that meets the objective would use the VT payload pointer to frequency justify the VT SPE to the frame rate of the STS SPE (which would be synchronized to the NE's clock), as described in [Section 3.5.2](#).

O3-51 [37] If the NE does not provide DS0 rearrangement capabilities for an incoming DS1, then it should recover clock from the DS1 and use that clock in the creation of the VT SPE.

Similarly, if the NE provides DS0 rearrangement capabilities for output DS1s, then it would need to provide slip buffers to align the DS0s from different VT SPEs, and the output DS1 would need to be synchronized to the NE's clock. However, if the DS0 rearrangement capabilities are not provided, the output DS1 could either be timed from the NE's clock, or its timing could be recovered from the payload (as is done for asynchronously mapped signals).

The signaling transport modes, DS1 framing bit transport capabilities and pulse density assurance techniques that VT PTE must support depend on the application. Additional criteria regarding which modes and techniques must be supported appear in the criteria documents for the individual types of NEs. Shown in the following sections are the general criteria for VT PTE that supports the byte-synchronous mapping.

[Section 6.2.1](#) contains the criteria for DS1 and VT1.5 maintenance signals (see [Figure 6-9](#)). In addition, SONET NEs that provide DS0 path termination or rearrangement capabilities need to support DS0 maintenance and trunk conditioning capabilities (see [Section 6.2.1.6](#), and [Figures 6-10](#) through [6-12](#)).

3.4.1.1.1.A DS1 Signaling and Framing Bit Transport

Two signaling transport modes are defined, as discussed below:

1. Signaling transfer mode

In the signaling transfer mode, the signaling information carried in the robbed bit positions of the DS1 signal is transferred to the S-bit positions within the VT1.5. In this mode the DS1 framing bit does not need to be transported in the F-bit position (e.g., the far-end VT PTE generates a new framing bit sequence).

CR3-52 [38] The VT PTE may be required to support the signaling transfer mode.

R3-53 [39v2] If the signaling transfer mode is being used, then the following apply.

- The signaling information carried in the robbed bit positions of the DS1 signal shall be copied to the corresponding S-bit positions within the VT1.5 when the DS1 is mapped. In the VT1.5 to DS1 direction, the signaling information carried by the S-bits within the VT1.5 shall be written over the appropriate robbed bit positions of the outgoing DS1 bit stream.
- For DS1s using the Superframe format, the robbed bit positions (from which the signaling information is copied) shall be set to '1' when the DS1 is mapped into the VT1.5.¹⁵
- The phase of the S-bits shall be indicated by the P-bits in the same byte as shown in [Figure 3-32](#) for 2-state, 4-state and 16-state signaling schemes.
- The VT PTE shall generate a new framing bit pattern for the outgoing DS1 bit stream.
- The process shall be performed such that the signaling information is transferred through unaltered (i.e., with no changes or phase distortions).

2. Clear mode

In the clear mode, the VT PTE does not transfer signaling information between the S-bits and any robbed bit signaling positions (which may or may not actually contain signaling information) in the DS1 signals. It is defined for cases where (for example):

- The DS1 framing bit is carried in the F-bit to identify the phase of the robbed bit signaling information contained in the mapped DS1.
- The signaling is carried in a separate channel (e.g., the Common Signaling Channel in some IDLC systems), so robbed bit signaling is not used in the DS1.

CR3-54 [40] The VT PTE may be required to support the clear mode.

CR3-55 [41] If the clear mode is being used, then the VT PTE may be required to be user provisionable (on a per-DS1 basis) to transport the DS1 frame bit in the F-bit position of the VT1.5.

15. For DS1s using the ESF format, the VT PTE may either set the robbed bit positions to '1', or leave them set to their existing values.

R3-56 [42] If the clear mode is being used and the DS1 frame bit is carried in the F-bit, then the framing bits in the incoming DS1 shall be placed in the transmitted F-bits, and the received F-bits shall be placed in the outgoing DS1 signal (with no change in phase relative to the DS0 channels). The phase of the F-bit shall be indicated by the P-bits in the same byte as shown in [Figure 3-32](#) for the SF and ESF formats.

R3-57 [43] If the clear mode is being used and the DS1 frame bit is not carried in the F-bit (i.e., the F-bit is undefined), then the VT PTE shall generate a new framing bit pattern for the outgoing DS1 bit stream.

In addition to the two signaling transport modes defined above for use on a per-DS1 basis, in some applications it may be desirable for the VT PTE to perform, or not perform, signaling transfer on a per-DS0 basis. In some applications, this capability may only need to be provided on a static (i.e., user-provisionable) basis.

CR3-58 [44] The VT PTE may be required to be user provisionable on a per-DS0 basis to either perform or not perform signaling transfer (i.e., to provide signaling transfer/clear DS0 transport).

In other applications, signaling transfer/clear DS0 transport may need to be provided on a dynamic basis. To accomplish this, a scheme has been defined in which the S-bits are set to a specific value for the DS0s for which signaling transfer is not to be performed, and all other values in the S-bits indicate that signaling transfer is to be performed. In this scheme, the assignment of a particular DS0 as a clear DS0 is made by two NEs (e.g., an IDT and an RDT in an IDLC system) that are able to communicate that assignment with each other and set the S-bits accordingly. Intermediate NEs then need to be able to adjust their operation based on the code in the S-bits. This scheme avoids the difficulty of provisioning each DS0 channel at a DS1 interface and also allows dynamic signaling transfer/clear DS0 assignment.

CR3-59 [45] The VT PTE may be required to be user provisionable (on either a per-DS0 or per-DS1 basis) to perform dynamic signaling transfer/clear DS0 transport.

R3-60 [46] If dynamic signaling transfer/clear DS0 transport is being used, then the code ABCD=1001 on the S-bits shall be interpreted to mean that signaling transfer is not to be performed on that DS0 channel in either direction of transmission. Signaling transfer shall be performed for DS0 channels whose associated S-bits do not contain the code ABCD=1001.

Note that the ABCD code used to indicate that signaling transfer is not to be performed on a particular DS0 (in both directions) only appears in the S-bits in the SONET signal. It does not appear in any form in the DS1 signal. Therefore, this scheme will only support a single byte-synchronous SONET to DS1 and DS1 to byte-synchronous SONET conversion between the two NEs that make the signaling transfer/clear DS0 assignments and set the S-bits to the appropriate codes.

3.4.1.1.1.B Line Codes and Pulse Density Assurance

The following criteria on DS1 line codes and pulse density assurance techniques are applicable:

-
- R3-61 [47]** The VT PTE shall accommodate both the Alternate Mark Inversion (AMI) line code and the Bipolar with Eight Zero Substitution (B8ZS) line code.
-

Although a DS1 signal that uses the AMI line code and appears at a DS1 interface is required to have no more than 15 consecutive zeros and at least N ones in each and every time window of $8 \times (N+1)$ digital time slots (see GR-499-CORE), the AMI line code itself does not provide any form of pulse density assurance. Several methods of assuring the necessary pulse density are available, including the B8ZS line code and Zero Code Suppression (ZCS).¹⁶ In most applications it is expected that SONET NEs will use B8ZS unless they are connected to NEs that do not support that line code. If a SONET NE is simply providing transport for DS1s created by other NEs, then it is the responsibility of those source NEs to perform ZCS (or to support B8ZS). However, if a SONET NE provides DS0 path termination or DS0 rearrangement capabilities, then it is a DS1 source (even though the DS1 that it creates may not appear at a DS1 interface until the VT1.5 is terminated at another SONET NE). Therefore, a source SONET NE may need to support ZCS if the DS1s that it creates are expected to be received by NEs that do not support B8ZS.

-
- CR3-62 [49]** VT PTE that supports DS0 path terminations or DS0 rearrangement capabilities may be required to support ZCS for pulse density assurance.

- R3-63 [52v2]** The choice between AMI or B8ZS, and of ZCS (if provided), shall be provisionable by the user on a per-DS1 interface basis.
-

3.4.1.1.2 Asynchronous Mapping for DS1

An asynchronous mapping of a DS1 into the payload capacity of a VT1.5 SPE is defined for clear-channel transport of DS1 signals that meet the DSX-1 requirements in GR-499-CORE. If the asynchronous DS1 mapping is supported, then the following criteria are applicable.

-
- R3-64 [53]** The asynchronous mapping of a DS1 into a VT1.5 SPE shall be as shown in [Figure 3-33](#).
-

The asynchronous DS1 mapping contains 771 information (I) bits, 6 stuff control (C) bits, 2 stuff opportunity (S) bits and 8 overhead communication channel (O) bits in each VT1.5 SPE. The remaining 13 bits of the VT1.5 payload capacity are fixed stuff (R) bits. The O-bits are reserved for future communication purposes. The values contained in the R- and O-bits are currently undefined.

16. ZCS, which is also called “bit stuffing” in some documents, involves the insertion (at the DS1 source) of a ‘1’ in bit 7 of any all-zeros DS0 byte.

R3-65 [54] In each VT1.5 SPE, two sets of stuff control bits (C_1 and C_2) shall be used to control the two stuff opportunities (S_1 and S_2). $C_1C_1C_1 = 000$ shall be used to indicate that S_1 is an information bit, while $C_1C_1C_1 = 111$ shall be used to indicate that S_1 is a stuff bit. The C_2 bits shall be used to control S_2 in the same way.

The value contained in the S-bits when they are stuff bits is undefined.

R3-66 [55] Majority vote shall be used to make the stuff decision in the desynchronizer for protection against single bit errors in the C-bits.

R3-67 [1152] The stuffing mechanism that generates the C-bits shall be implemented so that, given a desynchronizer whose characteristics are that of a second-order low-pass filter with a cutoff frequency of 175 Hz, the output jitter is less than 1.0 Unit Intervals peak-to-peak (UI_{pp}) and 0.3 Unit Intervals root-mean-squared (UI_{rms}), assuming no jitter or wander at the input of the synchronizer and no pointer adjustments.

R3-68 [56] The stuffing mechanism that generates the C-bits shall be implemented so that, given a desynchronizer with filtering characteristics equal to the DS1 jitter transfer mask shown in [Figure 5-27](#), the output jitter is less than $0.7 UI_{pp}$, assuming no jitter or wander at the input of the synchronizer and no pointer adjustments.

R3-69 [903] The stuffing mechanism that generates the C-bits shall be implemented so that, given a desynchronizer with filtering characteristics equal to the DS1 jitter transfer mask shown in [Figure 5-27](#), the overall jitter transfer (i.e., for the synchronizer/desynchronizer pair) is less than that same DS1 jitter transfer mask.

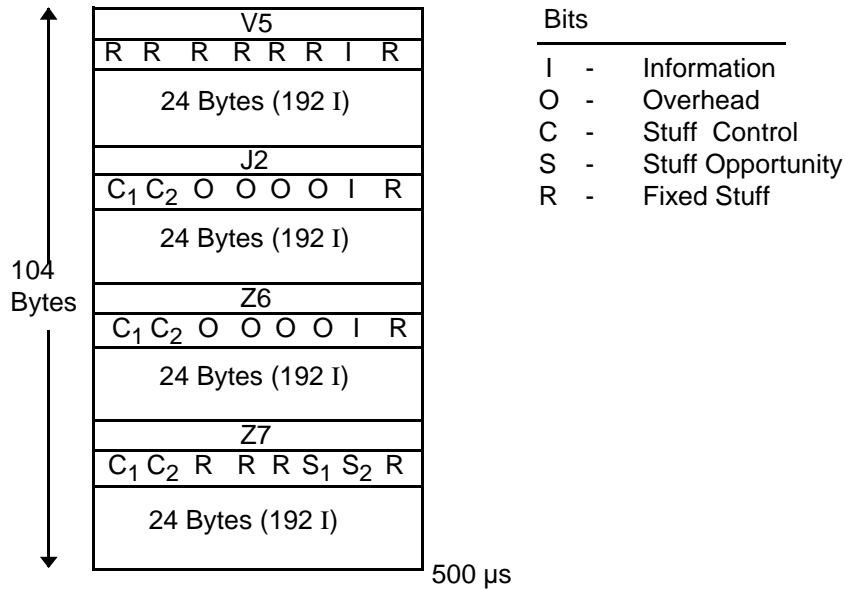
R3-70 [57] The DS1 interface shall accommodate both the AMI line code (assuming the DS1 source meets the zeros constraints in GR-499-CORE, see [Section 3.4.1.1.B](#)) and the B8ZS line code.

R3-71 [58] The choice of AMI or B8ZS shall be provisionable by the user on a per-DS1 interface basis.

Note that **R3-67 [1152]** was added to Issue 4 of this document to provide consistency with ANSI T1.105.02. The stuffing mechanism specifications in that standard do not address the case where jitter is applied to the input of the synchronizer (i.e., the case covered by **R3-69 [903]**) and, depending on the frequencies of various jitter components that are generated, may be either more or less stringent than **R3-68 [56]**.

[Section 6.2.1](#) contains the criteria for DS1 and VT1.5 maintenance signals (see [Figure 6-7](#)).

Figure 3-33 Asynchronous Mapping for DS1 Payload



3.4.1.1.3 Asynchronous Mapping for DS1C

An asynchronous mapping of a DS1C into the payload capacity of a VT3 SPE is defined for clear-channel transport of DS1C signals that meet the DSX-1C requirements in GR-499-CORE. If the asynchronous DS1C mapping is supported, then the following criteria are applicable.

R3-72 [59] The asynchronous mapping of a DS1C into a VT3 SPE shall be as shown in [Figure 3-34](#).

The asynchronous DS1C mapping contains 1574 information (I) bits, 12 stuff control (C) bits, 4 stuff opportunity (S) bits and 16 overhead communication channel (O) bits in each VT3 SPE. The remaining 58 bits of the VT3 payload capacity are fixed stuff (R) bits. The O-bits are reserved for future communication purposes. The values contained in the R- and O-bits are currently undefined.

R3-73 [60] Twice in each VT3 SPE, the two sets of stuff control bits (C₁ and C₂) shall be used to control the two stuff opportunities (S₁ and S₂). C₁C₁C₁ = 000 shall be used to indicate that S₁ is an information bit, while C₁C₁C₁ = 111 shall be used to indicate that S₁ is a stuff bit. The C₂ bits shall be used to control S₂ in the same way.

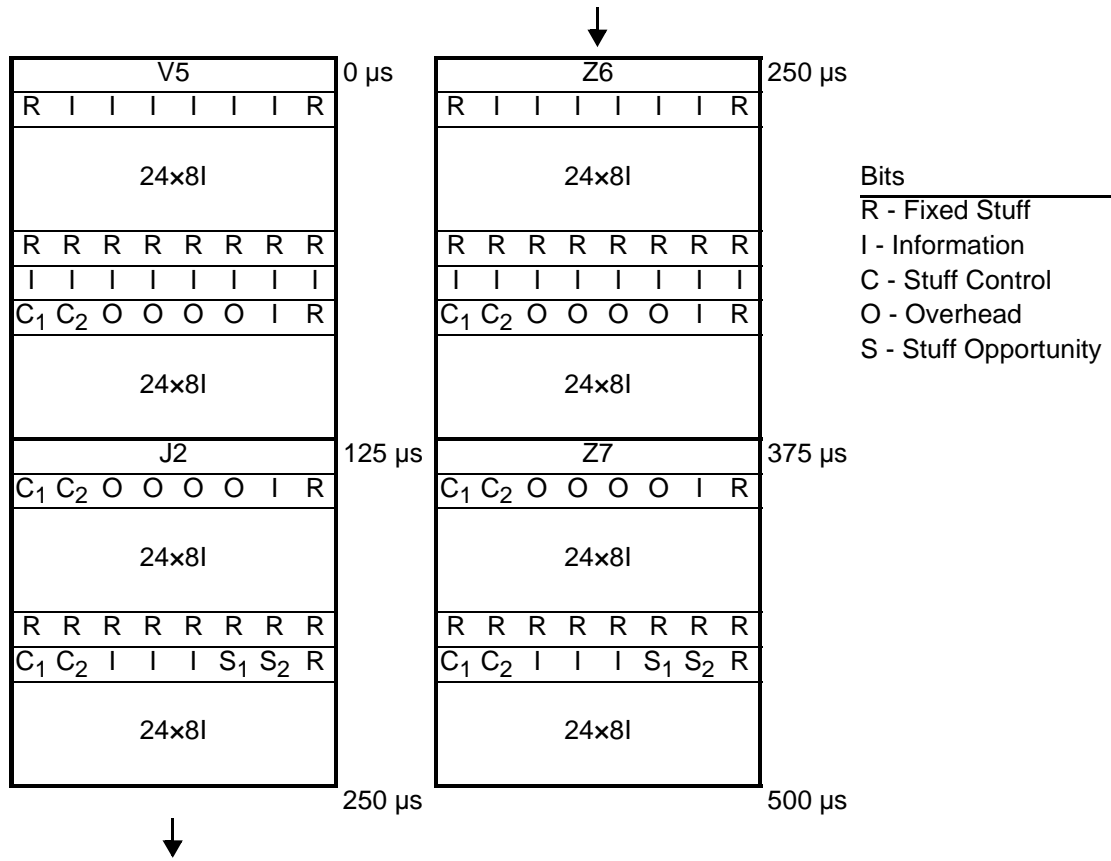
The value contained in the S-bits when they are stuff bits is undefined.

R3-74 [61] Majority vote shall be used to make the stuff decision in the desynchronizer for protection against single bit errors in the C-bits.

- R3-75 [62]** The stuffing mechanism that generates the C-bits shall be chosen so that, given a desynchronizer whose characteristics are that of a second-order low-pass filter with a cutoff frequency of 350 Hz, the output jitter is less than $1.0 U_{I_{pp}}$ and $0.3 U_{I_{rms}}$, assuming no jitter or wander at the input of the synchronizer and no pointer adjustments.
- R3-76 [904]** The stuffing mechanism that generates the C-bits shall be implemented so that, given a desynchronizer whose characteristics are that of a second-order low-pass filter with a cutoff frequency of 350 Hz, the overall jitter transfer (i.e., for the synchronizer/desynchronizer pair) is less than the DS1C jitter transfer mask in Section 7.3.2 of GR-499-CORE.
- R3-77 [63]** The DS1C interface shall accommodate both the AMI line code (assuming the DS1C source meets the ones density criteria from GR-499-CORE of at least 12.5% ones over any 150 consecutive bits) and the B8ZS line code.
- R3-78 [64]** The choice of AMI or B8ZS shall be provisionable by the user on a per-DS1C interface basis.
-

Section 6.2.1 contains the criteria for DS1C and VT3 maintenance signals (see Figure 6-7).

Figure 3-34 Asynchronous Mapping for DS1C Payload



3.4.1.1.4 Asynchronous Mapping for DS2

An asynchronous mapping of a DS2 into the payload capacity of a VT6 SPE is defined for clear-channel transport of DS2 signals that meet the DSX-2 requirements in GR-499-CORE. If the asynchronous DS2 mapping is supported, then the following criteria are applicable.

R3-79 [65] The asynchronous mapping of a DS2 into a VT6 SPE shall be as shown in [Figure 3-35](#).

The asynchronous DS2 mapping contains 3152 information (I) bits, 24 stuff control (C) bits, 8 stuff opportunity (S) bits and 32 overhead communication channel (O) bits in each VT6 SPE. The remaining 176 bits of the VT6 payload capacity are fixed stuff (R) bits. The O-bits are reserved for future overhead communication purposes. The values contained in the R- and O-bits are currently undefined (see [Section 3.2](#)).

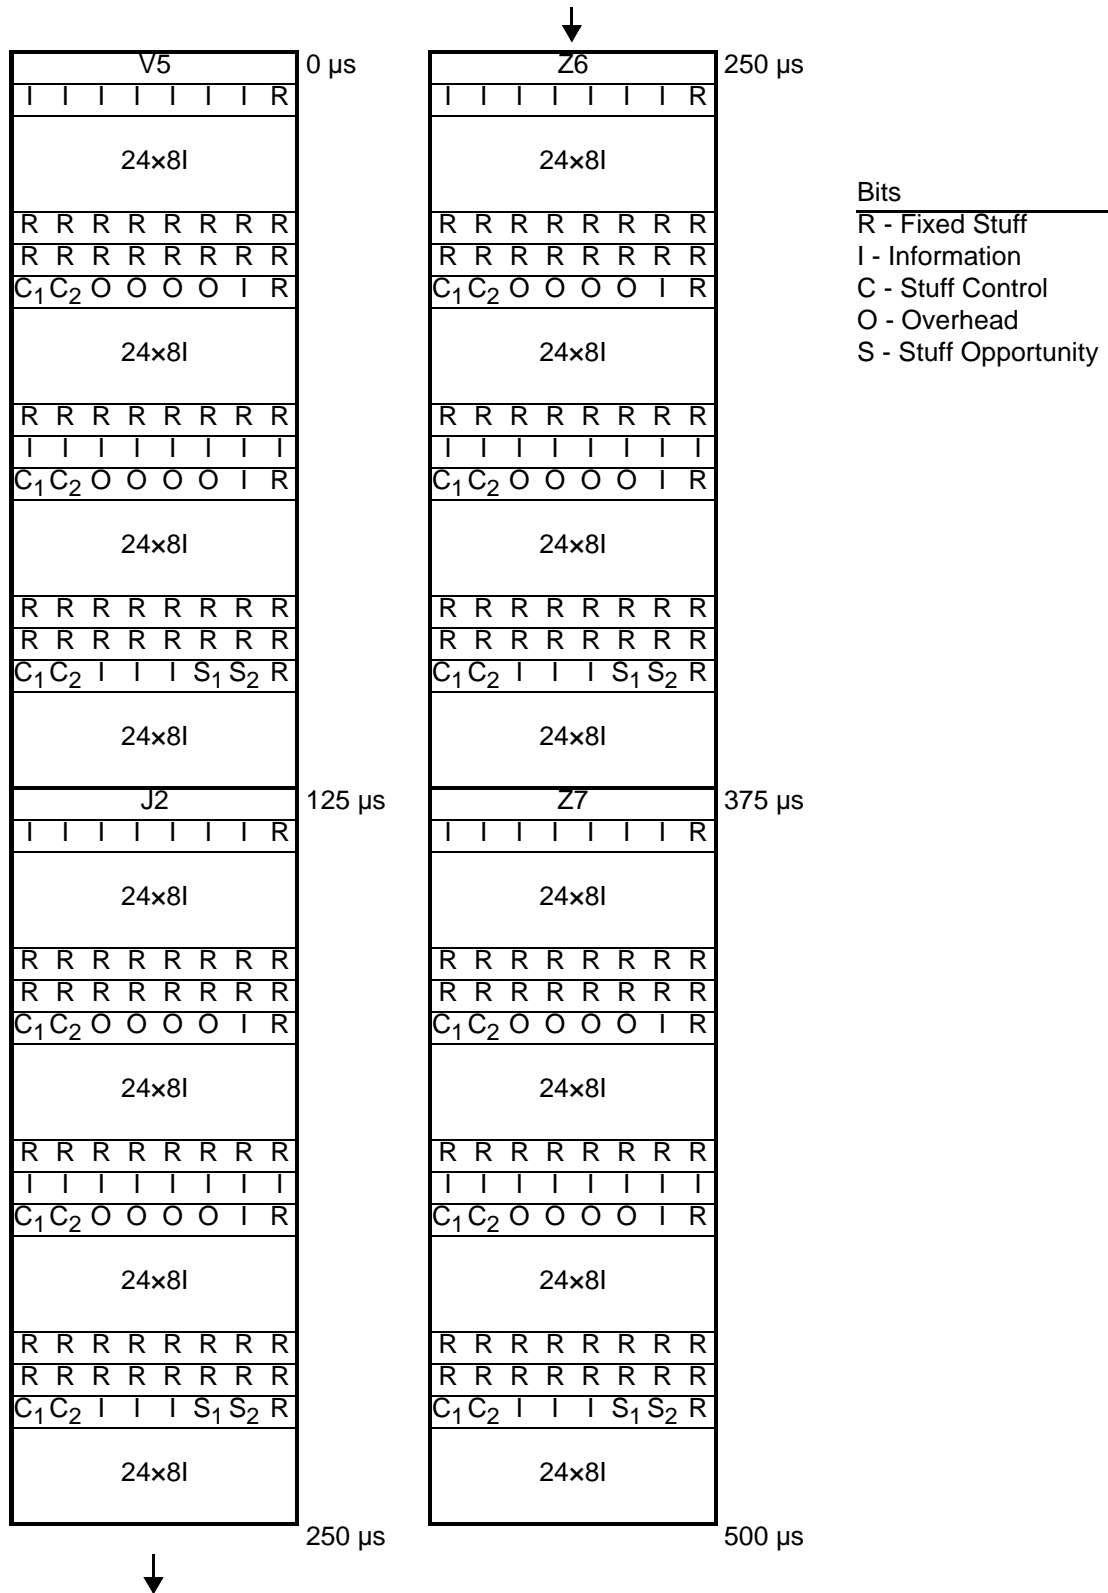
R3-80 [66] Four times in each VT6 SPE, the two sets of stuff control bits (C₁ and C₂) shall be used to control the two stuff opportunities (S₁ and S₂). C₁C₁C₁ = 000 shall be

used to indicate that S_1 is an information bit, while $C_1C_1C_1 = 111$ shall be used to indicate that S_1 is a stuff bit. The C_2 bits shall be used to control S_2 in the same way. The value contained in the S-bits when they are stuff bits is undefined.

- R3-81 [67]** Majority vote shall be used to make the stuff decision in the desynchronizer for protection against single-bit errors in the C-bits.
- R3-82 [68]** The stuffing mechanism that generates the C-bits shall be chosen so that, given a desynchronizer whose characteristics are that of a second-order low-pass filter with a cutoff frequency of 500 Hz, the output jitter is less than $1.0 U_{I_{pp}}$ and $0.3 U_{I_{rms}}$, assuming no jitter or wander at the input of the synchronizer and no pointer adjustments.
- R3-83 [905]** The stuffing mechanism that generates the C-bits shall be implemented so that, given a desynchronizer whose characteristics are that of a second-order low-pass filter with a cutoff frequency of 500 Hz, the overall jitter transfer (i.e., for the synchronizer/desynchronizer pair) is less than the DS2 jitter transfer mask in Section 7.3.2 of GR-499-CORE.

Section 6.2.1 contains the criteria for DS2 and VT6 maintenance signals (see Figure 6-7).

Figure 3-35 Asynchronous Mapping for DS2 Payload



3.4.1.2 Mappings That Apply to Multiple Sizes of VT SPEs

3.4.1.2.1 ATM to VT SPE Mapping

A method for mapping ATM cells (each of which consists of a 5-byte cell header and a 48-byte payload) into the payload capacity of a VTn or VTn-Xv SPE is defined. If this mapping is supported, then the following requirement is applicable.

R3-84 [1153] ATM cells shall be mapped into the VTn or VTn-Xv payload capacity by aligning the byte structure of every cell with the byte structure of the SPE. The entire payload capacity shall be filled with cells.

Filling the entire payload capacity of a nominal rate VTn or VTn-Xv SPE yields a transfer capacity as listed in [Table 3-6](#). Because the payload capacities of the various size VT SPEs are not necessarily integer multiples of the 53-byte ATM cell length, some cells may cross an SPE boundary.

For consistency with the ATM to STS SPE mapping discussed in [Section 3.4.2.2.1](#), an ATM cell payload scrambler is utilized. Details of this scrambler (a self-synchronizing scrambler with a generator polynomial of $x^{43}+1$) and various other topics such as the cell delineation process are contained in TR-NWT-001112, *Broadband-ISDN User to Network Interface and Network Node Interface Physical Layer Generic Criteria*.

3.4.1.2.2 HDLC-Framed Signal to VT SPE Mapping

This section contains criteria related to a mapping for HDLC-framed signals (e.g., those indicated by the VT path signal label codes '0A' and '0B') into the payload capacity of a VTn or VTn-Xv SPE. Unlike the HDLC-Framed Signal to STS SPE mapping discussed in [Section 3.4.2.2.2](#), scrambling of HDLC-framed signals is not required.

R3-85 [1154] The following apply if an HDLC-framed signal to VT SPE mapping is supported:

- The HDLC-framed signal shall be mapped into the VTn or VTn-Xv payload capacity by aligning the byte structure of every frame with the byte structure of the SPE
 - HDLC flags (i.e., '01111110' bytes) shall be used for interframe fill to account for the variable nature of the arrival of the HDLC frames
 - The entire VT payload capacity shall be filled with HDLC frames and HDLC flags (as necessary).
-

Filling the entire payload capacity of a nominal rate VT SPE yields a maximum transfer capacity as listed in [Table 3-6](#). Because HDLC frames are of variable length, a frame may cross an SPE boundary.

3.4.1.2.3 GFP to VT SPE Mapping

This section contains criteria related to a mapping of Generic Framing Procedure (GFP) frames into the payload capacity of a VTn or VTn-Xv SPE. If this mapping is supported, then the following requirement applies.

- R3-86 [1155]** GFP frames shall be mapped into the VTn or VTn-Xv payload capacity by aligning the byte structure of every frame with the byte structure of the SPE. The entire payload capacity shall be filled with frames.

Filling the entire payload capacity of a nominal rate VT SPE yields a transfer capacity as listed in [Table 3-6](#). Because GFP frames are of variable length (with no restrictions imposed by the mapping), some frames will cross an SPE boundary.

Similar to the ATM case discussed in [Section 3.4.2.2.1](#), a GFP frame consists of a GFP core header and a GFP payload area. Also similar to the ATM case, the GFP payload is scrambled using a self-synchronizing scrambler with a generator polynomial of $x^{43}+1$, and GFP Idle frames are inserted as necessary for rate synchronization purposes during the adaptation process. Therefore, neither of those functions is performed during the mapping process. Additional information and criteria concerning GFP and the GFP adaptation process are contained in ITU-T Rec. G.7041, *Generic framing procedure (GFP)*.

3.4.2 STS-1, STS-Nc, STS-1-Xv and STS-Nc-Xv Mappings

This section describes mappings that occupy the entire payload capacity of an STS-1, STS-Nc, STS-1-Xv or STS-3c-Xv SPE (see [Figures 3-4, 3-7, 3-9 and 3-10](#)). These are organized into mappings specific to certain size SPEs (e.g., the asynchronous DS3 to STS-1 SPE mapping), and mappings that apply to multiple size SPEs (e.g., the ATM mapping). For reference, [Table 3-7](#) lists various attributes of STS SPEs.

Table 3-7 STS SPE Attributes

Type	Columns	Rows	POH Column Number	Fixed Stuff Column Numbers	Number of Payload Capacity Columns	Nominal Capacity (Mb/s)
STS-1	87	9	1	30, 59	85	48.384
STS-1-Xv	X by 87	9	1 ^a	30, 59 ^a	X × 85	X × 48.384
STS-3c	261	9	1	None	260	149.76
STS-3c-Xv	X by 261	9	1 ^a	None ^a	X × 260	X × 149.76
STS-12c	1044	9	1	2 through 4	1040	599.04
STS-48c	4176	9	1	2 through 16	4160	2396.16
STS-192c	16704	9	1	2 through 64	16640	9584.64
STS-768c	66816	9	1	2 through 256	66560	38338.56

Note:

- a. This entry applies for each of the X constituent STS-1 or STS-3c SPEs.

3.4.2.1 Mappings into Specific Size STS SPEs

3.4.2.1.1 Asynchronous Mapping for DS3 into an STS-1 SPE

An asynchronous mapping for a DS3 into the payload capacity of an STS-1 SPE is defined for clear-channel transport of DS3 signals that meet the DSX-3 requirements in GR-499-CORE. If the asynchronous DS3 mapping is supported, then the following criteria are applicable.

R3-87 [69] The asynchronous mapping for a DS3 into an STS-1 SPE shall be as shown in [Figure 3-36](#).

The asynchronous DS3 mapping consists of 9 subframes every 125 μ s. Each subframe contains 621 information (I) bits, a set of 5 stuff control (C) bits, 1 stuff opportunity (S) bit and 2 overhead communication channel (O) bits. The remaining bits of the STS-1 payload capacity are fixed stuff (R) bits. The O-bits are reserved for future overhead communication purposes. The values contained in the R- and O-bits are currently undefined.

R3-88 [70] In each subframe, the set of five C-bits shall be used to control the S-bit. CCCCC = 00000 shall be used to indicate that the S-bit is an information bit, while CCCCC = 11111 shall be used to indicate that the S-bit is a stuff bit.

The value contained in the S-bit when it is a stuff bit is undefined.

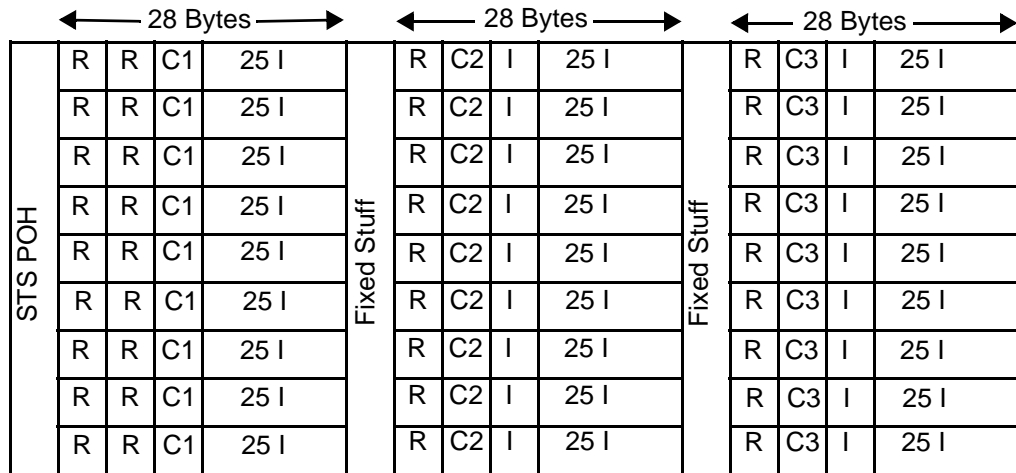
R3-89 [71] Majority vote shall be used to make the stuff decision in the desynchronizer for protection against single and double bit errors in the C-bits.

R3-90 [906] The stuffing mechanism that generates the C-bits shall be implemented so that, given a desynchronizer with filtering characteristics equal to the DS3 jitter transfer mask shown in [Figure 5-27](#), the output jitter is less than $0.4 U_{Ipp}$, assuming no jitter or wander at the input of the synchronizer and no pointer adjustments.


R3-91 [907] The stuffing mechanism that generates the C-bits shall be implemented so that, given a desynchronizer with filtering characteristics equal to the DS3 jitter transfer mask shown in [Figure 5-27](#), the overall jitter transfer (i.e., for the synchronizer/desynchronizer pair) is less than that same DS3 jitter transfer mask.

[Section 6.2.1](#) contains the criteria for DS3 and STS-1 path maintenance signals (see [Figure 6-5](#)).

Figure 3-36 Asynchronous Mapping for DS3 Payload



Bytes



- bits
- i: information (payload) bit
 - r: fixed stuff bit
 - c: stuff control bit
 - s: stuff opportunity bit
 - o: overhead communications channel bit

3.4.2.1.2 Asynchronous Mapping for DS4NA into an STS-3c SPE

An asynchronous mapping of a DS4NA (which has a nominal bit rate of 139.264 Mb/s) into the STS-3c payload capacity is defined for clear-channel transport of DS4NA signals that meet the DSX-4NA requirements in GR-499-CORE. If the asynchronous DS4NA mapping is supported, then the following criteria are applicable.

R3-92 [73] The asynchronous mapping of a DS4NA into an STS-3c SPE shall be as shown in [Figure 3-37](#).

Each 260-byte row of the STS-3c payload capacity is divided into 20 blocks of 13 bytes each (see [Figure 3-37](#) for the detailed structure of the blocks), and contains 1934 information (I) bits, a set of 5 stuff control (C) bits, 1 stuff opportunity (S) bit, 10 overhead communication channel (O) bits and 130 fixed stuff (R) bits. The O-bits are reserved for future overhead communications purposes. The values contained in the R- and O-bits are currently undefined.

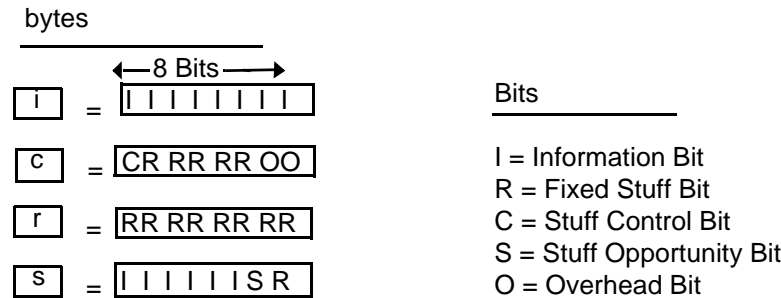

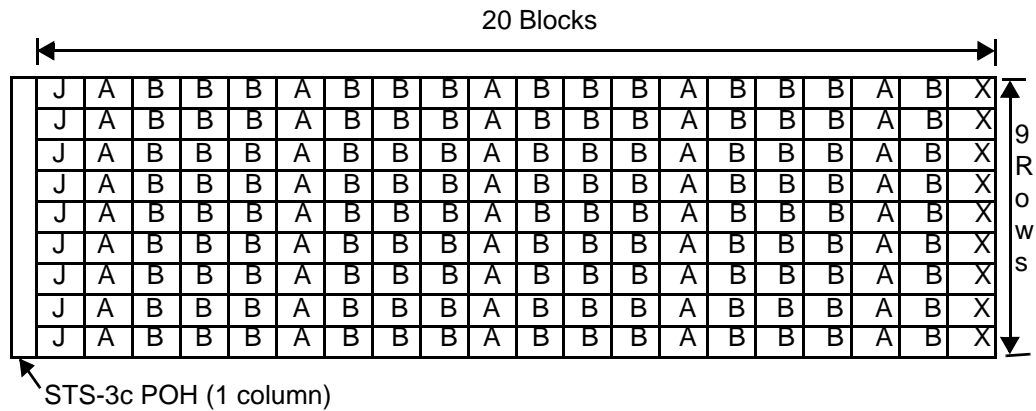
R3-93 [74] In each row, the set of five C-bits shall be used to control the S-bit.
CCCCC = 00000 shall be used to indicate that the S-bit is an information bit, while
CCCCC = 11111 shall be used to indicate that the S-bit is a stuff bit.

The value contained in the S-bit when it is a stuff bit is undefined.

R3-94 [75] Majority vote shall be used to make the stuff decision in the desynchronizer for protection against single and double bit errors in the C-bits.

Section 6.2.1 contains the criteria for DS4NA and STS-3c path maintenance signals.

Figure 3-37 Asynchronous Mapping for DS4NA Payload



3.4.2.1.3 Asynchronous Mapping for FDDI into an STS-3c SPE

An asynchronous mapping of a Fiber Distributed Data Interface (FDDI) signal into the STS-3c payload capacity is defined for clear-channel transport of FDDI signals. If this mapping is supported, then the following criteria are applicable.

-
- R3-95 [76]** The asynchronous mapping for a 125-Mb/s FDDI signal into an STS-3c SPE shall be as shown in [Figure 3-38](#).
-

Each 260-byte row of the STS-3c payload capacity is divided into 20 blocks of 13 bytes each (see [Figure 3-38](#) for the detailed structure of the blocks), and contains 1736 or 1735 information (I) bits, one set of 5 stuff control (C) bits, 1 stuff opportunity (S) bit, 2 overhead communication channel (O) bits and 336 or 337 fixed stuff (R) bits. The O-bits are reserved for future overhead communications purposes. The values contained in the R- and O-bits are currently undefined.

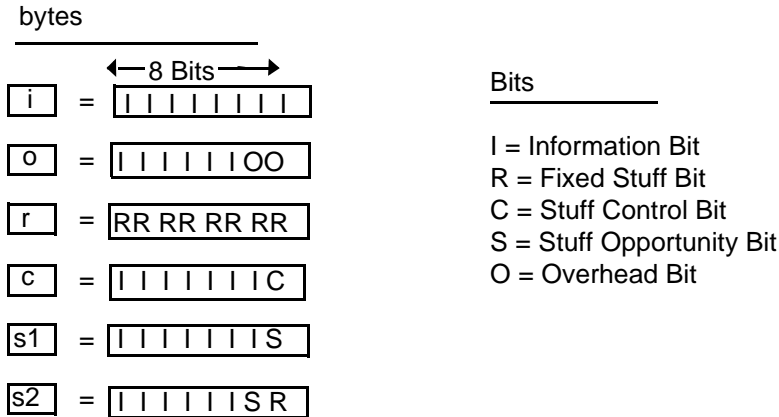

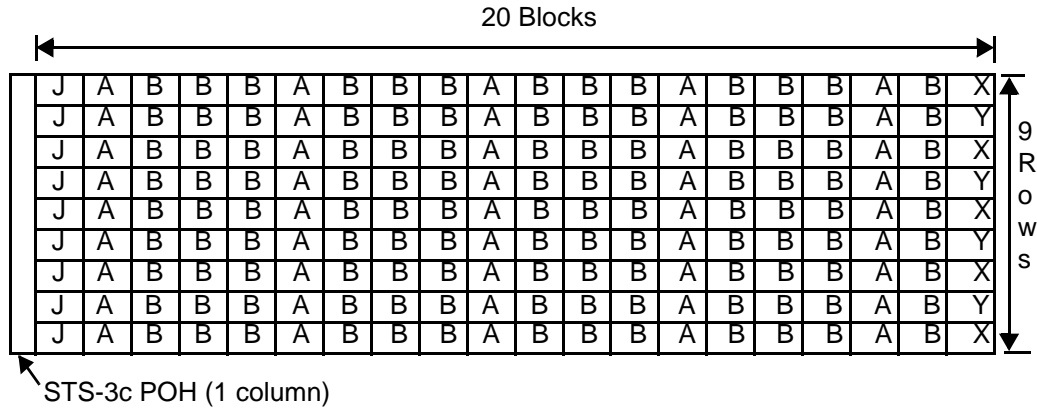
-
- R3-96 [77]** In each row, the set of five C-bits shall be used to control the S-bit. CCCCC = 00000 shall be used to indicate that the S-bit is an information bit, while CCCCC = 11111 shall be used to indicate that the S-bit is a stuff bit.
-

The value contained in the S-bit when it is a stuff bit is undefined.

-
- R3-97 [78]** Majority vote shall be used to make the stuff decision in the desynchronizer for protection against single and double bit errors in the C-bits.
-

[Section 6.2.1](#) contains the criteria for FDDI and STS-3c path maintenance signals.

Figure 3-38 Asynchronous Mapping for FDDI



3.4.2.1.4 DQDB Metropolitan Area Network (MAN) Mapping into an STS-3c SPE

A mapping for DQDB into an STS-3c SPE is defined in this section. In this mapping, DQDB layer slots are transported in the payload capacity of an STS-3c SPE. In addition, the mapping uses three of the STS POH bytes, as follows:

- The F2 (User Channel) and Z3 (Growth) bytes are used to carry the DQDB Layer Management (M1 and M2) bytes.
- Bits 1 and 2 of the H4 (Indicator) byte are used to carry the Link Status Signal (LSS).
- Bits 3 through 8 of the H4 byte are used to indicate the offset from the H4 byte to the beginning of the next 53-byte DQDB slot.

If the DQDB mapping is supported, then the following criteria apply.

R3-98 [80] DQDB slots shall be mapped into the STS-3c payload capacity by aligning the byte structure of every slot with the byte structure of the STS-3c SPE. The entire STS-3c payload capacity (i.e., 260 columns) shall be filled with slots, yielding a transfer capacity for DQDB slots of 149.760 Mb/s.

R3-99 [81] Bits 3 through 8 of the H4 byte shall contain a binary number in the range from '00000' (0) to '110100' (52) that indicates the offset between the H4 byte and the boundary of the first DQDB slot following the H4 byte.

The use of H4 to indicate the offset to the next DQDB slot is illustrated in [Figure 3-39](#), and the bit assignments for the H4 byte are shown in [Figure 3-40](#).

Because the STS-3c payload capacity is not an integer multiple of the DQDB slot length, some slots will cross the SPE boundary. Also, cell payload scrambling is used to provide security against payload information replicating (for example) the frame synchronous scrambling sequence used at the SONET Section layer.

Information about the status of the transmission link between the two adjacent Physical Layer Convergence Procedure (PLCP) entities connected to the ends of the transmission link is communicated by the LSS.

R3-100 [82] Bits 1 and 2 of the H4 byte shall be used to carry the LSS.

R3-101 [83] The F2 and Z3 bytes of the STS POH shall carry the DQDB M1 and M2 bytes, respectively.

The DQDB M1 and M2 bytes transport the DQDB layer node management information. These management bytes are generated at the head of a bus and are employed by the DQDB Layer Management Protocol Entity as IEEE 802.6, *Distributed Queue Dual Bus (DQDB) Subnetwork of a Metropolitan Area Network (MAN)*, describes.

Figure 3-39 DQDB Mapping into an STS-3c SPE

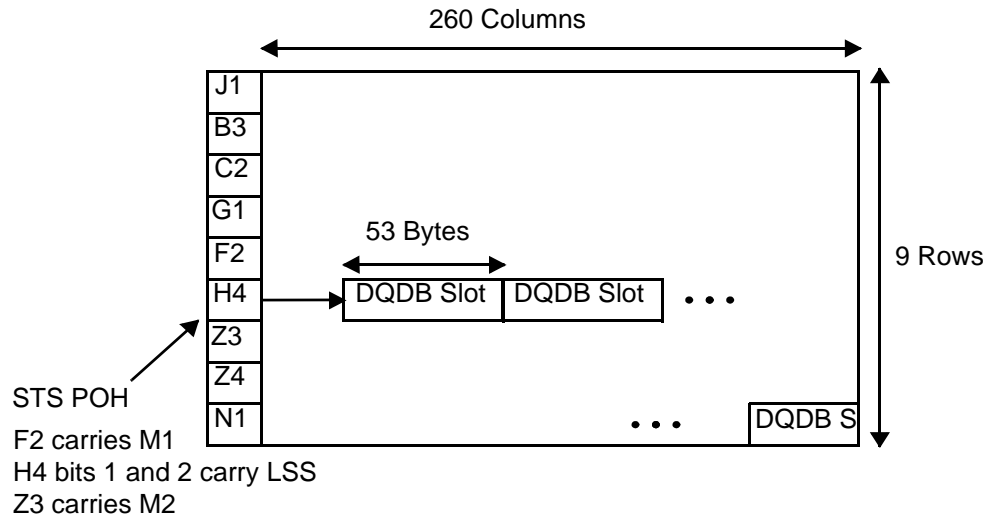


Figure 3-40 Bit Allocation for H4 Byte in DQDB Mapping

LSS		Slot Offset Indicator					
1	2	3	4	5	6	7	8

3.4.2.2 Mappings That Apply to Multiple Sizes of STS SPEs

3.4.2.2.1 ATM to STS SPE Mapping

A method for mapping ATM cells (each of which consists of a 5-byte cell header and a 48-byte payload) into the payload capacity of an STS-1, STS-Nc, STS-1-Xv or STS-3c-Xv SPE is defined. If this mapping is supported, then the following requirement is applicable.

R3-102 [72v2] ATM cells shall be mapped into the STS-1, STS-Nc, STS-1-Xv or STS-3c-Xv payload capacity by aligning the byte structure of every cell with the byte structure of the SPE. The entire payload capacity (e.g., 84 columns for an STS-1 SPE) shall be filled with cells.

Filling the entire payload capacity of a nominal rate STS SPE yields a transfer capacity as listed in [Table 3-7](#). Because the payload capacities of the various size STS SPEs are not integer multiples of the 53-byte ATM cell length, some cells will cross an SPE boundary.

Cell payload scrambling is used to provide security against payload information replicating (for example) the frame synchronous scrambling sequence used at the SONET Section layer. Details of the cell payload scrambler (a self-synchronizing scrambler with a generator polynomial of $x^{43}+1$) and various other topics such as the cell delineation process are contained in TR-NWT-001112.

3.4.2.2.2 HDLC-Framed Signal to STS SPE Mapping

This section contains criteria related to a mapping for HDLC-framed signals (e.g., those indicated by the STS path signal label codes '16' and '18') into the payload capacity of an STS-1, STS-Nc, STS-1-Xv or STS-3c-Xv SPE. Scrambling of HDLC-framed signals is necessary to provide security against payload information replicating (for example) the frame synchronous scrambling sequence used at the SONET Section layer. As described below, a self-synchronous scrambler with a generator polynomial of $x^{43}+1$ is used. Note that no initial seed is specified for this scrambler, and therefore the first 43 bits that are transmitted following a startup or reframe operation cannot be descrambled correctly.

R3-103 [1044v2] The following apply if an HDLC-framed signal mapping is supported.

- The HDLC-framed signal shall be mapped into the STS-1, STS-Nc, STS-1-Xv or STS-3c-Xv payload capacity by aligning the byte structure of every frame with the byte structure of the SPE.
- HDLC flags (i.e., '01111110' bytes) shall be used for interframe fill to account for the variable nature of the arrival of the HDLC frames.
- The entire STS payload capacity (e.g., 84 columns for an STS-1 SPE) shall be filled with HDLC frames and HDLC flags (as necessary).
- The HDLC-framed signal plus the interframe fill shall be scrambled before it is inserted into the STS payload capacity. In the reverse operation, after the STS path is terminated the payload shall be descrambled before it is passed to the HDLC layer.
- A self-synchronizing scrambler with a generator polynomial of $x^{43}+1$ shall be used.
- The most significant bit of each byte of the HDLC-framed signal and interframe fill (i.e., the bit that will be placed into bit 1 of a byte in the STS payload capacity, see [Figure 3-2](#)) shall enter the scrambler first, followed by the next most significant bit of that byte, etc.
- The scrambler shall run continuously (e.g., it shall not be reset for each SONET or HDLC frame).

Filling the entire payload capacity of a nominal rate STS SPE yields a maximum transfer capacity as listed in [Table 3-7](#). Because HDLC frames are of variable length, a frame may cross an SPE boundary.

Typically, HDLC frames include a CRC for error detection purposes. Two CRCs that are often used are a CRC-16 and a CRC-32. The use of the CRC-32 is preferred when the HDLC frames are carried over SONET, as that reduces the impact of the error multiplication effect caused by the $x^{43}+1$ scrambler.

O3-104 [1045] If an HDLC-framed signal mapping is supported and a CRC is applied over the HDLC payload signal, a CRC-32 should be used.

3.4.2.2.3 GFP to STS SPE Mapping

This section contains criteria related to a mapping of GFP frames into the payload capacity of an STS-1, STS-Nc, STS-1-Xv or STS-3c-Xv SPE. If this mapping is supported, then the following requirement applies.

R3-105 [1156] GFP frames shall be mapped into the STS-1, STS-Nc, STS-1-Xv or STS-3c-Xv payload capacity by aligning the byte structure of every frame with the byte structure of the SPE. The entire payload capacity (e.g., 84 columns for an STS-1 SPE) shall be filled with frames.

Filling the entire payload capacity of a nominal rate STS SPE yields a transfer capacity as listed in [Table 3-7](#). Because GFP frames are of variable length (with no restrictions imposed by the mapping), some frames will cross an SPE boundary.

Similar to the ATM case discussed in [Section 3.4.2.2.1](#), a GFP frame consists of a GFP core header and a GFP payload area. Also similar to the ATM case, the GFP payload is scrambled using a self-synchronizing scrambler with a generator polynomial of $x^{43}+1$, and GFP Idle frames are inserted as necessary for rate synchronization purposes during the adaptation process. Therefore, neither of those functions is performed during the mapping process. Additional information and criteria concerning GFP and the GFP adaptation process are contained in ITU-T Rec. G.7041.

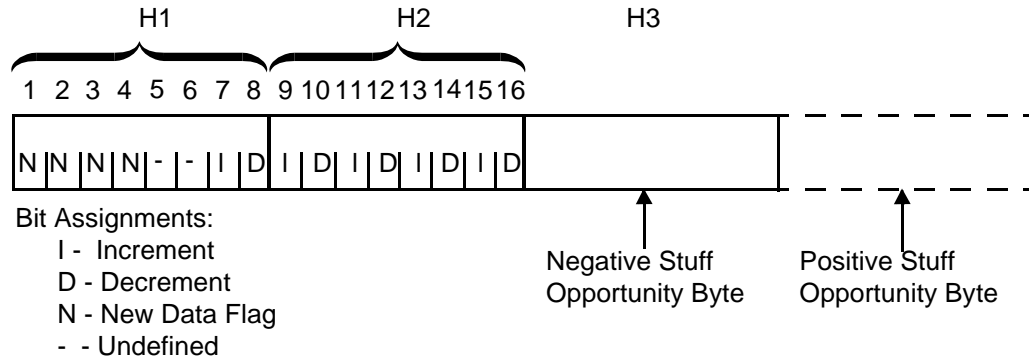
3.5 Payload Pointers

3.5.1 STS Payload Pointer


The STS payload pointer provides a method of allowing flexible and dynamic alignment of the STS SPE within the STS envelope capacity, independent of the actual contents of the SPE. Dynamic alignment means that the STS SPE is allowed to float within the STS envelope capacity. Thus, the pointer is able to accommodate differences not only in the phases of the STS SPE and the transport overhead, but in the frame rates as well.

The payload pointer is contained in the H1 and H2 bytes of the LOH, and designates the location of the byte where the STS SPE begins. These two bytes can be viewed as one word, as shown in Figure 3-41. Bits 1 through 4 of the pointer word carry the New Data Flag, and bits 7 through 16 carry the pointer value. Bits 5 and 6 of the STS pointer word are undefined, and therefore the criteria in Section 3.2 are applicable.

Figure 3-41 STS Payload Pointer (H1, H2) Coding



Normal Values



To Indicate:
 Set NDF: Invert 4 N-Bits
 Negative Stuff: Invert 5 D-Bits
 Positive Stuff: Invert 5 I-Bits

3.5.1.1 Pointer Value

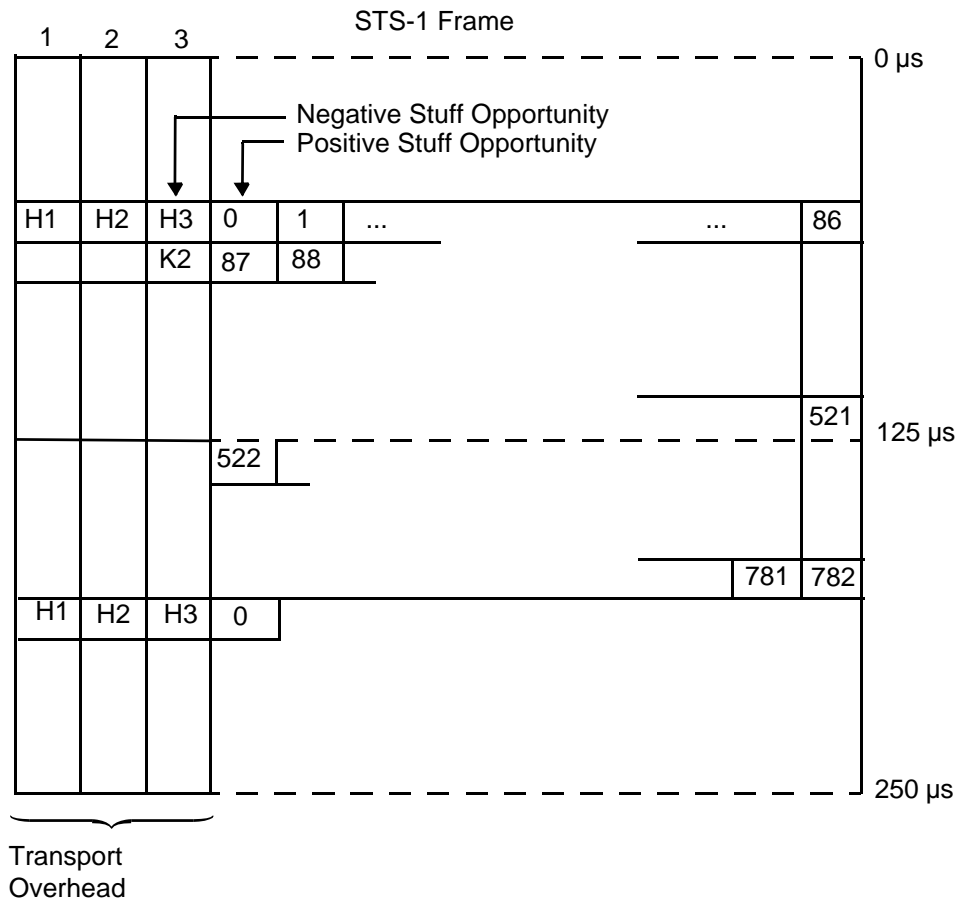
Bits 7 through 16 of the pointer word carry the pointer value, which indicates the offset between the pointer word and the first byte of the STS SPE (i.e., the J1 byte). The transport overhead bytes are not counted in the offset. For example, a pointer value of 0 indicates that the STS SPE starts in the byte location that immediately follows the H3 byte, while an offset of 87 indicates that it starts immediately after the K2 byte location.

R3-106 [85] The pointer value shall be a binary number with a range of 0 to 782, and shall indicate the offset between the pointer word and the first byte of the STS SPE (as shown in Figure 3-42).

Note that in the case of an STS-Nc SPE, there are at least two ways to view the relationship between the pointer value and the offset from the pointer word to the J1 byte. For example, it could be considered that there is a one-to-one correspondence, and that only the STS-Nc envelope capacity bytes that are associated with the first STS-1 of the STS-Nc are counted in determining the offset.

Alternatively, all of the bytes in the STS-Nc envelope capacity could be counted in determining the offset, and the NE could then transmit a pointer value equal to the offset divided by N. The net result is the same in either view, and (for example) a pointer value of 87 indicates that the J1 byte is located in the first byte of the STS-Nc envelope capacity following the K2 byte locations (i.e., the same as for an STS-1 SPE).

Figure 3-42 STS Pointer Offset Numbering



3.5.1.2 STS Frequency Justification

When the frame rate of the STS SPE is less than that of the transport overhead, the alignment of the SPE is periodically slipped back in time (using a positive stuff byte) and the pointer value is incremented by one. Similarly, when the frame rate of the STS SPE is greater than that of the transport overhead, the alignment of the SPE is periodically advanced in time (using a “negative stuff” byte) and the pointer value is decremented by one. In both cases, subsequent pointers contain the new offset.

R3-107 [86] When there is a frequency offset between the frame rate of the transport overhead and that of the STS SPE, the pointer value shall be incremented or

decremented as needed (although see rule 7 of **R3-117 [96v2]**), accompanied by a corresponding positive or negative stuff byte.

- R3-108 [87]** A pointer increment operation shall be indicated by inverting bits 7, 9, 11, 13 and 15 (the I-bits) of the pointer word. The positive stuff byte shall appear immediately after the H3 byte in the frame containing the inverted I-bits, as shown in **Figure 3-43**.

The value contained in the positive stuff byte is undefined.

-
- R3-109 [88]** A pointer decrement operation shall be indicated by inverting bits 8, 10, 12, 14 and 16 (the D-bits) of the pointer word. The H3 byte shall be used as the negative stuff byte, (i.e., it is used to carry an SPE byte in the frame containing the inverted D-bits), as shown in **Figure 3-44**.

The following criteria are applicable to the receiving STS pointer processor, and provide protection against errors on the I- and D-bits. The method in the objective is an extension of the majority vote method for I- and D-bits separately, and enhances performance during error bursts.

-
- R3-110 [89]** If **O3-111 [90]** (the “8 of 10” objective) is not met, then the increment decision shall be made by a majority vote of the I-bits, and the decrement decision shall be made by a majority vote of the D-bits.¹⁷

- O3-111 [90]** The increment/decrement decision should be made at the receiver by a match of 8 or more of the 10 I- and D-bits to either the increment or decrement indication.

As an example of the “8 of 10” objective, suppose a pointer processor receives a pointer word that has (due to transmission errors) three of the I-bits and two of the D-bits inverted from the previous pointer value. In that case, six of the 10 I- and D-bits would be correct for an increment operation (i.e., the three inverted I-bits plus the three noninverted D-bits), while four of the 10 would be correct for a decrement operation (i.e., the two noninverted I-bits plus the two inverted D-bits). If the pointer processor met the “8 of 10” objective, it would not interpret the pointer to be indicating either an increment or decrement operation.

17.If a majority of the I-bits and a majority of the D-bits are detected to be inverted in the same pointer word, there are several ways that a pointer processor that does not meet the “8 of 10” objective can react. For example, this requirement could be literally interpreted to mean that the pointer processor must consider the pointer value to be both incremented and decremented (i.e., it would consider the contents of the H3 byte to be part of the SPE, but would ignore the byte following H3). However, the intent of the requirement is that the pointer processor not consider the pointer value to be either incremented or decremented (which is consistent with the “8 or 10” objective).

Figure 3-43 Positive STS Pointer Adjustment Operation (Increment)





Figure 3-44 Negative STS Pointer Adjustment Operation (Decrement)



3.5.1.3 New Data Flag (NDF)

Bits 1 through 4 of the pointer word (i.e., the N-bits) carry an NDF, which can either be “normal” or “set”, and which allows an arbitrary change of the pointer value due to a change in the payload.

R3-112 [91] A normal NDF shall be indicated (during normal operation) by a '0110' code in the N-bits (see [Figure 3-41](#)). The NDF shall be set by inverting the N-bits to '1001.' The new alignment of the STS SPE shall be indicated by the pointer value accompanying the set NDF and takes effect at the offset indicated.

R3-113 [92] The decoding at the pointer processor shall be performed by majority voting (i.e., the NDF shall be detected as being set if three or four of the N-bits match the '1001' code). If a set NDF is detected, then the coincident pointer value shall replace the current value at the offset indicated by the new pointer value.

3.5.1.4 Concatenation Indicator

Concatenation indicators contained in the payload pointers of the second through Nth STS-1s in an STS-Nc are used to show that those STS-1s each contain part of the STS-Nc SPE (rather than individual STS-1 SPEs).

R3-114 [93] The first STS-1 within an STS-Nc shall have a normal pointer word.

R3-115 [94] All subsequent STS-1s within the STS-Nc shall have their pointer values (i.e., bits 7 through 16) set to all ones and their N-bits set to '1001' (i.e., set NDFs).

This value of the pointer word (i.e., '1001XX11 11111111') is the concatenation indicator, and does not indicate an offset.

R3-116 [95] A pointer processor in an NE that is transmitting or receiving an STS-Nc SPE shall perform the operations indicated by the pointer in the first STS-1 of the STS-Nc on all N of the STS-1s in that STS-Nc.

This means that if a decrement operation is transmitted or detected on the pointer in the first STS-1 of the STS-Nc, then the H3 bytes in all N of the STS-1s in the STS-Nc are considered to be negative stuff bytes. Similarly, if an increment operation is transmitted or detected, then the first N bytes of the STS-Nc envelope capacity following the last H3 byte are considered positive stuff bytes.

3.5.1.5 STS Payload Pointer Generation Rules

The pointer generation criteria from [Sections 3.5.1.1](#) through [3.5.1.4](#) are summarized in this section as a set of pointer generation rules. These rules also contain several additional requirements statements that do not appear in the preceding sections.

R3-117 [96v2] The STS payload pointer shall be generated according to the following rules.

1. During normal operation, a normal NDF is sent (i.e., the N-bits are set to '0110'), and the pointer value locates the start of the STS SPE within the STS envelope capacity.
2. The pointer value shall only be changed by the operations in rules [4](#), [5](#) or [6](#).

3. If an STS-Nc SPE is being transmitted, a normal pointer word is generated for the first STS-1 only. The concatenation indicator is generated in the other pointers. All operations indicated by the pointer in the first STS-1 apply to each STS-1 in the STS-Nc.
4. If a positive stuff is needed, the current pointer value is sent with the I-bits inverted, and the subsequent positive stuff opportunity is considered an undefined byte. Subsequent pointers contain the previous pointer value incremented by one.
5. If a negative stuff is needed, the current pointer value is sent with the D-bits inverted, and the subsequent negative stuff opportunity is overwritten with an SPE byte. Subsequent pointers contain the previous pointer value decremented by one.
6. If the alignment of the SPE changes for any reason other than rules 4 or 5, the new pointer value shall be sent accompanied by a set NDF. The set NDF only appears in the first frame that contains the new value. The new SPE begins at the first occurrence of the offset indicated by the new pointer value.
7. No increment or decrement operation shall be performed for three frames following any of the operations in rules 4, 5 and 6.¹⁸
8. For a nonterminated STS path, an incoming all-ones pointer word shall be regenerated or relayed with no more than a three-frame delay. When a non-all-ones pointer word is subsequently received, the downstream pointer shall be generated based on the pointer generation and interpretation criteria summarized in this requirement (**R3-117 [96v2]**) and **R3-119 [97]**.¹⁹

Rule 8 is applicable to LTE that processes STS pointers. If all of the LTE in a long chain of SONET NEs takes the full three frames allowed in Rule 8, then the detection and termination of an AIS-P defect at the STS PTE could be significantly delayed. Therefore the following objective is applicable.

O3-118 [908] LTE that processes STS pointers should regenerate or relay an incoming all-ones pointer word with no more than a one-frame delay.

Note that LTE cannot perform increment or decrement operations while it is performing all-ones pointer relay, and therefore the incoming STS SPE associated with the all-ones pointers cannot be expected to be passed through unaltered (see [Section 6.2.1.2.2](#)).

-
18. Note that in the pointer interpretation rules in the following section, there is no rule or requirement equivalent to pointer generation rule 7. If a pointer processor detects an increment or decrement operation within three frames after another pointer change operation (e.g., due to transmission errors), it can either ignore that operation or interpret it as a valid operation.
 19. Since the STS path is not terminated, the STS AIS detection and generation criteria in [Section 6.2.1.2.2](#) are not the applicable criteria for determining the value of the pointer word to be sent downstream.

3.5.1.6 STS Payload Pointer Interpretation

The pointer interpretation criteria from [Sections 3.5.1.1](#) through [3.5.1.4](#) are summarized in this section as a set of pointer interpretation rules. These rules also contain several additional requirements statements that do not appear in the preceding sections.


R3-119 [97] The STS payload pointer shall be interpreted according to the following rules.

1. During normal operation, the pointer value locates the start of the STS SPE within the STS envelope capacity.
 2. Any variation from the current pointer value shall be ignored unless a consistent new value is received three times consecutively, or the variation is one of the operations in [rules 4, 5 or 6](#). Any consistent new value received three times in succession shall replace the current value at the offset indicated by the new pointer value.
 3. If the pointer word contains the concatenation indicator, then the operations performed on that STS-1 are identical to those performed on the first STS-1 within the STS-Nc. [Rules 4 and 5](#) do not apply to this pointer word.
 4. If an increment is detected, then the byte following H3 shall be considered a positive stuff byte, and the current pointer value shall be incremented by one.
 5. If a decrement is detected, then H3 shall be considered a negative stuff byte, and the current pointer value shall be decremented by one.
 6. If a set NDF is detected, then the coincident pointer value replaces the current value at the offset indicated by the new pointer value.
-

3.5.2 VT Payload Pointer

Analogous to the STS payload pointer, the VT payload pointer provides a method of allowing flexible and dynamic alignment of the VT SPE within the VT superframe (and therefore within the STS SPE), independent of the actual contents of the VT SPE. The VT payload pointer is contained in the V1 and V2 bytes, and designates the location of the byte where the VT SPE begins (i.e., the V5 byte). The V1 and V2 bytes can be viewed as one word, as shown in [Figure 3-45](#). Bits 1 through 4 of the pointer word carry the NDF, bits 5 and 6 indicate the size of the VT, and bits 7 through 16 carry the pointer value.

Figure 3-45 VT Payload Pointer (V1, V2) Coding




Bit Assignments:

- I - Increment
- D - Decrement
- N - New Data Flag

To Indicate:

- Set NDF: Invert 4 N-Bits
- Negative Stuff: Invert 5 D-Bits
- Positive Stuff: Invert 5 I-Bits

Normal Values

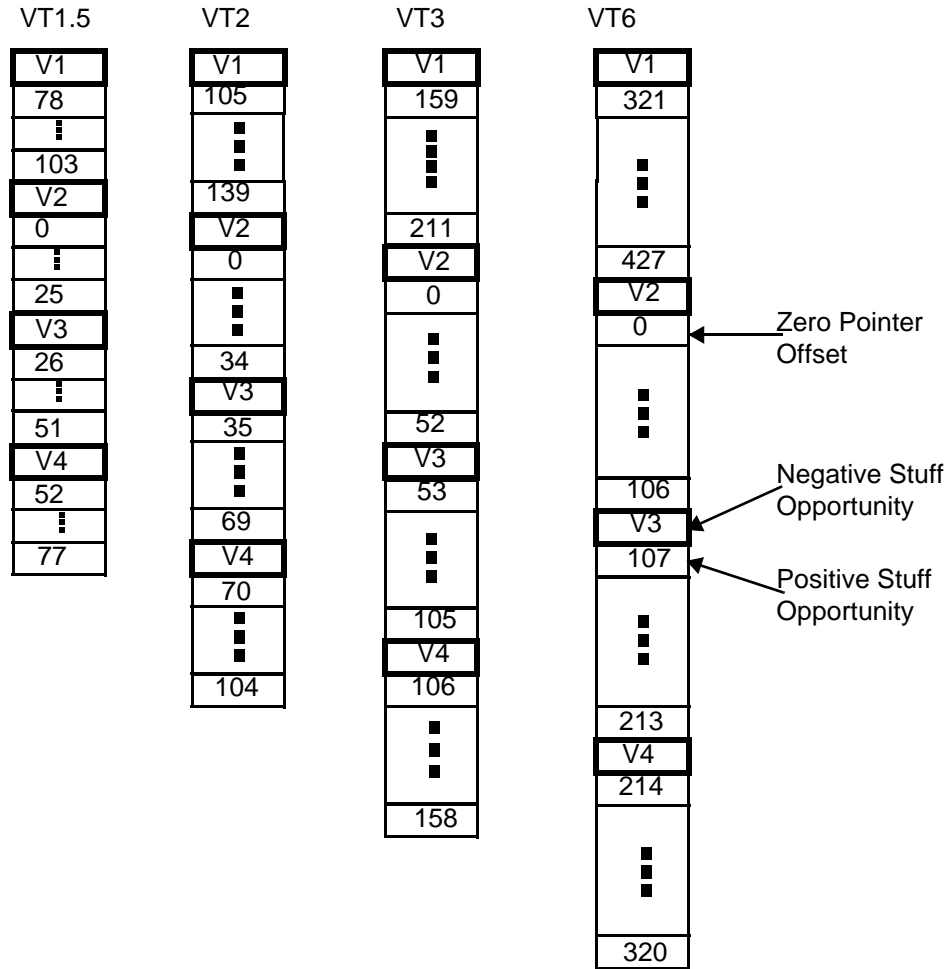


3.5.2.1 VT Pointer Value

Bits 7 through 16 of the pointer word carry the pointer value, which indicates the offset between the pointer word and the first byte of the VT SPE. The V1 through V4 bytes are not counted in the offset. The range of valid offsets is a function of the size of the VT, as shown in [Figure 3-46](#).

R3-120 [98] The pointer value shall be a binary number with a range of 0 to 103 (VT1.5), 0 to 139 (VT2), 0 to 211 (VT3) or 0 to 427 (VT6), and shall indicate the offset between the pointer word and the first byte of the VT SPE (as shown in [Figure 3-46](#)).

Figure 3-46 VT Pointer Offsets



3.5.2.2 VT Frequency Justification

When the frame rate of the VT SPE is greater than or less than that of the STS SPE, the alignment of the VT SPE is periodically slipped forward or back in time (using a negative or positive stuff byte) and the pointer value is adjusted by one.

R3-121 [99] When there is a frequency offset between the frame rate of the STS SPE and that of the VT SPE, the pointer value shall be incremented or decremented as needed (although see rule 6 of **R3-130 [108v2]**), accompanied by a corresponding positive or negative stuff byte.

R3-122 [100] A pointer increment operation shall be indicated by inverting bits 7, 9, 11, 13 and 15 (the I-bits) of the pointer word. The positive stuff byte shall appear immediately after the V3 byte in the VT superframe containing the inverted I-bits, as shown in **Figure 3-46**.

The value contained in the positive stuff byte is undefined.

- R3-123 [101]** A pointer decrement operation shall be indicated by inverting bits 8, 10, 12, 14 and 16 (the D-bits) of the pointer word. The V3 byte shall be used as the negative stuff byte, (i.e., it is used to carry an SPE byte in the VT superframe containing the inverted D-bits), as shown in [Figure 3-46](#).

The following criteria are applicable to the receiving VT pointer processor and provide protection against errors on the I- and D-bits.

- R3-124 [102]** If **O3-125 [103]** (the “8 of 10” objective) is not met, then the increment decision shall be made by a majority vote of the I-bits, and the decrement decision shall be made by a majority vote of the D-bits.
- O3-125 [103]** The increment/decrement decision should be made at the receiver by a match of 8 or more of the 10 I- and D-bits to either the increment or decrement indication.

3.5.2.3 VT Size Indicator

Bits 5 and 6 of the VT pointer word are used to indicate the size of the VT.

- R3-126 [104]** Bits 5 and 6 of the VT payload pointer shall indicate the size of the VT using the code ‘00’ (VT6), ‘01’ (VT3), ‘10’ (VT2) or ‘11’ (VT1.5).

The VT size codes are also listed in [Table 3-8](#), along with the VT pointer value ranges applicable for each size of VT.

Table 3-8 VT Size Indicator

Size Bits 5 and 6	Designation	VT Pointer Range
00	VT6	0 through 427
01	VT3	0 through 211
10	VT2	0 through 139
11	VT1.5	0 through 103

3.5.2.4 New Data Flag

Bits 1 through 4 of the pointer word carry the NDF. The NDF allows an arbitrary change of the pointer value, and also allows an arbitrary change of the size of the VTs in a VT group (i.e., due to a change in the payload). If there is a change in the size of one VT in a VT group, then there is implicitly a simultaneous change in the size of all of the VTs in that group.

-
- R3-127 [105]** A normal NDF shall be indicated (during normal operation) by a '0110' code in the N-bits (see [Figure 3-45](#)). The NDF shall be set by inverting the N-bits to '1001'. The new alignment of the VT SPE shall be indicated by the pointer value accompanying the set NDF and takes effect at the offset indicated.
- R3-128 [106]** The decoding at the pointer processor shall be performed by majority voting (i.e., the NDF shall be detected as being set if three or four of the N-bits match the '1001' code). If a set NDF is detected, then the coincident pointer value shall replace the current value at the offset indicated by the new pointer value.
- R3-129 [107]** If a new size of VT is transmitted, then all 1 to 4 (depending on the new size) of the VT payload pointers in the VT group shall simultaneously indicate a set NDF and the same new size. The new size shall take effect immediately.²⁰
-

3.5.2.5 VT Payload Pointer Generation Rules

The pointer generation criteria from [Sections 3.5.2.1](#) through [3.5.2.4](#) are summarized in this section as a set of pointer generation rules. These rules also contain several additional requirements statements that do not appear in the preceding sections.

-
- R3-130 [108v2]** The VT payload pointer shall be generated according to the following rules.
1. During normal operation, a normal NDF is sent (i.e., the N-bits are set to '0110'), the size bits indicate the size of the VT, and the pointer value locates the start of the VT SPE within the VT envelope capacity.
 2. The pointer value shall only be changed by the operations in rules [3](#), [4](#) or [5](#).
 3. If a positive stuff is needed, the current pointer value is sent with the I-bits inverted, and the subsequent positive stuff opportunity is considered an undefined byte. Subsequent pointers contain the previous pointer value incremented by one.
 4. If a negative stuff is needed, the current pointer value is sent with the D-bits inverted, and the subsequent negative stuff opportunity is overwritten with an SPE byte. Subsequent pointers contain the previous pointer value decremented by one.
 5. If the alignment of the SPE changes for any reason other than rules [3](#) or [4](#), the new pointer value shall be sent accompanied by a set NDF. The set NDF only appears in the first VT superframe that contains the new value. The new SPE begins at the first occurrence of the offset indicated by the new pointer value.
 6. No increment or decrement operation shall be performed for three VT superframes following any of the operations in rules [3](#), [4](#) and [5](#).²¹

²⁰Note that when AIS-V is being transmitted, the entire VT superframe contains all ones, including the N-bits and size bits in the VT pointer word. Since the NDF and size bits are needed in all of the VTs of the new size to indicate to the receiving pointer processor that a new size of VTs is being transmitted, this requirement implies that AIS-V cannot be transmitted in the first VT superframes containing the new VTs.

7. For a nonterminated VT path, an incoming all-ones pointer word shall be regenerated or relayed with no more than a three-superframe delay. When a non-all-ones pointer word is subsequently received, the downstream pointer shall be generated based on the pointer generation and interpretation criteria summarized in this requirement (**R3-130 [108v2]**) and **R3-132 [109]**.²²
8. If the size of the VTs within a VT group is to change, then the NDFs in all of the VTs of the new size (in that VT group) are set simultaneously.

Rule 7 is applicable to STS PTE that processes VT pointers. If all of the STS PTE in a long chain of SONET NEs takes the full three VT superframes allowed in Rule 7, then the detection and termination of an AIS-V defect at the VT PTE could be significantly delayed. Therefore the following objective is applicable.

-
- O3-131 [909]** STS PTE that processes VT pointers should regenerate or relay an incoming all-ones pointer word with no more than a one-superframe delay.
-

Note that STS PTE cannot perform increment or decrement operations while it is performing all-ones pointer relay, and therefore the incoming VT SPE associated with the all-ones pointers cannot be expected to be passed through unaltered (see [Section 6.2.1.2.3](#)).

3.5.2.6 VT Payload Pointer Interpretation Rules

The pointer interpretation criteria from [Sections 3.5.2.1](#) through [3.5.2.4](#) are summarized in this section as a set of pointer interpretation rules. These rules also contain several additional requirements statements that do not appear in the preceding sections.

-
- R3-132 [109]** The VT payload pointer shall be interpreted according to the following rules.
1. During normal operation, the pointer value locates the start of the VT SPE within the VT envelope capacity.
 2. Any variation from the current pointer value shall be ignored unless a consistent new value is received three times consecutively, or the variation is one of the operations in rules 3, 4 or 5. Any consistent new value received three times in succession shall replace the current value at the offset indicated by the new pointer value.

21. Note that in the pointer interpretation rules in the following section, there is no rule or requirement equivalent to pointer generation rule 6. If a pointer processor detects an increment or decrement operation within three VT superframes after another pointer change operation (e.g., due to transmission errors), it can either ignore that operation or interpret it as a valid operation.
22. Since the VT path is not terminated, the AIS-V detection and generation criteria in [Section 6.2.1.2.3](#) are not the applicable criteria for determining the value of the pointer word to be sent downstream.

3. If an increment is detected, then the byte following V3 shall be considered a positive stuff byte, and the current pointer value shall be incremented by one.
4. If a decrement is detected, then the V3 byte shall be considered a negative stuff byte, and the current pointer value shall be decremented by one.
5. If a set NDF is detected, then the coincident pointer value replaces the current value at the offset indicated by the new pointer value.
6. If the equipment has the capability to correctly process different VT sizes based on the received VT size bits, and a set NDF and an arbitrary new size of VT are received simultaneously in all of the VTs within a VT group, then the coincident pointer values and sizes shall replace the current pointer values and sizes at the offsets indicated in the new pointers.
7. If the equipment has the capability to correctly process different VT sizes based on the received VT size bits, then any variation from the current VT size shall be ignored unless consistent valid pointers indicative of a new VT size are received three times consecutively in all of the (new) VTs within a VT group, or the variation is the operation in rule 6. The VT size associated with such pointers received three times in succession shall replace the current size immediately.

As indicated above, the applicability of VT pointer interpretation rules 6 and 7 depend on the functionality of the receiving equipment. Some equipment may be capable of correctly processing different VT sizes based on the received VT size bits. For example, STS PTE that processes the incoming VT pointers and passes VTs (in groups) to other STS PTE within the same NE for transmission on an outgoing SONET signal could detect changes in the VT size bits and change its pointer processing functions accordingly. If such a capability is provided, then rules 6 and 7 would be applicable. Other equipment may not be capable of processing different size VTs correctly using the VT size bits. For example, a SONET ADM that is only equipped to support DS1 and DS3 tributary interfaces would only be expected to be capable of processing VT1.5s. If that ADM received VTs of any other size, it would be expected to declare LOP-V alarms for all of the affected VT1.5s. In such a case, the LOP defect detection criteria in Section 6.2.1.1.3 would be the applicable criteria, and rules 6 and 7 would not apply.

In addition, the phrase “valid pointers indicative of a new VT size” in pointer interpretation rule 7 can be interpreted various ways, and is intended to allow for any existing designs while encouraging the development of new designs that can detect incoming VT size changes that are not detected via rule 6 (e.g., due to transmission errors). The intended interpretation is that the pointer words in at least one of the VTs of the new size should contain normal NDFs, consistent and valid size bits, and a constant in-range pointer value (i.e., a normal pointer word) for three consecutive VT superframes. The remaining VTs could have either consistent normal pointer words or consistent all-ones pointer words (i.e., AIS) in their corresponding VT superframes.

4 Physical Layer

This section describes optical and electrical parameters that enable multi-supplier compatibility at the Physical layer for SONET signals with bit rates up through OC-768. [Section 4.1](#) provides an overview of the Physical layer classifications, while [Section 4.2](#) contains the optical parameter requirements for the application categories defined in [Section 4.1](#). [Section 4.3](#) describes a methodology that may be used for engineering a single-mode fiber optic system, and [Section 4.4](#) contains the SONET electrical interface specifications.

4.1 Physical Layer Classifications

SONET signals may be transported by either electrical or optical means, and SONET NEs may have either electrical or optical interfaces (or both). While technical considerations limit the feasibility of electrical transport to short distances and relatively low bit rates, the flexibility of optical transmission allows a wide range of possible applications and corresponding implementations. However, to simplify the development of compatible multi-supplier SONET optical systems, it is desirable to define a relatively small set of application categories and corresponding sets of optical interface specifications. This allows for some degree of optical device commonality among the various applications and a balancing of economic and technical considerations in developing compatible multi-supplier SONET equipment.

In this document, the following broad application categories are used:

- Short Reach (SR) optical interfaces refer to optical sections having system loss budgets from 0 dB to 4 or 7 dB. Depending on the SONET hierarchical level, SR transmitters may be either Light-Emitting Diodes (LEDs), low-power (e.g., 50 μ W or -13 dBm) Multi-Longitudinal Mode (MLM) lasers or, in the case of very high bit-rate signals, high-power (e.g., 500 μ W or -3 dBm) Single Longitudinal Mode (SLM) lasers.
- Intermediate Reach (IR) optical interfaces refer to optical sections with system loss budgets from 0, 3 or 6 dB to 11 or 12 dB. Typically, low-power SLM or MLM laser transmitters are used at the lower bit rates, while high-power SLM lasers and (in a few cases) optical amplifiers are used at the higher bit rates.
- Long Reach (LR) optical interfaces refer to optical sections with system loss budgets from 10, 11 or 16 dB to 22, 24 or 28 dB (depending on the bit rate). Typical of long-haul telecommunications systems, LR interfaces are based on high-power MLM or SLM lasers. In addition, at the OC-192 and OC-768 bit rates optical amplifiers are generally utilized at the transmitter and/or receiver, and some form of dispersion accommodation may be necessary.
- Very-long Reach (VR) optical interfaces refer to optical sections with system loss budgets from 22 to 33 dB. Also typical of long-haul telecommunications systems, VR interfaces (which are currently defined only at bit rates of OC-12, OC-48 and OC-192) are based on SLM lasers, may utilize optical amplifiers at the transmitter and/or receiver, and at the OC-192 bit rate may require some form of dispersion accommodation.

- Ultra-long Reach (UR) optical interfaces refer to optical sections with system loss budgets from 33 to 44 dB. Similar to LR and VR interfaces, UR interfaces (which are currently defined only at the OC-12 and OC-48 bit rates) are based on SLM lasers, and utilize optical amplifiers at the transmitter and/or receiver.
- Electrical interfaces, compatible with DSX-type interfaces for the asynchronous hierarchy, are defined for the STS-1 and STS-3 levels.

In addition, for each of the broad optical application categories defined above, it is possible to consider the use of transmitters operating in either the 1310-nm or 1550-nm wavelength region, and either dispersion-unshifted single-mode fiber [also referred to as EIA/TIA Class IVa fiber, G.652 fiber or Conventional Single Mode Fiber (C-SMF)] or Dispersion-Shifted single-mode Fiber (DSF, also referred to as EIA/TIA Class IVb or G.653 fiber).¹ In this document, optical parameters are specified for most levels of the current SONET hierarchy (i.e., all OC-N levels from OC-1 through OC-768, with the exception of OC-24) in three to five of the broad application categories defined above, and for various combinations of wavelengths and fiber types. These combinations are summarized in [Table 4-1](#), which for the user's convenience also indicates if a particular application is assumed to include the use of one or more optical amplifiers, and if any type of dispersion accommodation is required. In the applications that use optical amplifiers, those amplifiers may be any of the following types.

- A separate unit placed either after the optical transmitter in a stand-alone Booster Amplifier (BA) configuration, or prior to the optical receiver in a stand-alone Pre-Amplifier (PA) configuration
- A BA integrated into the transmitter package to produce an Optically Amplified Transmitter (OAT)
- A PA integrated into the receiver package to produce an Optically Amplified Receiver (OAR).

In addition, optical amplifiers can be placed at intermediate locations, in which case they would typically be referred to as Line Amplifiers (LAs). However, at this time the applications listed in [Table 4-1](#) are limited to systems without LAs. Application codes for systems with LAs may be included in a future issue of this document.

1. GR-20-CORE, *Generic Requirements for Optical Fiber and Optical Fiber Cable*, contains generic requirements for optical fiber and optical fiber cables. In addition, applications that utilize multimode fiber as defined in ITU-T Rec. G.651, *Characteristics of a 50/125 μ m multimode graded index optical fibre cable*, or Non-Zero Dispersion Shifted Fiber (NZ-DSF) as defined in ITU-T Rec. G.655, *Characteristics of a non-zero dispersion-shifted single-mode optical fibre and cable*, are specified in several standards documents but are not considered here.

Table 4-1 Application Categories

Application	Nominal Wavelength (nm)	Fiber Type	Applicable Bit Rate (OC-X)	Dispersion Accommodation Method	Type of Amplification Assumed
SR-1	1310	C-SMF	All ^a	-	-
SR-2	1550	C-SMF	192, 768	-	-
IR-1	1310	C-SMF	1 to 192 ^b	-	-
IR-2	1550	C-SMF	1 to 192 ^b	-	-
			768	PDC	PA/OAR
IR-3	1550	DSF	192	-	-
			768	PDC	PA/OAR
LR-1	1310	C-SMF	1 to 192 ^b	-	-
LR-2	1550	C-SMF	1, 3, 12, 48	-	-
			192 (LR-2)	PCH ^c	-
			192 (LR-2a)	PDC ^d	PA/OAR
			192 (LR-2b)	SPM ^e	BA/OAT
			192 (LR-2c)	PCH	PA/OAR
LR-3	1550	DSF	768	PDC	BA/OAT, PA/OAR
			1, 3, 12, 48	-	-
			192	-	BA/OAT
VR-1	1310	C-SMF	768	PDC	BA/OAT, PA/OAR
			12	-	PA/OAR
VR-2	1550	C-SMF	12	-	PA/OAR
			48	-	BA/OAT
			192 (VR-2a)	PDC	BA/OAT, PA/OAR
			192 (VR-2b)	PDC and SPM	BA/OAT, PA/OAR
VR-3	1550	DSF	12	-	PA/OAR
			48	-	BA/OAT
			192	-	BA/OAT, PA/OAR
UR-2	1550	C-SMF	12, 48	-	BA/OAT, PA/OAR
UR-3	1550	DSF	12, 48	-	BA/OAT, PA/OAR

Notes:

- Not Applicable
- a. X = 1, 3, 12, 48, 192 and 768
- b. X = 1, 3, 12, 48 and 192
- c. Pre-chirp (PCH)
- d. Passive Dispersion Compensation (PDC)
- e. Self Phase Modulation (SPM)

In general, the optical parameter criteria in this document are based on specifications contained in the following standards documents:

- ANSI T1.105.06, *Synchronous Optical Network (SONET): Physical Layer Specifications*
- ITU-T Rec. G.691, *Optical interfaces for single channel STM-64 and other SDH systems with optical amplifiers*
- ITU-T Rec. G.693, *Optical interfaces for intra-office systems*
- ITU-T Re. G.957, *Optical interfaces for equipments and systems relating to the synchronous digital hierarchy*
- ITU-T Rec. G.959.1, *Optical transport network physical layer interfaces.*

In addition, it should be noted that the standards listed above also contain specifications for a number of other applications and/or hierarchical levels that have not been included here. In particular:

- ANSI T1.105.06 includes (via references to ANSI T1.646, *Network and Customer Installation Interfaces - Broadband ISDN: Common Criteria*) specifications for OC-1 and OC-3 SR-0 applications that utilize a nominal 1310-nm source wavelength and multimode fiber
- ITU-T Rec. G.693 provides specifications for a number of STM-64/OC-192 Very Short Reach (VSR) applications, as well as several applications whose primary differences from the OC-192 SR applications defined in this document are that reduced maximum chromatic dispersion values or different types of fibers are specified
- ITU-T Rec. G.693 also defines a number of STM-256/OC-768 SR applications that are specified to utilize different types of fibers, or for which larger values of the maximum optical path attenuation parameter and either higher transmitter output power levels or more sensitive receivers (but the same maximum chromatic dispersion values) are specified
- Rec. G.959.1 defines a number of additional single-channel Non-Return to Zero 10 Gb/s [NRZ 10G, which includes Optical channel Transport Unit 2 (OTU2) and STM-64/OC-192] and NRZ 40G (i.e., OTU3 and STM-256/OC-768) applications, including:
 - Two whose primary differences from the OC-192 IR-2 and IR-3 applications provided in this document are that slightly lower transmitted and received power levels [appropriate for Avalanche Photodiode (APD) receivers instead of Positive-Intrinsic-Negative (PIN) receivers] are specified
 - An NRZ 10G SR application for 1550-nm signals carried on G.653 fibers (i.e., an OC-192 SR-3 application)
 - Several NRZ 10G and NRZ 40G applications for 1550-nm signals carried on G.655 fibers.
- ITU-T Rec. G.959.1 also provides specifications for several single-channel NRZ 2.5 G (i.e., OTU1 and STM-16/OC-48) applications that are similar to OC-48 applications provided in this document, but with modifications to certain parameter values in support of a system BER objective of 10^{-12} instead of 10^{-10} (see [Section 4.2.1](#))

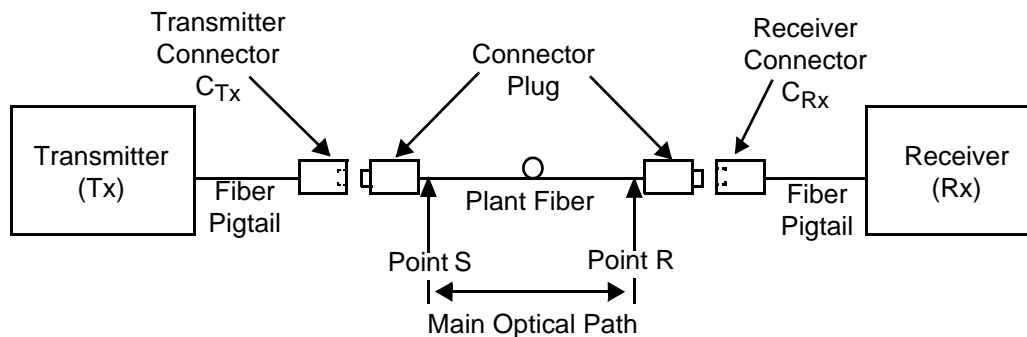
In addition, ANSI has defined optical parameters for the OC-24 and Optical VT Group (OVTG) rates, and electrical parameters for the VT1.5 rate, while ITU-T Rec. G.959.1 contains specifications for several multi-channel interfaces and single-channel Return to Zero 40 Gb/s (RZ 40G) applications. The need to include some or all of these additional applications and hierarchical levels in future issues of this document is under study.

4.2 Optical Parameter Definitions and Interface Requirements

4.2.1 General

To specify SONET optical parameters for the various application categories, optical system interfaces can be represented as shown in Figure 4-1. Point S is a reference point on the optical fiber just after the transmitter (Tx) optical connector (C_{Tx}), and Point R is a reference point on the optical fiber just before the receiver (Rx) optical connector (C_{Rx}). Points S and R provide a convenient separation of the optical link into a transmitter subsection, a receiver subsection, and an optical connection subsection. Optical parameters are specified for the transmitter at point S, for the receiver at point R, and for the optical connection between points S and R. Any additional connectors within the fiber plant, such as those at a fiber distribution frame, are considered part of the main optical path. Note that in cases where a BA is utilized, point S is defined to be downstream of that BA. Similarly, if a PA is utilized, point R is defined to be upstream of that PA. In addition, in cases where a PDC component is utilized (see Section 4.2.5.3.3.A), that component is assumed to be located upstream of point S or downstream of point R.

Figure 4-1 Optical System Interfaces (Points S and R)



All of the parameter values specified are worst-case, end-of-life values and are to be met over the ranges of operating conditions described in Section 7.1.1.² The parameters are specified relative to an optical system design objective of a BER not worse than 1×10^{-10} (OC-1 through OC-48 SR, IR and LR applications) or 1×10^{-12} (all remaining applications) for the extreme case of optical path attenuation and dispersion conditions for each application. Separate equipment margins are not specified, and it is assumed that transmitters, receivers and the optical cable plant individually meet the specifications that appear in the following sections. The intent of these interface specifications is to enable multi-supplier compatibility within

application codes. In some cases, the use of these specifications may lead to more conservative optical section designs than could be obtained through supplier-proprietary interfaces, the use of statistical design approaches for engineering an optical link, or in environments more constrained than those that [Section 7.1.1](#) describes. Note that in [Tables 4-6](#) through [4-25](#) (in which the various parameter values for each application are listed) parameters that are not specified or that are not appropriate are indicated with the notation “NA” (Not Applicable).

4.2.2 Optical Line Coding

R4-1 [110] For all SONET optical system interfaces described, binary NRZ optical line coding shall be used.

RZ optical line coding may have some advantages for certain applications, particularly at very high bit rates (e.g., OC-768), but is currently For Further Study (FFS).

4.2.3 Transmitter

This section contains definitions and criteria covering such transmitter parameters as the ranges of operating wavelengths or frequencies and coupled transmitter output powers, maximum spectral widths, and minimum extinction and side-mode suppression ratios. In addition, it includes criteria that are primarily applicable to OC-192 and OC-768 transmitters concerning source frequency chirp factors and maximum spectral power density. Finally, transmitter pulse shapes are specified by a mask for the eye diagram at point S.

In addition to being applicable under the specified worst-case operating conditions and at the transmitter’s end-of-life (see [Section 4.2.1](#)), the criteria referred to above and contained in the following subsections apply to a fully modulated transmitter that is exposed to the specified worst-case Optical Return Loss (ORL) at point S. While it is not expected to occur in the network, for testing purposes this includes the case in which all of the returned optical power occurs as a result of a single-point reflection. In addition, see [Section 4.2.5.1](#) for a related objective concerning the impact of a “worse than worst-case” reflection on system performance.

Depending on the attenuation and dispersion characteristics and the hierarchical level of an application, feasible transmitter devices may include LEDs, MLM lasers and/or SLM lasers. In addition, an SLM laser may be either directly modulated or

-
2. While these criteria apply at the equipment’s “end of life”, they are often used in the evaluation of new (i.e., beginning-of-life) equipment. In some cases the measured parameter values are not expected to vary appreciably over time, and in those cases allowances for aging are not a significant issue. However, in other cases significant changes are likely to occur. For example, the difference between the sensitivity of a new receiver and its end-of-life counterpart is typically expected to be in the range of 2 to 4 dB. Therefore, results of a sensitivity test performed on a new receiver that are 0 to 4 dB “better” than the specified value may indicate that although it (the new receiver) currently “conforms” to the requirement, it might not meet that requirement at some point within its nominal lifetime.

indirectly modulated [e.g., via an external Mach Zehnder (MZ) or Electro-Absorption (EA) modulator]. This document indicates one or more nominal device types or modulation methods for each application, and in some cases where two device types are listed, the applicable values of certain parameters depend on the particular type of device that is utilized.³ However, this indication is not intended to be a requirement on the source type or modulation method, and any device meeting the transmitter characteristics specified may be substituted for the nominal device type without degradation in system performance. For example, an SLM laser can (assuming it meets the applicable criteria) be used in any application where an LED or MLM laser is indicated.

4.2.3.1 Spectral Characteristics

This section provides definitions and criteria concerned with the parameters listed in Tables 4-6 through 4-25 that are related to the spectral characteristics of the optical signal.

λ_T – The central operating wavelength. This is the wavelength at which the effective optical power resides (e.g., see TIA-455-127, *FOTP-127, Spectral Characterization of Multimode Laser Diodes*).

λ_{Tmin} , λ_{Tmax} – The minimum and maximum central operating wavelengths. These values define the allowable range of transmitter central wavelengths, and are generally centered around some nominal value (λ_{Tnom}).

ν_{nom} – The nominal central frequency. This parameter is used along with $\Delta\nu_{max}$ (and instead of the central wavelength parameters listed above) to define the acceptable operating wavelength/frequency range for OC-768 signals. In general, a value of ν_{nom} can be converted to a λ_{Tnom} value using the equation:

$$\nu = c/\lambda \quad (4-1)$$

where c is the speed of light in a vacuum (2.99792458×10^8 m/s).

$\Delta\nu_{max}$ – The maximum central frequency deviation. This parameter defines the allowable difference between the nominal central frequency and the actual central frequency, and includes all of the processes that affect the instantaneous value of the central frequency over a measurement interval appropriate to the bit rate of the signal (e.g., source chirp, information bandwidth, broadening due to self-phase modulation).

$\Delta\lambda_{rms}$ – The root-mean-square (rms) spectral width. This is the spectral width specified for LEDs and MLM lasers (see TIA-455-127), and is defined to include modes out to and including those 20 dB down from the peak mode. $\Delta\lambda_{rms}$ does not apply to SLM transmitters.

3. In cases where the source type is listed as SLM, the modulation method may be either direct or indirect, with direct modulation expected to predominate at the relatively lower bit rates and shorter target distances (see Table 4-4). Also, in cases where Indirect Modulation (IM) is listed, the source type is implied to be SLM.

$\Delta\lambda_{20}$ – The full spectral width measured 20 dB down from the maximum of the central wavelength peak. This is the spectral width specified for SLM transmitters. $\Delta\lambda_{20}$ does not apply to MLM lasers or to LEDs.

$SMSR_{min}$ – Minimum acceptable value of the Side-Mode Suppression Ratio (*SMSR*). *SMSR* is defined as the ratio (in dB) of the average optical power in the dominant longitudinal mode (*M1*) of an SLM laser to the optical power in the most significant side mode (*M2*), and is calculated as follows:

$$SMSR = 10\log_{10} \left(\frac{M1}{M2} \right) \text{dB.} \quad (4-2)$$

α – The source frequency chirp factor (α). This parameter is defined as follows:

$$\alpha = \frac{\frac{d\phi}{dt}}{\frac{1}{2P} \left(\frac{dP}{dt} \right)} \quad (4-3)$$

where ϕ is the optical phase of the signal and P is the signal power. In general, α is a function of time and is (expected to be) specified in terms of a range of acceptable values when measured at a particular time (i.e., during a particular transition that occurs periodically in the SONET signal). Additional information related to chirp can be found in Appendix IV of ITU-T Rec. G.691.

SPD_{max} – Maximum acceptable value of the Spectral Power Density (*SPD*). This is the highest time-averaged optical power within any 10 MHz interval in the signal spectrum, and is limited in order to control Brillouin scattering for high-power sources with narrow linewidths. Sources with narrow linewidths are common for OC-192 and OC-768 systems (which often require indirectly modulated SLM lasers in order to minimize chromatic dispersion).

-
- R4-2 [115v3]** The central wavelength and chirp factor (if applicable) of the transmitter shall be within the appropriate ranges listed in Tables 4-6 through 4-25.
- R4-3 [1157]** The central frequency of an IR or LR OC-768 transmitter shall be within the range defined by the values of the nominal central frequency and maximum central frequency deviation parameters listed in Table 4-14 or 4-20.
- R4-4 [116v2]** The spectral width and maximum spectral power density (if applicable) of the transmitter shall be less than or equal to the appropriate values listed in Tables 4-6 through 4-25.
- R4-5 [117v3]** The side-mode suppression ratio (if applicable) of the transmitter shall be greater than or equal to the appropriate value listed in Tables 4-6 through 4-25.
-

Note that for some applications that are expected to utilize MLM sources (e.g., OC-12 IR-1 in Table 4-11), two possible central wavelength ranges and corresponding maximum spectral width and D_{SRmax} (see Section 4.2.5.3) values are given. Compared to the typical case where a single set of these three parameters is

specified, this is expected to increase the portion of the manufactured sources that can be used. That is, it permits the use of transmitters having central wavelengths that are relatively far from the zero-dispersion wavelength of the applicable fiber (in cases where the spectral width is comparatively small), as well as transmitters having relatively wide spectral widths (in cases where the central wavelength is comparatively close to the fiber's zero-dispersion wavelength).

4.2.3.2 Coupled Transmit Power

P_{Tmax} , P_{Tmin} – The maximum and minimum coupled transmitter powers. These values define the allowed range of average optical power, P_T , at point S for a pseudo-random data sequence (see TIA-526.2, *OFSTP-2, Effective Transmitter Output Power Coupled into Single-Mode Fiber Optic Cable*).

- R4-6 [1158]** The average optical power coupled from the transmitter into the fiber (and appearing at point S) shall be within the appropriate range listed in Tables 4-6 through 4-25.

4.2.3.3 Extinction Ratio

r_{emin} – Minimum extinction ratio, r_e . This is the minimum acceptable value of the ratio (in dB) of the average optical power at the center of a logical one (A) to the average optical power at the center of a logical zero (B).

$$r_e = 10 \log_{10} \left[\frac{A}{B} \right] \text{ dB.} \quad (4-4)$$


- R4-7 [1159]** The extinction ratio of the transmitter shall be greater than or equal to the appropriate value listed in Tables 4-6 through 4-25.

4.2.3.4 Mask of the Eye Diagram

This section contains the eye diagram requirement that applies to SONET transmitters, and discusses the mask and measurement filter used in determining a transmitter's conformance to that requirement. In addition, it contains related information and criteria regarding the rise times of pulses in very high power signals.


- R4-8 [119v3]** Transmit pulses, referenced to a noise filter with transfer characteristics as given in Equation 4-5, shall fall within the eye diagram mask as defined in Figure 4-2 or 4-3 (as appropriate).

Figure 4-2 SONET Eye Diagram Mask (OC-1 through OC-12)



Rates	X1	X2	Y1
OC-1 and OC-3	0.15	0.35	0.20
OC-12	0.25	0.40	0.20

Figure 4-3 SONET Eye Diagram Mask (OC-48 through OC-768)



Rates	$\Delta X1$	Y1	Y2	Y3	Y4
OC-48	0.20	0.25	0.75	0.25	0.25
OC-192 unamplified, 1310 nm	0.20	0.25	0.75	0.40	0.25
OC-192 unamplified, 1550 nm	0.20	0.25	0.75	0.25	0.25
OC-192 (a and c) ¹	FFS	FFS	FFS	FFS	FFS
OC-192 (none and b) ^{1, 2}	0.20	$\Delta+0.25$	$\Delta+0.75$	0.25	0.25
OC-768	0.20	0.25	0.75	0.25	0.25

1. “a and c” and “none and b” refer to the dispersion accommodation technique used in amplified OC-192 systems. Thus, the “a and c” entries apply to the LR-2a, LR-2c and VR-2a application categories, and the “none and b” entries apply to the LR-2b, LR-3, VR-2b and VR-3 application categories.
2. With Δ variable, $-0.25 < \Delta < +0.25$

Note that in evaluating an eye diagram result, it is important to consider not only the eye opening, but also the overshoot and undershoot limitations of the transmitted signal. Also note that for systems employing dispersion accommodation techniques based on predistortion of the signal (which are currently defined only at the OC-192 bit rate), the eye mask can only be defined between points with undistorted signals. These points, however, do not coincide with the main optical path interfaces, and thus may not be accessible. As listed in table in Figure 4-3, the parameter values to be used in the eye mask definition for such systems are FFS.

As indicated in **R4-8 [119v3]**, a transmitter's eye diagram is determined with respect to a low-pass reference filter (which need not represent the noise filter of the transmission equipment optical receiver). The specified reference filter for qualifying transmitted pulses through an eye diagram measurement is a fourth-order Bessel-Thomson filter with a transfer function given by:

$$H(p) = (1/105) [105 + 105y + 45 y^2 + 10 y^3 + y^4] \quad (4-5)$$

where,

$$\begin{aligned} p &= j \omega / \omega_r, \\ y &= 2.1140p, \\ \omega_r &= 2\pi f_r, \\ f_r &= 0.75 f_0, \\ f_0 &= \text{bit rate.} \end{aligned}$$

For this filter, the nominal attenuation at the reference frequency, f_r , is 3 dB. The corresponding attenuation and group delay distortion at various frequencies are given in **Table 4-2**. For OC-192 and OC-768, the optical reference receiver function is defined as the total frequency response of any combination of photodetector, low-pass filter and oscilloscope functional elements, together with any interconnection of those elements. The allowable deviation of the nominal attenuation values to account for manufacturing tolerances of the optical reference receiver components are indicated in **Table 4-3**.

Table 4-2 Attenuation and Group Delay Distortion as a Function of Frequency

f/f_0	f/f_r	Attenuation (dB)	Group Delay Distortion (UI)
0.15	0.2	0.1	0
0.3	0.4	0.4	0
0.45	0.6	1.0	0
0.6	0.8	1.9	0.002
0.75	1.0	3.0	0.008
0.9	1.2	4.5	0.025
1.0	1.33	5.7	0.044
1.05	1.4	6.4	0.055
1.2	1.6	8.5	0.10
1.35	1.8	10.9	0.14
1.5	2.0	13.4	0.19
2.0	2.67	21.5	0.30

Table 4-3 Tolerance Values of the Attenuation of the Optical Reference Receiver

Bit Rate	Tolerance, Δa (dB)	
	For f/f_r in the Range of 0.001 to 1.0	For f/f_r in the Range of 1.0 to 2.0
OC-1	± 0.3	± 0.3 to $\pm 2.0^a$
OC-3	± 0.3	± 0.3 to $\pm 2.0^a$
OC-12	± 0.3	± 0.3 to $\pm 2.0^a$
OC-48	± 0.5	± 0.5 to $\pm 3.0^a$
OC-192	± 0.85	± 0.85 to $\pm 4.0^a$
OC-768	FFS	FFS

Notes:

- a. Intermediate values of Δa should be interpolated linearly on a logarithmic frequency scale.

In addition to the eye diagram mask discussed above, for systems that are subject to SPM due to high power levels, specification of the minimum rise time is necessary to avoid the possibility of SPM breakdown.

R4-9 [1106] The rise time (10% to 90% value of the single pulse) for systems with transmitter power levels of +12 to +15 dBm shall be greater than or equal to 30 ps.

For lower power levels such as +10 to +13 dBm, this value and the interaction with the signal chirp is FFS.

4.2.4 Receiver

Receiver characteristics specified in this section are receiver sensitivity (P_{Rmin}) and overload (P_{Rmax}), receiver reflectance, and the allowable power penalty from the combined effects of dispersion, multiple reflections and jitter.

4.2.4.1 Receiver Sensitivity and Overload

P_{Rmin} , P_{Rmax} – The receiver sensitivity and receiver overload optical power levels. These values define the minimum and maximum levels of the average received power at point R, P_R , to achieve a 10^{-10} BER for SR, IR and LR applications at rates up through OC-48 or a 10^{-12} BER for all remaining applications (see note “a” in Tables 4-6 through 4-25).

In general, the specified limits on P_{Rmax} and P_{Rmin} must be met without FEC, and in the presence of receiver connector degradations. They must also be met for any received signal that has eye diagram, extinction ratio and spectral characteristics that meet the criteria for the receiver’s particular application (including a signal with all “worst-case” characteristics), and for a received signal that is jittered at any level up to that allowed by the applicable jitter generation requirement (i.e.,

R5-273 [357v2]). On the other hand, they do not have to be met in the presence of dispersion or multiple reflections in the optical path, or jitter in excess of that allowed by **R5-273 [357v2]**. These latter effects are addressed separately, in the allocation of the maximum optical path power penalty and the jitter tolerance criteria.

R4-10 [120v2] The sensitivity of an OC-N receiver shall be better than or equal to the applicable value of P_{Rmin} given in Tables 4-6 through 4-25 [i.e., the measured value of P_{Rmin} (in units of dBm) shall be at least as negative as the listed value].

R4-11 [122v2] The overload power level of an OC-N receiver shall equal or exceed the applicable value of P_{Rmax} given in Tables 4-6 through 4-25.

In addition to the aging issue discussed in footnote 2 on page 4–6, various characteristics of the signal used during receiver sensitivity and overload tests need to be considered in evaluating the results of those tests. The reason for this is that it is difficult to generate signals with the worst-case characteristics discussed above, and therefore many receivers are likely to be tested using signals having “better than worst-case” characteristics. For some parameters the impact of this is expected to be minor (e.g., in the absence of dispersion, the spectral width and central wavelength of the test signal are expected to have relatively little impact in most applications). However, for other parameters it can mean that the measured sensitivity or overload power needs to be better than the listed value of P_{Rmin} or P_{Rmax} by some non-negligible margin. This is discussed in limited detail in Annex A of ITU-T Rec. G.691 for the case of a receiver sensitivity test signal with a better than worst-case eye diagram and extinction ratio.

For most of the SONET signal rates and the application categories defined in Section 4.1, a receiver’s conformance to **R4-11 [122v2]** permits worst-case engineering of the optical link without the use of optical attenuators for the extreme cases of minimum attenuation and maximum output power. Note however, that the use of attenuators may be necessary in certain applications. In addition, to reduce the chance of damage to an optical receiver in situations where an optical transmitter might be connected to the receiver with minimal attenuation in the optical path (e.g., in a laboratory situation where a short fiber is used to loop back an NE’s output OC-N signal), the following requirement applies.

R4-12 [1107v2] Optical receivers available after the release of Issue 3 of this document (in September 2000) shall meet one of the following:

- The receiver shall be able to tolerate continuous exposure to an optical input signal having a power level equal to the maximum transmitter output power allowed for that application (P_{Tmax}) plus at least 1 dB of margin.
- The receiver shall be labeled to indicate the maximum optical input power level to which it can be continuously exposed without damage.

In **R4-12 [1107v2]**, to tolerate or not be damaged by continuous exposure to a signal with a particular high power level means that the receiver does not display any apparent degradation in its receiver sensitivity and overload power levels after receiving the signal for some test period. (For example, if the receiver sensitivity

was measured to be -34.1 dBm before a 5-minute high power test period, then it should be approximately -34.1 dBm after that test period.) It does not mean that the receiver has to be able to accept and process that signal.

4.2.4.2 Receiver Reflectance

Maximum Receiver Reflectance – The maximum reflectance allowable from the receiver as measured at reference point R.

R4-13 [112] The receiver reflectance shall be less than (more negative than) the value listed under “Max. Receiver Reflectance” in Tables 4-6 through 4-25.

The receiver reflectance value of -27 dB that is specified for some applications in Tables 4-6 through 4-25 is intended to ensure acceptable penalties due to multiple reflections for all likely system configurations involving components such as multiple connectors and splices. Systems employing fewer or higher-performance optical components produce fewer multiple reflections and, consequently, are able to tolerate receivers having higher reflectance values. As an extreme example, if only two connectors exist in the system, a -14 dB receiver reflectance is considered acceptable. Therefore, at the option of each network provider and depending on the specific application, maximum receiver reflectance values lower (more negative) than -14 dB in Tables 4-6 through 4-25 may be increased to as high as -14 dB.

4.2.4.3 Optical Power Penalty

An optical power penalty is an apparent change in a receiver’s sensitivity as a result of distortions to the optical signal waveform. In most cases it is manifested as a shift of the receiver’s BER versus received power curve towards higher input power levels. This is defined as a positive power penalty. Negative power penalties may occur in some situations, but are generally small. (A negative penalty indicates that a less-than-perfect waveform has been improved in some way by the distortions.) Ideally, the BER curve is simply translated; however, shape variations can occur and may indicate the emergence of a BER floor.

This section is concerned with optical power penalties that occur as a result of some baseline signal’s transmission over a “conforming” optical path.⁴ It is not concerned with the optical power penalty that might be observed if a system were to be tested with and without the transmitter exposed to the specified worst-case ORL at point S. That particular case is covered by the combination of the criteria in Section 4.2.4.1 that indicate that a receiver must display some minimum sensitivity when it is receiving the worst-case signal that a transmitter is allowed to generate, and those

-
4. Ideally, the characteristics of the baseline signal used in the power penalty tests discussed in this section would be the same as those specified for the signal that is supposed to be used in determining the conformance of a receiver to its sensitivity requirement (i.e., a “worst-case conforming” signal, see Section 4.2.4.1). However, it is recognized that this may not always be possible. In such cases, the power penalties need to be calculated using the apparent sensitivities of the receiver measured when the incoming signals are undistorted and distorted versions of the same baseline signal.

in Section 4.2.3 that indicate that a transmitter must generate a worst-case or better signal in the presence of the minimum-allowed ORL. Thus, degradations to the transmitter's output that are caused by the minimum-allowed ORL are supposed to be accounted for in the waveform used in the receiver sensitivity tests. In addition, this section is not concerned with the optical power penalty that might occur as a result of a worse-than-worst-case reflection (see O4-18 [114v2] in Section 4.2.5.1).

Among the phenomena that can cause degradation of an optical signal as it traverses an optical path are chromatic dispersion, multiple reflections and Polarization Mode Dispersion (PMD). Additional information related to each of these processes appears in Section 4.2.5, and in general their impacts can appear at the receiver in forms such as pulse broadening and increased jitter.⁵

P_O – Optical path power penalty (in dB). P_O accounts for the total degradation of the optical waveform along the optical path between points S and R.

In actual implementations it is not expected that the worst-case conditions for each type of disturbing process will occur simultaneously. Therefore, for purposes of conformance verification, the power penalties observed at a receiver are required to be less than the specified values under the worst-case conditions for each process tested separately. Note that in most applications P_O is specified to be 1 dB and is applicable for each type of disturbing process. However in some high dispersion applications P_O is specified to be 2 or 3 dB. In those cases the 2 or 3 dB value is applicable for chromatic dispersion, while the maximum allowable power penalty in the other cases is set at 1 dB.

-
- R4-14 [121]** An OC-N receiver shall display a power penalty of no more than P_O in the presence of the worst-case chromatic dispersion and maximum chromatic dispersion deviation (if applicable) for that application (based on the values specified in Tables 4-6 through 4-25).
- R4-15 [113v2]** An OC-N receiver shall display a power penalty of no more than 1.0 dB in the presence of the worst-case combination of multiple discrete reflectances (including receiver reflectance) and minimum optical path attenuation (based on the values specified in Tables 4-6 through 4-25).
- R4-16 [1108v2]** An OC-N receiver shall display a power penalty of no more than 1.0 dB in the presence of the worst-case Digital Group Delay (i.e., DGD_{max}) values specified in Tables 4-6 through 4-25.
-

Since power penalties are defined in terms of the changes in the receiver's apparent sensitivity, they need to be determined at the BER at which sensitivity is defined for that receiver (i.e., at 1×10^{-10} or 1×10^{-12} , depending on the application). Also, the margin by which a receiver meets its receiver sensitivity requirement could be a useful factor in evaluating the importance or impact of any relatively large power penalty result. For example, a receiver that meets the receiver sensitivity requirement with a margin of 8 dB and displays a nonconforming chromatic

5. The jitter on a signal may also increase beyond the level at which the receiver's sensitivity is supposed to be tested as a result of jitter enhancement or accumulation (e.g., through a string of regenerators).

dispersion power penalty of 1.5 dB could be considered “better” for use in at least some applications than a receiver that meets the receiver sensitivity requirement by a margin of 3 dB and has a (conforming) chromatic dispersion power penalty of 0.5 dB.

4.2.5 Optical Path

The primary parameters associated with the optical path are reflections and ORL, attenuation, chromatic dispersion, and DGD caused by PMD. Typically, the first three of these are important for all systems, while DGD is primarily significant for OC-192 and OC-768 systems. In addition, for high input powers such as those that occur in systems utilizing BAs/OATs, nonlinear effects may occur and degrade (or in certain cases, improve) the system performance. In those cases the potential effects include SPM, Stimulated Raman Scattering (SRS) and Stimulated Brillouin Scattering (SBS).

In general, the optical path parameter values in Tables 4-6 through 4-25 are expected to be sufficient to cover current and near-term future SONET installations in a typical network for intra-office, inter-office short-haul (including loop feeder) and inter-office long/very long/ultra long-haul applications. These parameter values serve a dual purpose of setting limits on the optical path that a system is expected to operate over and specifying the optical path characteristics that the system must tolerate (see Section 4.2.4.3).

4.2.5.1 Reflections and ORL

Refractive index discontinuities along the optical path may cause reflections. If not controlled, reflections can degrade system performance through their disturbing effect on the operation of the laser or through multiple reflections that lead to interferometric noise at the receiver. While the performance degradation from single-point reflections on the laser can be reduced by optical isolation of the transmitter, multiple reflection effects are much more difficult to predict and control because they depend in a complicated fashion on the number, spacing and reflectance (R) values of the discrete reflectors as well as on the optical source type and its spectral characteristics. In general, reflection-induced degradations increase with system bit rate and optical source coherence, and minimizing individual reflectances reduces the impact of both single and multiple reflections on system performance. By enforcing reflectance requirements on individual components placed in the fiber optic transmission span (and by requiring system performance to have a tolerance to specified reflectance or ORL values, see Sections 4.2.3 and 4.2.4.3), the effects of fiber optic system reflection noise can be minimized.

R4-17 [111v2] To ensure the capability for upgrading SONET transport systems to high bit rates, all splices, connectors, attenuators, couplers, Wavelength Division Multiplexing (WDM) components, PDC devices and optical fiber amplifiers intended for installation in new facilities shall meet the reflectance requirements specified in the following documents:

GR-765-CORE, *Generic Requirements for Single Fiber Single-Mode Optical Splices and Splicing Systems*

GR-326-CORE, *Generic Requirements for Singlemode Optical Connectors and Jumper Assemblies*

GR-910-CORE, *Generic Requirements for Fiber Optic Attenuators*

GR-1209-CORE, *Generic Requirements for Passive Optical Components*

GR-2854-CORE, *Generic Requirements for Fiber Optic Dispersion Compensators*

GR-1312-CORE, *Generic Requirements for Optical Fiber Amplifiers and Proprietary Dense Wavelength-Division Multiplexed Systems.*

Note that some of the documents listed in **R4-17 [111v2]** may pre-date the development of OC-192 and OC-768 systems. In those cases, the OC-48 reflectance criteria apply for the high bit-rate systems until such time as the documents are updated to be inclusive of those systems.

In addition to individual component reflectance values, another important reflectance-related characteristic of the optical path is its ORL. For SONET systems, that parameter is defined to be the ratio (in dB) of the optical power arriving downstream at the system interface (P_1) to the optical power reflected back upstream to the same interface (P_2). This includes the reflected power contributions from all system components downstream from the interface.

$$\text{ORL} = 10 \log_{10} \left[\frac{P_1}{P_2} \right] \text{ dB} \quad (4-6)$$

As indicated previously, specific requirements relate to discrete reflections along the optical path, to values of the system ORL (which accounts for both discrete reflections and distributed reflections along the fiber), and to reflections from the receiver. Values for each of these are given in Tables 4-6 through 4-25. For embedded facilities that may have ORL values or components with reflectances worse than these specifications, it is an individual network provider's decision to determine the necessity of replacing components or reducing their number to upgrade each facility to the highest anticipated bit rate and desired application. Also, to determine the reflectances between points S and R, and the ORL at those same points, measurements are taken at the end-faces of the connector plugs of the plant fiber shown in Figure 4-1. Since such measurements do not include the performance of the respective connectors in the operational system, it is typically assumed that the connectors C_{TX} and C_{RX} have nominal reflectance values for the specific connector types used.

Finally, to accommodate embedded facilities that may have reflection performance worse than the current recommendations, it is desirable that SONET system transmitters and receivers tolerate higher levels of reflections than would be generated by the discrete reflections discussed above.

O4-18 [114v2] For all applications in Tables 4-6 through 4-25, the optical system should display a power penalty of no more than 1 dB in the presence of a -8.5 dB reflection.

In general, it is expected that most systems that conform to **O4-18 [114v2]** will do so by having transmitters that are able to tolerate a -8.5 dB reflection (i.e., that generate a worst-case or better signal in the presence of such a reflection), and in those cases it would appear to be appropriate for the objective to be moved to [Section 4.2.3](#) and revised to be specific to transmitters. However, it is also possible that a system's transmitter will generate a worse than worst-case signal when exposed to a -8.5 dB reflection, and that its receiver will be designed to tolerate the additional degradations.

4.2.5.2 Attenuation

Fiber attenuation results from a combination of physical phenomena (e.g., Rayleigh scattering and waveguide attenuation) that act to absorb the optical energy propagating in the fiber or to scatter it back upstream or out through the fiber cladding. This document specifies application categories that are based on the maximum attenuation values summarized in [Table 4-4](#). That table also lists, for both 1310-nm and 1550-nm signals, target distances and approximate maximum ranges that correspond to the maximum attenuation values. Note that the maximum range values are estimates calculated using assumed fiber attenuation coefficients (including splices, connectors, etc.) of 0.55 dB/km for 1310-nm signals and 0.275 dB/km for 1550-nm signals, and are provided for information purposes only (i.e., they are not used here for specification purposes). Attenuation coefficients for actual fibers may vary significantly from these assumed values, and thus the actual maximum ranges for systems deployed in the network may also vary significantly from those that appear in [Table 4-4](#). Also note that the target distances are less than the maximum ranges for most OC-1 through OC-48 applications, the OC-192 SR-1 application and the OC-768 SR-1 and SR-2 applications, and that the two values are equal for all VR and UR applications and the remaining OC-192 and OC-768 applications.

In general, the maximum attenuation values listed in [Table 4-4](#) and [Tables 4-6 through 4-25](#) are worst-case values including losses from splices, connectors, attenuators or other optical devices, and any additional cable margin to cover allowances for future modifications (e.g., additional splices) to the cable plant, fiber cable performance variations due to environmental factors and any degradation of connectors or other optical devices between points S and R. In addition, each maximum attenuation value is consistent with the minimum transmitter power (P_{Tmin}), minimum received power (P_{Rmin}) and optical path power penalty (P_O) values listed for that application. Similarly, the minimum attenuation values are generally consistent with the corresponding maximum transmitter power (P_{Tmax}) and maximum received power (P_{Rmax}) values. Finally, to provide flexibility in implementing multi-supplier SONET systems, in a number of cases some overlap in attenuation ranges is provided between adjacent applications (e.g., between SR and IR applications).

Table 4-4 Maximum Attenuation and Range Values

Nominal Wavelength Region	Application Category	Maximum Attenuation (dB)	Target Distance (km)	Approximate Max. Range (km)
1310 nm	OC-1 to OC-48 SR-1	7	2 ^a	12
	OC-192 and OC-768 SR-1	4	2 ^b	7
	OC-1 to OC-48 IR-1	12	15 ^a	21
	OC-192 IR-1	11	20 ^c	20
	OC-1 and OC-3 LR-1	28	40 ^a	50
	OC-12 and OC-48 LR-1	24	40 ^a	42
	OC-192 LR-1	22	40 ^c	40
	OC-12 VR-1	33	60 ^d	60
1550 nm	OC-192 SR-2	7	25 ^c	25
	OC-768 SR-2	4	2 ^b	14
	OC-1 to OC-48 IR-2	12	15 ^a	42
	OC-192 and OC-768 IR-2 & 3	11	40 ^c	40
	OC-1 and OC-3 LR-2 & 3	28	60 or 80 ^a	100
	OC-12 and OC-48 LR-2 & 3	24	60 or 80 ^a	85
	OC-192 and OC-768 LR-2&3	22	80 ^{c,d}	80
	OC-12 to 192 VR-2 & 3	33	120 ^d	120
	OC-12 and OC-48 UR-2 & 3	44	160 ^d	160

Notes:

- a. See ANSI T1.105.06 and ITU-T Rec. G.957.
- b. See ITU-T Rec. G.693.
- c. See ITU-T Rec. G.959.1.
- d. See ITU-T Rec. G.691.

4.2.5.3 Chromatic Dispersion

In general, the chromatic dispersion criteria in this document are specified in terms of the following parameters.

D_{SRmin} , D_{SRmax} – The minimum and maximum values of chromatic dispersion (in ps/nm) between points S and R that can be accommodated by a transmitter and receiver.

Note however, that in most applications (e.g., applications that are not specified to utilize any type of dispersion accommodation) a system is expected to operate properly in the absence of any chromatic dispersion, and therefore D_{SRmin} is equal

to zero and is either not listed or shown to be NA in the corresponding table. In addition, in all of the applications where D_{SRmin} is expected to have a non-zero value, the particular values are currently listed as being FFS.

In the case of D_{SRmax} , two different methods were used to obtain the values that appear in this document. For a number of OC-1 through OC-48 applications these values were calculated using the method described in Section 4.2.5.3.1, while for OC-12 VR and UR, OC-48 LR-2, VR and UR, and all OC-192 and OC-768 applications they were calculated as described in Section 4.2.5.3.2. In addition, the dispersion accommodation methods specified for use in certain OC-192 and OC-768 applications are discussed in Section 4.2.5.3.3.

4.2.5.3.1 Chromatic Dispersion for Most Applications at OC-1 through OC-48

Chromatic dispersion specifications for a number of OC-1, OC-3, OC-12 and OC-48 applications were derived through consideration of the allowable pulse broadening as a fraction of the timeslot width, ε ,

$$\varepsilon = 10^{-6} \times B \times D_{SRmax} \times \delta\lambda \quad (4-7)$$

where B is the system bit rate (in Mb/s) and $\delta\lambda$ is the spectral width (in nm). For LEDs and MLM lasers, $\delta\lambda = \Delta\lambda_{rms}$, while for SLM lasers (assuming Gaussian line shapes) $\delta\lambda = \Delta\lambda_{20} / 6.07$.

For a given dispersion power penalty, ε has a maximum value. Specifying a value for ε allows calculation of a D_{SRmax} value that is consistent with the specified maximum transmitter spectral width. For Intersymbol Interference (ISI) induced power penalties associated with LEDs and SLM lasers, the value $\varepsilon = 0.306$ was assumed. To account for mode-partition noise in MLM lasers, the value $\varepsilon = 0.115$ was assumed. These values are consistent with a total 1-dB dispersion power penalty predicted by empirical and analytic models.

Note that in almost every application where an SLM laser is expected to be deployed, the value of D_{SRmax} calculated as described above (and using an assumed maximum value of $\Delta\lambda_{20} = 1$ nm) is either significantly greater than or significantly less than the maximum amount of dispersion that could occur in that application (based on the target distance or attenuation-based maximum range and the worst-case fiber chromatic dispersion coefficient for the specified combination of wavelength range and type of fiber). In the former cases, as well as in a number of similar cases for applications where MLM lasers are expected to be deployed, D_{SRmax} is indicated to be NA and the system can be considered to be limited by attenuation rather than by dispersion. On the other hand, the latter cases correspond to the applications that would be dispersion limited without some modification to reduce that effect (e.g., a reduced range of allowed central wavelengths, or the use of indirect modulation in order to reduce the spectral width of the signal to a value well below 1 nm), and in those cases D_{SRmax} is calculated as described in Section 4.2.5.3.2. Finally, in the few remaining cases where the calculated value of D_{SRmax} is relatively close to the maximum amount of dispersion that could occur, there is concern that the calculated values of D_{SRmax} may not adequately account

for the dispersion effects resulting from laser chirp. The need to modify the epsilon model as well as the values of D_{SRmax} and $\Delta\lambda_{20}$ to accommodate chirp-induced dispersion power penalties in those cases is under study.

4.2.5.3.2 Chromatic Dispersion for High Bit Rate/Long Range Applications

Chromatic dispersion specifications for OC-12 VR and UR, OC-48 LR-2, VR and UR, and all OC-192 and OC-768 applications are based on the following equation:

$$D_{SRmax} = L_T \times D_{max} \quad (4-8)$$

where L_T is the target distance given in Table 4-4 (in km), and D_{max} is the maximum fiber chromatic dispersion coefficient (in ps/nm-km). For the purpose of computing D_{SRmax} , D_{max} is assumed to be 3.3 ps/nm-km for signals in the 1310-nm region carried on C-SMF or signals in the 1550-nm region carried on DSF, and 20 ps/nm-km for signals in the 1550-nm region carried on C-SMF (i.e., 3.3 ps/nm-km for “-1” or “-3” application categories, and 20 ps/nm-km for “-2” application categories). These values are worst-case values used to provide added tolerance to the systems, thus allowing better utilization of installed fiber plant.

4.2.5.3.3 Chromatic Dispersion Accommodation

In general, chromatic dispersion can be a major issue for OC-192 and OC-768 systems. This is particularly true for systems operating in the 1550-nm window over C-SMF. One method of reducing the effects of chromatic dispersion is to utilize narrow linewidth sources (i.e., indirectly modulated lasers). However even for an indirectly modulated, zero-chirp laser transmitter, the maximum transmission distance (with a 1 dB power penalty) is only about 50 to 60 km for an OC-192 signal (and considerably less for an OC-768 signal). In order to extend the range beyond this limit, other methods, termed dispersion accommodation methods, are necessary.

Dispersion accommodation methods can be divided into two major categories. These are the passive methods (e.g., PDC) and the active methods [e.g., PCH, SPM, Dispersion Supported Transmission (DST)]. Three of these methods (i.e., PDC, PCH and SPM) are addressed in the following subsections, while a discussion of the DST method appears in ITU-T Rec. G.691 (and is currently not expected to be added to a future issue of this document). In addition, in some applications the use of a combination of methods (i.e., a “hybrid” method) may be necessary, and in the future adaptive dispersion compensation may be utilized.

4.2.5.3.3.A Passive Dispersion Compensation

A PDC device is a length of fiber [e.g., Dispersion Compensating Fiber (DCF)] or device (e.g., Fiber Bragg Grating) introduced at either the transmitter (i.e., upstream of point S) or receiver (i.e., downstream of point R) to cause dispersion of the opposite polarity from that which occurs in the main optical path. Thus, if a PDC

device is designed to compensate a particular length of C-SMF, then a fiber link composed of that device and that length of C-SMF will appear to have near-zero dispersion. PDC devices are described further in GR-2854-CORE, and can be designed to compensate any particular length of C-SMF.

In general, for proper operation of a system that utilizes PDC, the relationship between the chromatic dispersion of the optical path at any particular time (D_P) and the actual dispersion compensation provided by the PDC device at that time (D_C) needs to satisfy the following inequality:

$$|D_P + D_C| < D_{r\ max} \quad (4-9)$$

where $D_{r\ max}$ is the maximum residual dispersion that the system is able to tolerate. In addition, D_P and D_C can be related to the path dispersion value determined at the time of installation (D_I) using the equations $D_P = D_I + \delta_P$ and $D_C = -(D_I + \delta_C)$.⁶ Substituting for D_P and D_C in Equation 4-9 yields:

$$|\delta_P - \delta_C| < D_{r\ max} \quad (4-10)$$

In general, effects such as measurement inaccuracies, temperature changes, repairs, aging, the granularity of the possible PDC settings, and intentional under (or over) compensation can cause δ_P and δ_C to have non-zero initial values and to vary over time. According to ITU-T Rec. G.959.1, it is possible to assume reasonable values for $D_{r\ max}$ and δ_C , and to use those values to calculate a maximum allowable magnitude for δ_P . In that case δ_P becomes an additional dispersion-related parameter to be specified (in addition to the two parameters defined in Section 4.2.5.3) and is defined as follows:

δ_P – The maximum chromatic dispersion deviation. This parameter defines the maximum allowable difference between the actual value of the chromatic dispersion of the optical path at any particular time and the path dispersion value determined at the time of installation.

Also according to Rec. G.959.1, an appropriate value for $D_{r\ max}$ in OC-768 applications (i.e., in the only applications in which δ_P is currently expected to be specified) is about 30 ps/nm. On the other hand, at the time that this document was being prepared, Rec. G.959.1 did not indicate an appropriate assumed value for δ_C or specify δ_P for any of the types of applications that have been included here. Therefore, δ_P is listed as being FFS in all of the tables in which it appears (i.e., in Tables 4-14 and 4-20).

As alluded to above, in designing a system that requires dispersion compensation, it is not necessary to choose a PDC device that will result in a total dispersion value of approximately zero. Instead, a system can be designed to be under (or over) compensated. In the under-compensation case, if an uncompensated system is

6. Note that in determining the equations to be used in relating D_P and D_C to D_I in this document, it was assumed that the magnitudes of δ_P and δ_C should generally be small, and that a positive value of δ_P or δ_C should mean that the magnitude of the corresponding dispersion parameter (D_P or D_C) had increased. This second assumption does not appear in Rec. G.959.1 (either explicitly or implicitly), and therefore Equation 4-10 appears as “ $|\delta_P| + |\delta_C| < D_{r\ max}$ ” in that document.

dispersion-limited at a length L_d of C-SMF, the addition of a PDC device designed to compensate a length L_c of C-SMF could extend the dispersion-limited distance to $(L_c + L_d)$.

The use of PDC has the advantage that it is passive, has low polarization sensitivity and a large optical passband, and does not require any special instrumentation to install. A few of the drawbacks are the need to store the extra fiber or other device, the cost, and the extra attenuation associated with the PDC device (which generally must be compensated for by providing additional gain in the BA/OAT or PA/OAR).

4.2.5.3.3.B *Pre-Chirp*

This method relies on optical frequency shifts generated at the transmitter during the signal's rising and/or falling edges to provide pulse compression. Typically, in the types of systems considered in this document, a negative chirp is introduced. PCH may provide all the dispersion accommodation needed in a particular application, or it may be used in conjunction with PDC and/or SPM.

4.2.5.3.3.C *Self Phase Modulation*

In a high power system, pulse compression caused by SPM can provide some or all of the necessary dispersion accommodation. The basis for this pulse compression is a nonlinear effect that causes the fiber to have a different index of refraction for a logical '1' than for a logical '0'. The changes in the refractive index that occur during the signal's rising and falling edges cause frequency shifts in the signal, with a downward frequency shift (i.e., a red shift) at the rising edge and an upward shift (i.e., a blue shift) at the falling edge. For a signal in the 1550-nm region traversing C-SMF, lower frequency components travel slower than higher frequency components. Therefore, the frequency shifts result in pulse compression, which delays the onset of pulse broadening caused by chromatic dispersion.

Since SPM is only significant when the optical power level is high (e.g., greater than about +10 dBm), the amount of dispersion accommodation that can be provided via SPM depends on the optical power profile along the transmission fiber. In addition, the maximum optical power level that can be specified for use in an application that utilizes SPM is limited by other non-linear effects (e.g., SRS, SBS) and safety concerns. Thus, in applications requiring a large amount of dispersion accommodation, another dispersion accommodation method (e.g., PDC) will generally need to be used in conjunction with SPM.

Related to SPM, it should be noted that if the pulse shape and the peak pulse power meet additional criteria, then the pulse becomes a soliton or a soliton-like pulse. Soliton or soliton-like pulses can be engineered to propagate much farther than conventional NRZ pulses, and therefore there has been interest in the possible definition of new application codes for soliton and/or soliton-like RZ signal formats.

4.2.5.4 Polarization Effects

4.2.5.4.1 Polarization Mode Dispersion

PMD is a fiber property that causes different states of polarization to propagate with different group velocities. With conventional (i.e., non-polarization-preserving) fiber, random coupling between the fundamental polarization states of the optical signal occurs along the length of the fiber, leading to a randomly varying DGD between signal intensities that are transmitted in each polarization state. This results in pulse distortion and a power penalty at the optical receiver. When measured over a period of time, the DGD value exhibits statistical variations, with a probability density function that is often assumed to follow the Maxwellian distribution (if the fiber lengths involved are greater than several kilometers). The PMD value of a fiber, which is typically expressed in picoseconds, is the average value of this statistical variable. In a typical communications fiber or optical cable, the PMD effect is wavelength-dependent and fluctuates in time due to such factors as temperature changes and mechanical vibrations.

The PMD coefficients of most current fibers tend to be small enough that PMD effects are only of significant concern in very long distance or high bit-rate applications. However some fibers, especially older fibers, can have very high PMD coefficients, thereby causing problems even in relatively short connections. GR-2947-CORE, *Generic Requirements for Portable Polarization Mode Dispersion (PMD) Test Sets*, contains criteria for portable PMD test sets that can be used to measure the PMD characteristics of new or existing fiber.

Currently, PMD specifications are based on the following parameter.

DGD_{max} – The maximum differential group delay that the system must tolerate with a power penalty of no more than 1 dB.

Note that in all cases where a value of DGD_{max} is listed in Tables 4-6 through 4-25, that value corresponds to 0.3 bit period. According to ITU-T Rec. G.691, for a well-designed receiver a penalty of 1.0 dB for that amount of DGD corresponds to a penalty of 0.1 to 0.2 dB for a DGD of 0.1 bit period.

Due to the statistical nature of DGD, the instantaneous DGD and its power penalty for the optical path may exceed the specified maximum values with a certain probability. The exact statistics of the DGD for the optical path will depend on the statistics of the PMD for fiber cables and optical components in the field environment. The corresponding values of power penalty will also depend on factors such as the optical transmitter and receiver designs.

The probability of the instantaneous DGD exceeding any given value can be obtained from its probability density function. When the DGD statistics are Maxwellian and the system DGD_{max} is known, the equivalent mean DGD can be obtained by dividing by the ratio of maximum to mean DGD that corresponds to an acceptable probability of exceeding that DGD_{max} . Some example ratios derived from Maxwellian statistics are given in Table 4-5.

Table 4-5 Maxwellian DGD Probabilities

“Acceptable” Probability of Exceeding DGD_{max}	DGD_{max}/DGD_{mean}
4.2×10^{-5}	3.0
7.7×10^{-7}	3.5
7.4×10^{-9}	4.0

Statistical specifications of PMD for various types of optical fiber have been developed in IEC SC86A and SC86C. Such specifications enable system designers to calculate the probability that the DGD value will exceed a certain threshold over time.

Note that in many cases chromatic dispersion becomes a concern before DGD. However in some applications (particularly applications where 1550-nm signals are transmitted over DSF), DGD may become the overriding limitation. Also note that methods to compensate for DGD (other than simply regenerating the signal at an intermediate location) are currently being investigated.

4.2.5.4.2 Polarization Dependent Loss

In addition to DGD effects, it should be noted that the random coupling of polarization states in optical fiber can also lead to a time-dependent link loss through interactions with other in-line optical elements. Components such as isolators and optical amplifiers introduce Polarization Dependent Loss (PDL) that converts polarization fluctuations into amplitude modulation of the optical signal. However, this document currently does not contain any criteria related to PDL.

4.2.6 Optical Parameter Tables

Tables 4-6 through 4-25 contain values for the various optical parameters discussed in Sections 4.2.3 through 4.2.5 and applications defined in Section 4.1. In the tables, parameter values that require further investigations are indicated by “FFS”. In addition, those parameters that do not apply to a particular application either are not listed or are indicated by “NA”.

Table 4-6 SR-1 OC-1 Through OC-48 Optical Parameters

Parameter ^a	OC-1	OC-3	OC-12	OC-48	Units
Equivalent Standards Code	T1.105.06, OC-1 SR-1	Rec. G.957, I-1	Rec. G.957, I-4	Rec. G.957, I-16	
Transmitter	MLM or LED	MLM/LED	MLM/LED	MLM	
$\lambda_{Tmin} - \lambda_{Tmax}$	1260 – 1360	1260 – 1360	1261 – 1360	1266 – 1360	nm
$\Delta\lambda_{rms}$	80	40/80	14.5/35	4	nm
$\Delta\lambda_{20}$	NA	NA	NA	NA	nm
SMSR _{min}	NA	NA	NA	NA	dB
P _{Tmax}	- 14	- 8	- 8	- 3	dBm
P _{Tmin}	- 23	- 15	- 15	- 10	dBm
r _{emin}	8.2	8.2	8.2	8.2	dB
Optical Path					
System ORL _{min} ^b	NA	NA	NA	24	dB
DSR _{max}	NA	18/25	13/14	12	ps/nm
Attenuation	0 – 7 ^c	0 – 7	0 – 7	0 – 7	dB
Max. Reflectance between S and R	NA	NA	NA	- 27	dB
Receiver					
P _{Rmax}	- 14	- 8	- 8	- 3	dBm
P _{Rmin}	- 31	- 23	- 23	- 18	dBm
P _O	1	1	1	1	dB
Max. Receiver Reflectance	NA	NA	NA	- 27 ^d	dB

Notes:

- a. All parameter values in this table are specified relative to an optical system design objective of a BER not worse than 10^{-10} (without FEC).
- b. Also see **O4-18 [114v2]** on page 4–18.
- c. The version of ANSI T1.105.06 that was available at the time that this document was being prepared (i.e., T1.105.06-2002) specified a wider attenuation range of 0 to 12 dB for this application category. However, a value of 12 dB for the maximum attenuation is inconsistent with the values of P_{Tmin} , P_{Rmin} and P_O provided here and in the standard.
- d. See the discussion in [Section 4.2.4.2](#).

Table 4-7 SR OC-192 Optical Parameters

Parameter ^a	SR-1	SR-2	Units
Equivalent Standards Code	Rec. G.693, VSR2000-2R1	Rec. G.959.1, P1I1-2D2	
Transmitter	SLM	IM	
$\lambda_{Tmin} - \lambda_{Tmax}$	1290 – 1330	1500 – 1580	nm
$\Delta\lambda_{20}$	1.0	FFS ^c	nm
SMSR _{min}	30	30	dB
α	NA	FFS ^c	-
SPD _{max}	NA	FFS	mW/MHz
P _{Tmax}	-1	-1	dBm
P _{Tmin}	-6	-5	dBm
r _{emin}	6	8.2	dB
Optical Path			
System ORL _{min} ^b	14	24	dB
D _{SRmax}	6.6	500 ^d	ps/nm
Attenuation	0 – 4	0 – 7	dB
Max. Reflectance between S and R	-27	-27	dB
DGD _{max}	30	30	ps
Receiver			
P _{Rmax}	-1	-1	dBm
P _{Rmin}	-11	-14	dBm
P _O	1	2	dB
Max. Receiver Reflectance	-14	-27 ^e	dB

Notes:

- a. All parameter values in this table are specified relative to an optical system design objective of a BER not worse than 10^{-12} (without FEC).
- b. Also see **O4-18 [114v2]** on page 4–18.
- c. It is expected that when values are eventually specified for these parameters, it will not be possible to meet the criteria using a directly modulated source.
- d. For a 2-dB dispersion penalty.
- e. See the discussion in [Section 4.2.4.2](#).

Table 4-8 SR OC-768 Optical Parameters

Parameter ^a	SR-1	SR-2	Units
Equivalent Standards Code	Rec. G.693, VSR2000-3R1	Rec. G.693, VSR2000-3R2	
Transmitter	IM	IM	
$\lambda_{Tmin} - \lambda_{Tmax}$	1290 – 1330	1530 – 1565	nm
$\Delta\lambda_{20}$	FFS	FFS ^c	nm
SMSR _{min}	35	35	dB
α	NA	FFS ^c	-
SPD _{max}	FFS	FFS	mW/MHz
P _{Tmax}	+3	+3	dBm
P _{Tmin}	0	0	dBm
r _{emin}	8.2	8.2	dB
Optical Path			
System ORL _{min} ^b	24	24	dB
D _{SRmax}	6.6	40 ^d	ps/nm
Min. – Max. Attenuation	0 – 4	0 – 4	dB
Max. Reflectance between S and R	-27	-27	dB
DGD _{max}	7.5	7.5	ps
Receiver			
P _{Rmax}	+3	+3	dBm
P _{Rmin}	-5	-6	dBm
P _O	1	2	dB
Max. Receiver Reflectance	-27 ^e	-27 ^e	dB

Notes:

- a. All parameter values in this table are specified relative to an optical system design objective of a BER not worse than 10^{-12} (without FEC).
- b. Also see **O4-18 [114v2]** on page 4–18.
- c. It is expected that when values are eventually specified for these parameters, it will not be possible to meet the criteria using a directly modulated source.
- d. For a 2-dB dispersion penalty.
- e. See the discussion in [Section 4.2.4.2](#).

Table 4-9 IR OC-1 Optical Parameters

Parameter ^a	IR-1	IR-2		Units
Equivalent Standards Code	T1.105.06, OC-1 IR-1	T1.105.06, OC-1 IR-2		
Transmitter	MLM	MLM	SLM	
$\lambda_{Tmin} - \lambda_{Tmax}$	1260 – 1360	1430 – 1580	1430 – 1580	nm
$\Delta\lambda_{rms}$	23	4 ^c	NA	nm
$\Delta\lambda_{20}$	NA	NA	1	nm
SMSR _{min}	NA	NA	30	dB
P _{Tmax}	– 8	– 8	– 8	dBm
P _{Tmin}	– 15	– 15	– 15	dBm
r _{emin}	8.2	8.2	8.2	dB
Optical Path				
System ORL _{min} ^b	NA	NA	NA	dB
D _{SRmax}	NA	555 ^c	NA	ps/nm
Attenuation	0 – 12	0 – 12	0 – 12	dB
Max. Reflectance between S and R	NA	NA	NA	dB
Receiver				
P _{Rmax}	– 8	– 8		dBm
P _{Rmin}	– 28	– 28		dBm
P _O	1	1		dB
Max. Receiver Reflectance	NA	NA		dB

Notes:

- a. All parameter values in this table are specified relative to an optical system design objective of a BER not worse than 10^{-10} .
- b. Also see **O4-18 [114v2]** on page 4–18.
- c. The $D_{SRmax} = 317$ ps/nm entry that appeared in the various versions of this document before Issue 4 (as well as in the 2002 version of ANSI T1.105.06) was inconsistent with Equation 4-7 and the specified values of $\Delta\lambda_{rms} = 4$ nm and $\varepsilon = 0.115$ (see Section 4.2.5.3.1). While that inconsistency could also have been eliminated by changing the $\Delta\lambda_{rms}$ entry to 7 nm (possibly increasing the portion of the manufactured sources that could be used in equipment intended for this application), it was felt that the change shown here (which may allow the use of IR-2 equipment over somewhat longer optical paths) would be of more value. This decision is subject to change based on future standards work.

Table 4-10 IR OC-3 Optical Parameters

Parameter ^a	IR-1	IR-2		Units
Equivalent Standards Code	Rec. G.957, S-1.1	Rec. G.957, S-1.2		
Transmitter	MLM	MLM	SLM	
$\lambda_{Tmin} - \lambda_{Tmax}$	1261 – 1360	1430 – 1576	1430 – 1580	nm
$\Delta\lambda_{rms}$	7.7	2.5	NA	nm
$\Delta\lambda_{20}$	NA	NA	1	nm
$SMSR_{min}$	NA	NA	30	dB
P_{Tmax}	– 8	– 8	– 8	dBm
P_{Tmin}	– 15	– 15	– 15	dBm
r_{emin}	8.2	8.2	8.2	dB
Optical Path				
System ORL_{min}^b	NA	NA	NA	dB
DSR_{max}	96	296	NA	ps/nm
Attenuation	0 – 12	0 – 12	0 – 12	dB
Max. Reflectance between S and R	NA	NA	NA	dB
Receiver				
P_{Rmax}	– 8	– 8		dBm
P_{Rmin}	– 28	– 28		dBm
P_O	1	1		dB
Max. Receiver Reflectance	NA	NA		dB

Notes:

- a. All parameter values in this table are specified relative to an optical system design objective of a BER not worse than 10^{-10} .
- b. Also see **O4-18 [114v2]** on page 4–18.

Table 4-11 IR OC-12 Optical Parameters

Parameter ^a	IR-1	IR-2	Units
Equivalent Standards Code	Rec. G.957, S-4.1	Rec. G.957, S-4.2	
Transmitter	MLM	SLM	
$\lambda_{Tmin} - \lambda_{Tmax}$	1274 – 1356 or 1293 – 1334 ^c	1430 – 1580	nm
$\Delta\lambda_{rms}$	2.5 (or 4.0) ^c	NA	nm
$\Delta\lambda_{20}$	NA	1	nm
$SMSR_{min}$	NA	30	dB
P_{Tmax}	– 8	– 8	dBm
P_{Tmin}	– 15	– 15	dBm
r_{emin}	8.2	8.2	dB
Optical Path			
System ORL _{min} ^b	NA	24	dB
D_{SRmax}	74 (or 46) ^c	NA	ps/nm
Attenuation	0 – 12	0 – 12	dB
Max. Reflectance between S and R	NA	– 27	dB
Receiver			
P_{Rmax}	– 8	– 8	dBm
P_{Rmin}	– 28	– 28	dBm
P_O	1	1	dB
Max. Receiver Reflectance	NA	– 27 ^d	dB

Notes:

- a. All parameter values in this table are specified relative to an optical system design objective of a BER not worse than 10^{-10} .
- b. Also see **O4-18 [114v2]** on page 4–18.
- c. Transmitters having central wavelengths within the narrower of these ranges (i.e., from 1293 to 1334 nm) are allowed a wider maximum spectral width of 4.0 nm. In addition, in that case the applicable value of D_{SRmax} is 46 ps/nm.
- d. See the discussion in [Section 4.2.4.2](#).

Table 4-12 IR OC-48 Optical Parameters

Parameter ^a	IR-1	IR-2	Units
Equivalent Standards Code	Rec. G.957, S-16.1	Rec. G.957, S-16.2	
Transmitter	SLM	SLM	
$\lambda_{Tmin} - \lambda_{Tmax}$	1260 – 1360	1430 – 1580	nm
$\Delta\lambda_{rms}$	NA	NA	nm
$\Delta\lambda_{20}$	1	<1 ^d	nm
$SMSR_{min}$	30	30	dB
P_{Tmax}	0	0	dBm
P_{Tmin}	-5	-5	dBm
r_{emin}	8.2	8.2	dB
Optical Path			
System ORL_{min} ^b	24	24	dB
DSR_{max}	NA	Note d	ps/nm
Attenuation	0 – 12	0 – 12	dB
Max. Reflectance between S and R	-27	-27	dB
Receiver			
P_{Rmax}	0	0	dBm
P_{Rmin}	-18	-18	dBm
P_O	1	1	dB
Max. Receiver Reflectance	-27 ^c	-27 ^c	dB

Notes:

- a. All parameter values in this table are specified relative to an optical system design objective of a BER not worse than 10^{-10} (without FEC).
- b. Also see **O4-18 [114v2]** on page 4-18.
- c. See the discussion in [Section 4.2.4.2](#).
- d. Currently, the methods for estimating chirp power penalties, and thus the appropriate values of $\Delta\lambda_{20}$ and DSR_{max} for this application category, are under study. If agreements are reached (e.g., in ITU-T), it is expected that this document will be revised to reflect those agreements. In addition, it is expected that the value of $\Delta\lambda_{20}$ will be less than 1 nm.

Table 4-13 IR OC-192 Optical Parameters

Parameter ^a	IR-1	IR-2	IR-3	Units
Equivalent Standards Code	Rec. G.959.1, P1S1-2D1	Rec. G.959.1, P1S1-2D2b	Rec. G.959.1, P1S1-2D3b	
Transmitter	SLM	IM	SLM	
$\lambda_{Tmin} - \lambda_{Tmax}$	1290 – 1330	1530 – 1565	1530 – 1565	nm
$\Delta\lambda_{20}$	FFS	FFS ^c	FFS	nm
SMSR _{min}	30	30	30	dB
α	NA	FFS ^c	FFS	-
SPD _{max}	FFS	FFS	FFS	mW/MHz
P _{Tmax}	+5	+2	+2	dBm
P _{Tmin}	+1	-1	-1	dBm
r _{emin}	6	8.2	8.2	dB
Optical Path				
System ORL _{min} ^b	14	24	24	dB
D _{SRmax}	70	800 ^d	130	ps/nm
Min. – Max. Attenuation	6 – 11	3 – 11	3 – 11	dB
Max. Reflectance between S and R	-27	-27	-27	dB
DGD _{max}	30	30	30	ps
Receiver				
P _{Rmax}	-1	-1	-1	dBm
P _{Rmin}	-11	-14	-13	dBm
P _O	1	2	1	dB
Max. Receiver Reflectance	-14	-27 ^e	-27 ^e	dB

Notes:

- All parameter values in this table are specified relative to an optical system design objective of a BER not worse than 10^{-12} (without FEC).
- Also see **O4-18 [114v2]** on page 4–18.
- It is expected that when values are eventually specified for these parameters, it will not be possible to meet the criteria using a directly modulated source.
- For a 2-dB dispersion penalty.
- See the discussion in [Section 4.2.4.2](#).

Table 4-14 IR OC-768 Optical Parameters

Parameter ^a	IR-2	IR-3	Units
Equivalent Standards Code	Rec. G.959.1, P1S1-3C2	Rec. G.959.1, P1S1-3C3	
Transmitter	IM	IM	
v_{nom}	192.1	192.1	THz
Δv_{max}	±40	±40	GHz
$SMSR_{min}$	35	35	dB
SPD_{max}	FFS	FFS	mW/MHz
P_{Tmax}	+3	+3	dBm
P_{Tmin}	-3	-3	dBm
r_{emin}	8.2	8.2	dB
Optical Path			
System ORL_{min}^b	24	24	dB
D_{SRmax}	800 ^c	140 ^{c,d}	ps/nm
δ_p	FFS ^e	FFS ^e	ps/nm
Min. – Max. Attenuation	0 – 11	0 – 11	dB
Max. Reflectance between S and R	-27	-27	dB
DGD_{max}	7.5	7.5	ps
Receiver			
P_{Rmax}	+3	+3	dBm
P_{Rmin}	-17	-17	dBm
P_O	3	3	dB
Max. Receiver Reflectance	-27 ^f	-27 ^f	dB

Notes:

- a. All parameter values in this table are specified relative to an optical system design objective of a BER not worse than 10^{-12} (without FEC).
- b. Also see **04-18 [114v2]** on page 4–18.
- c. For a 3-dB dispersion penalty.
- d. For IR-3 applications, the value of D_{SRmax} calculated as described in [Section 4.2.5.3.2](#) is 132 ps/nm. Consistent with ITU-T Rec. G.959.1, that value is rounded down to 130 ps/nm in [Table 4-13](#) (for OC-192) and is rounded up to 140 ps/nm here (for OC-768).
- e. No value has been specified for this parameter, and therefore it needs to be agreed upon by joint engineering between the service provider and the NE supplier.
- f. See the discussion in [Section 4.2.4.2](#).

Table 4-15 LR OC-1 Optical Parameters

Parameter ^a	LR-1		LR-2	LR-3		Units
Equivalent Standards Code	T1.105.06, LR-1		T1.105.06, LR-2	T1.105.06, LR-3		
Transmitter	MLM	SLM	SLM	MLM	SLM	
$\lambda_{Tmin} - \lambda_{Tmax}$	1280-1335	1280-1335	1480-1580	1480-1580	1480-1580	nm
$\Delta\lambda_{rms}$	13	NA	NA	5	NA	nm
$\Delta\lambda_{20}$	NA	1	1	NA	1	nm
SMSR _{min}	NA	30	30	NA	30	dB
P _{Tmax}	0	0	0	0	0	dBm
P _{Tmin}	-5	-5	-5	-5	-5	dBm
r _{emin}	10	10	10	10	10	dB
Optical Path						
System ORL _{min} ^b	NA	NA	20	NA	NA	dB
D _{SRmax}	171	NA	NA	444	NA	ps/nm
Min. – Max. Attenuation	10 – 28	10 – 28	10 – 28	10 – 28	10 – 28	dB
Max Reflectance between S and R	NA	NA	-25	NA	NA	dB
Receiver						
P _{Rmax}	-10		-10	-10		dBm
P _{Rmin}	-34		-34	-34		dBm
P _O	1		1	1		dB
Max. Receiver Reflectance	NA		-25 ^c	NA		dB

Notes:

- a. All parameter values in this table are specified relative to an optical system design objective of a BER not worse than 10^{-10} .
- b. Also see **O4-18 [114v2]** on page 4–18.
- c. See the discussion in [Section 4.2.4.2](#).

Table 4-16 LR OC-3 Optical Parameters

Parameter ^a	LR-1		LR-2	LR-3		Units
Equivalent Standards Code	Rec. G.957, L-1.1		Rec. G.957, L-1.2	Rec. G.957, L-1.3		
Transmitter	MLM	SLM	SLM	MLM	SLM	
$\lambda_{Tmin} - \lambda_{Tmax}$	1263-1360	1263-1360	1480-1580	1523-1577 or 1534-1566 ^d	1480-1580	nm
$\Delta\lambda_{rms}$	3	NA	NA	2.5 (or 3) ^d	NA	nm
$\Delta\lambda_{20}$	NA	1	1	NA	1	nm
$SMSR_{min}$	NA	30	30	NA	30	dB
P_{Tmax}	0	0	0	0	0	dBm
P_{Tmin}	-5	-5	-5	-5	-5	dBm
r_{emin}	10	10	10	10	10	dB
Optical Path						
System ORL_{min} ^b	NA	NA	20	NA	NA	dB
DSR_{max}	246	NA	NA	296 (or 246) ^d	NA	ps/nm
Min. – Max. Attenuation	10 – 28	10 – 28	10 – 28	10 – 28	10 – 28	dB
Max Reflectance between S and R	NA	NA	-25	NA	NA	dB
Receiver						
P_{Rmax}	-10	-10	-10	-10	-10	dBm
P_{Rmin}	-34	-34	-34	-34	-34	dBm
P_O	1	1	1	1	1	dB
Max. Receiver Reflectance	NA	NA	-25 ^c	NA	NA	dB

Notes:

- All parameter values in this table are specified relative to an optical system design objective of a BER not worse than 10^{-10} .
- Also see **O4-18 [114v2]** on page 4-18.
- See the discussion in [Section 4.2.4.2](#).
- Transmitters having central wavelengths within the narrower of these ranges (i.e., from 1534 to 1566 nm) are allowed a wider maximum spectral width of 3.0 nm. In addition, in that case the applicable value of DSR_{max} is 246 ps/nm.

Table 4-17 LR OC-12 Optical Parameters

Parameter ^a	LR-1		LR-2	LR-3	Units
Equivalent Standards Code	Rec. G.957, L-4.1		Rec. G.957, L-4.2	Rec. G.957, L-4.3	
Transmitter	MLM	SLM	SLM	SLM	
$\lambda_{Tmin} - \lambda_{Tmax}$	1296–1330 or 1300–1325 ^c	1280–1335	1480–1580	1480–1580	nm
$\Delta\lambda_{rms}$	1.7 (or 2.0) ^c	NA	NA	NA	nm
$\Delta\lambda_{20}$	NA	1	<1 ^d	1	nm
SMSR _{min}	NA	30	30	30	dB
P _{Tmax}	+2	+2	+2	+2	dBm
P _{Tmin}	– 3	– 3	– 3	– 3	dBm
r _{emin}	10	10	10	10	dB
Optical Path					
System ORL _{min} ^b	20	20	24	20	dB
D _{SRmax}	109 (or 92) ^c	NA	Note d	NA	ps/nm
Min. – Max. Attenuation	10 – 24	10 – 24	10 – 24	10 – 24	dB
Max. Reflectance between S and R	– 25	– 25	– 27	– 25	dB
Receiver					
P _{Rmax}	– 8		– 8	– 8	dBm
P _{Rmin}	– 28		– 28	– 28	dBm
P _O	1		1	1	dB
Max. Receiver Reflectance	– 14		– 27 ^e	– 14	dB

Notes:

- a. All parameter values in this table are specified relative to an optical system design objective of a BER not worse than 10^{-10} .
- b. Also see **O4-18 [114v2]** on page 4–18.
- c. Transmitters having central wavelengths within the narrower of these ranges (i.e., from 1300 to 1325 nm) are allowed a wider maximum spectral width of 2.0 nm. In addition, in that case the applicable value of D_{SRmax} is 92 ps/nm.
- d. Currently, the methods for estimating chirp power penalties, and thus the appropriate values of $\Delta\lambda_{20}$ and D_{SRmax} for this application category, are under study. If agreements are reached (e.g., in ITU-T), it is expected that this document will be revised to reflect those agreements. In addition, it is expected that the value of $\Delta\lambda_{20}$ will be less than 1 nm.
- e. See the discussion in [Section 4.2.4.2](#).

Table 4-18 LR OC-48 Optical Parameters

Parameter ^a	LR-1	LR-2	LR-3	Units
Equivalent Standards Code	Rec. G.957, L-16.1	Rec. G.957, L-16.2	Rec. G.957, L-16.3	
Transmitter	SLM	SLM	SLM	
$\lambda_{Tmin} - \lambda_{Tmax}$	1280 – 1335	1500 – 1580	1500 – 1580	nm
$\Delta\lambda_{rms}$	NA	NA	NA	nm
$\Delta\lambda_{20}$	1	<1 ^e	<1 ^e	nm
SMSR _{min}	30	30	30	dB
P _{Tmax}	+3	+3	+3	dBm
P _{Tmin}	– 2	– 2	– 2	dBm
r _{emin}	8.2	8.2	8.2	dB
Optical Path				
System ORL _{min} ^b	24	24	24	dB
DSR _{max}	NA	1200 – 1600 ^{f, g}	Note e	ps/nm
Min. – Max. Attenuation	10 ^c – 24	10 ^c – 24	10 ^c – 24	dB
Max. Reflectance between S and R	– 27	– 27	– 27	dB
Receiver				
P _{Rmax}	– 9	– 9	– 9	dBm
P _{Rmin}	– 27	– 28	– 27	dBm
P _O	1	2	1	dB
Max. Receiver Reflectance	– 27 ^d	– 27 ^d	– 27 ^d	dB

Notes:

- a. All parameter values in this table are specified relative to an optical system design objective of a BER not worse than 10⁻¹⁰ (without FEC).
- b. Also see **O4-18 [114v2]** on page 4–18.
- c. The lower attenuation limit of 10 dB may require use of optical attenuators, lower maximum output power, increased minimum overload, or a combination thereof.
- d. See the discussion in [Section 4.2.4.2](#).
- e. Currently, the methods for estimating chirp power penalties, and thus the appropriate value of $\Delta\lambda_{20}$ for this application category, are under study. If agreements are reached (e.g., in ITU-T), it is expected that this document will be revised to reflect those agreements. In addition, it is expected that the value of $\Delta\lambda_{20}$ will be less than 1 nm.
- f. For a 2-dB dispersion penalty.
- g. These values correspond to the approximate worst-case dispersion for 80 km of C-SMF and a wavelength range of 1500 to 1580 nm. For the purpose of chromatic dispersion power penalty tests, the receiver would need to be able to accommodate the larger of these values (i.e., 1600 ps/nm).

Table 4-19 LR OC-192 Optical Parameters

Parameter ^a	LR-1	LR-2	LR-2a and LR-2c	LR-2b	LR-3	Units
Equivalent Standards Code	Rec. G.959.1, P1L1-2D1	Rec. G.959.1, P1L1-2D2	Rec. G.691, L-64.2a and L-64.2c	Rec. G.691, L-64.2b	Rec. G.691, L-64.3	
Transmitter	SLM	IM	IM	IM	IM	
$\lambda_{Tmin} - \lambda_{Tmax}$	1290–1320	1530–1565	1530–1565	1530–1565	1530–1565	nm
$\Delta\lambda_{20}$	FFS	FFS ^e	FFS ^e	FFS ^e	FFS	nm
SMSR _{min}	30	30	FFS	FFS	FFS	dB
α	NA	FFS ^e	FFS ^e	FFS ^e	FFS	-
SPD _{max}	FFS	FFS	FFS	FFS	FFS	mW/MHz
P _{Tmax}	+7	+4	+2	+13	+13	dBm
P _{Tmin}	+3	0	-2	+10	+10	dBm
r _{emin}	6	9	10	8.2	8.2	dB
Optical Path						
System ORL _{min} ^b	24	24	24	24	24	dB
D _{SRmin} –D _{SRmax}	NA – 130	FFS–1600 ^f	FFS–1600 ^f	FFS–1600 ^f	NA – 260	ps/nm
Min. – Max. Attenuation	16 ^c – 22	11 – 22	11 – 22	16 ^c – 22	16 ^c – 22	dB
Max. Reflectance between S and R	-27	-27	-27	-27	-27	dB
DGD _{max}	30	30	30	30	30	ps
Receiver						
P _{Rmax}	-9	-7	-9	-3	-3	dBm
P _{Rmin}	-20	-24	-26	-14	-13	dBm
P _O	1	2	2	2	1	dB
Max. Receiver Reflectance	-27 ^d	-27 ^d	-27 ^d	-27 ^d	-27 ^d	dB

Notes:

- All parameter values in this table are specified relative to an optical system design objective of a BER not worse than 10^{-12} (without FEC).
- Also see **O4-18 [114v2]** on page 4–18.
- Systems with attenuation levels between 11 and 16 dB will need to use this category with an optical attenuator.
- See the discussion in [Section 4.2.4.2](#).
- It is expected that when values are eventually specified for these parameters, it will not be possible to meet the criteria using a directly modulated source.
- For a 2-dB dispersion penalty.

Table 4-20 LR OC-768 Optical Parameters

Parameter ^a	LR-2	LR-3	Units
Equivalent Standards Code	Rec. G.959.1, P1L1-3A2	Rec. G.959.1, P1L1-3A3	
Transmitter	IM	IM	
v_{nom}	192.1	192.1	THz
Δv_{max}	±40	±40	GHz
$SMSR_{min}$	35	35	dB
SPD_{max}	FFS	FFS	mW/MHz
P_{Tmax}	+8	+8	dBm
P_{Tmin}	+5	+5	dBm
r_{emin}	10	10	dB
Optical Path			
System ORL_{min}^b	24	24	dB
D_{SRmax}	1600 ^c	280 ^{c,d}	ps/nm
δ_p	FFS ^e	FFS ^e	ps/nm
Min. – Max. Attenuation	11 – 22	11 – 22	dB
Max. Reflectance between S and R	-27	-27	dB
DGD_{max}	7.5	7.5	ps
Receiver			
P_{Rmax}	-3	-3	dBm
P_{Rmin}	-20	-20	dBm
P_O	3	3	dB
Max. Receiver Reflectance	-27 ^f	-27 ^f	dB

Notes:

- a. All parameter values in this table are specified relative to an optical system design objective of a BER not worse than 10^{-12} (without FEC).
- b. Also see **04-18 [114v2]** on page 4-18.
- c. For a 3-dB dispersion penalty.
- d. For OC-192 and OC-768 LR-3 applications, the value of D_{SRmax} calculated as described in [Section 4.2.5.3.2](#) is 264 ps/nm. Consistent with ITU-T Recommendations G.691 and G.959.1, that value is rounded down to 260 ps/nm in [Table 4-19](#) (for OC-192) and is rounded up to 280 ps/nm here (for OC-768).
- e. No value has been specified for this parameter, and therefore it needs to be agreed upon by joint engineering between the service provider and the NE supplier.
- f. See the discussion in [Section 4.2.4.2](#).

Table 4-21 VR OC-12 Optical Parameters

Parameter ^a	VR-1	VR-2	VR-3	Units
Equivalent Standards Code	Rec. G.691, V-4.1	Rec. G.691, V-4.2	Rec. G.691, V-4.3	
Transmitter	SLM	SLM	SLM	
$\lambda_{Tmin} - \lambda_{Tmax}$	1290 – 1330	1530 – 1565	1530 – 1565	nm
$\Delta\lambda_{20}$	FFS	FFS	FFS	nm
SMSR _{min}	FFS	FFS	FFS	dB
α	NA	NA	NA	-
SPD _{max}	FFS	FFS	FFS	mW/MHz
P _{Tmax}	+4	+4	+4	dBm
P _{Tmin}	0	0	0	dBm
r _{emin}	10	10	10	dB
Optical Path				
System ORL _{min} ^b	24	24	24	dB
D _{SRmax}	200	2400	400	ps/nm
Min. – Max. Attenuation	22 – 33	22 – 33	22 – 33	dB
Max. Reflectance between S and R	-27	-27	-27	dB
DGD _{max}	480	480	480	ps
Receiver				
P _{Rmax}	-18	-18	-18	dBm
P _{Rmin}	-34	-34	-34	dBm
P _O	1	1	1	dB
Max. Receiver Reflectance	-27 ^c	-27 ^c	-27 ^c	dB

Notes:

- All parameter values in this table are specified relative to an optical system design objective of a BER not worse than 10^{-12} .
- Also see **O4-18 [114v2]** on page 4-18.
- See the discussion in [Section 4.2.4.2](#).

Table 4-22 VR OC-48 Optical Parameters

Parameter ^a	VR-2	VR-3	Units
Equivalent Standards Code	Rec. G.691, V-16.2	Rec. G.691, V-16.3	
Transmitter	IM	SLM	
$\lambda_{Tmin} - \lambda_{Tmax}$	1530 – 1565	1530 – 1565	nm
$\Delta\lambda_{20}$	FFS	FFS	nm
SMSR _{min}	FFS	FFS	dB
α	FFS	FFS	-
SPD _{max}	FFS	FFS	mW/MHz
P _{Tmax}	+13	+13	dBm
P _{Tmin}	+10	+10	dBm
r _{emin}	8.2	8.2	dB
Optical Path			
System ORL _{min} ^b	24	24	dB
D _{SRmax}	2400 ^c	400	ps/nm
Min. – Max. Attenuation	22 – 33	22 – 33	dB
Max. Reflectance between S and R	-27	-27	dB
DGD _{max}	120	120	ps
Receiver			
P _{Rmax}	-9	-9	dBm
P _{Rmin}	-25	-24	dBm
P _O	2	1	dB
Max. Receiver Reflectance	-27 ^d	-27 ^d	dB

Notes:

- a. All parameter values in this table are specified relative to an optical system design objective of a BER not worse than 10^{-12} (without FEC).
- b. Also see **O4-18 [114v2]** on page 4-18.
- c. For a 2-dB dispersion penalty.
- d. See the discussion in [Section 4.2.4.2](#).

Table 4-23 VR OC-192 Optical Parameters

Parameter ^a	VR-2a	VR-2b	VR-3	Units
Equivalent Standards Code	Rec. G.691, V-64.2a	Rec. G.691, V-64.2b	Rec. G.691, V-64.3	
Transmitter	IM	IM	IM	
$\lambda_{Tmin} - \lambda_{Tmax}$	1530 – 1565	1530 – 1565	1530 – 1565	nm
$\Delta\lambda_{20}$	FFS ^c	FFS ^c	FFS ^c	nm
SMSR _{min}	FFS	FFS	FFS	dB
α	FFS ^c	FFS ^c	FFS ^c	-
SPD _{max}	FFS	FFS	FFS	mW/MHz
P _{Tmax}	+13	+15	+13	dBm
P _{Tmin}	+10	+12	+10	dBm
r _{emin}	10	8.2	8.2	dB
Optical Path				
System ORL _{min} ^b	24	24	24	dB
D _{SRmin} – D _{SRmax}	FFS – 2400 ^d	FFS – 2400 ^d	NA – 400	ps/nm
Min. – Max. Attenuation	22 – 33	22 – 33	22 – 33	dB
Max. Reflectance between S and R	-27	-27	-27	dB
DGD _{max}	30	30	30	ps
Receiver				
P _{Rmax}	-9	-7	-9	dBm
P _{Rmin}	-25	-23	-24	dBm
P _O	2	2	1	dB
Max. Receiver Reflectance	-27 ^e	-27 ^e	-27 ^e	dB

Notes:

- a. All parameter values in this table are specified relative to an optical system design objective of a BER not worse than 10^{-12} (without FEC).
- b. Also see **O4-18 [114v2]** on page 4–18.
- c. It is expected that when values are eventually specified for these parameters, it will not be possible to meet the criteria using a directly modulated source.
- d. For a 2-dB dispersion penalty.
- e. See the discussion in [Section 4.2.4.2](#).

Table 4-24 UR OC-12 Optical Parameters

Parameter ^a	UR-2	UR-3	Units
Equivalent Standards Code	Rec. G.691, U-4.2	Rec. G.691, U-4.3	
Transmitter	SLM	SLM	
$\lambda_{Tmin} - \lambda_{Tmax}$	1530 – 1565	1530 – 1565	nm
$\Delta\lambda_{20}$	FFS	FFS	nm
$SMSR_{min}$	FFS	FFS	dB
α	NA	NA	-
SPD_{max}	FFS	FFS	mW/MHz
P_{Tmax}	+15	+15	dBm
P_{Tmin}	+12	+12	dBm
r_{emin}	10	10	dB
Optical Path			
System ORL _{min} ^b	24	24	dB
D_{SRmax}	3200 ^c	530	ps/nm
Min. – Max. Attenuation	33 – 44	33 – 44	dB
Max. Reflectance between S and R	-27	-27	dB
DGD_{max}	480	480	ps
Receiver			
P_{Rmax}	-18	-18	dBm
P_{Rmin}	-34	-33	dBm
P_O	2	1	dB
Max. Receiver Reflectance	-27 ^d	-27 ^d	dB

Notes:

- a. All parameter values in this table are specified relative to an optical system design objective of a BER not worse than 10^{-12} .
- b. Also see **O4-18 [114v2]** on page 4-18.
- c. For a 2-dB dispersion penalty.
- d. See the discussion in [Section 4.2.4.2](#).

Table 4-25 UR OC-48 Optical Parameters

Parameter ^a	UR-2	UR-3	Units
Equivalent Standards Code	Rec. G.691, U-16.2	Rec. G.691, U-16.3	
Transmitter	IM	IM	
$\lambda_{Tmin} - \lambda_{Tmax}$	1530 – 1565	1530 – 1565	nm
$\Delta\lambda_{20}$	FFS ^c	FFS ^c	nm
SMSR _{min}	FFS	FFS	dB
α	FFS ^c	FFS ^c	-
SPD _{max}	FFS	FFS	mW/MHz
P _{Tmax}	+15	+15	dBm
P _{Tmin}	+12	+12	dBm
r _{emin}	10	10	dB
Optical Path			
System ORL _{min} ^b	24	24	dB
D _{SRmax}	3200 ^d	530	ps/nm
Min. – Max. Attenuation	33 – 44	33 – 44	dB
Max. Reflectance between S and R	-27	-27	dB
DGD _{max}	120	120	ps
Receiver			
P _{Rmax}	-18	-18	dBm
P _{Rmin}	-34	-33	dBm
P _O	2	1	dB
Max. Receiver Reflectance	-27 ^e	-27 ^e	dB

Notes:

- All parameter values in this table are specified relative to an optical system design objective of a BER not worse than 10^{-12} (without FEC).
- Also see **O4-18 [114v2]** on page 4–18.
- It is expected that when values are eventually specified for these parameters, it will not be possible to meet the criteria using a directly modulated source.
- For a 2-dB dispersion penalty.
- See the discussion in [Section 4.2.4.2](#).

4.3 Engineering of a Single-Mode Fiber Optic Transmission System

This section describes a methodology that may be used for engineering a single-mode fiber optic transmission system. The design approach this section describes is derived from TIA-559, *Single-Mode Fiber Optic System Transmission Design*, which was originally intended to apply to single-supplier systems operating up to about 0.5 Gb/s in the 1310-nm region over dispersion unshifted single-mode fiber.

Parameters that allow the design of regenerator sections are discussed below. Some of the parameters in this section are defined relative to measurements specified in EIA/TIA documents (denoted by FOTP or OFSTP, some of which may be in draft form). [Section 4.3.1](#) describes the transmission design information for terminal and regenerator station equipment, while [Section 4.3.2](#) describes that for outside-plant cable. Worksheets 1 through 4 in Appendix B are examples of forms that could be used in gathering the transmission design information. [Section 4.3.3](#) discusses how this information is used in the design and analysis of regenerator sections, which users or system integrators can perform.

In this section, the transmission design parameters are specified by worst-case values.

R4-19 [123] Suppliers shall provide worst-case values of transmission design parameters⁷ requested as part of system documentation.

CR4-20 [124] A network provider may require that suppliers guarantee the worst-case values over the lifetime of their system components.

4.3.1 Terminal Equipment Transmission Design Information

This section describes the transmission information for terminal and regenerator station fiber optic equipment that a supplier and system integrator must provide. Worksheets 1 through 3 contain a summary of terminal equipment transmission design information.

R4-21 [125] The supplier shall provide general terminal equipment information in Worksheet 1.

R4-22 [126] The supplier and the system integrator shall provide terminal equipment parameters under normal operating and short-term emergency conditions in Worksheets 2 and 3, respectively.

R4-23 [127v2] If the terminal equipment has multiple options, the information of the worksheets shall be provided for each option.

7. Some of the parameters are to be provided by the “system integrator” responsible for the overall system design. This may or may not be the supplier.

4.3.1.1 General System Information

Terminal Equipment Identification – The terminal equipment identification uniquely characterizes the type of product and provides a traceable indicator for determining the product specifications, features, issue or revision, and manufacturer (e.g., CLEI™).

Optical Line Rate – $N \times 51.840$ Mb/s, where $N = 1, 3, 12, 24, 48, 192$ or 768 .

4.3.1.2 Transmitter Information

General – The unit providing the transmitter function is identified by a unique descriptor from which the following information can be determined, using the appropriate documentation:

- Manufacturer
- Terminal equipment association
- System design application (e.g., single-mode long-reach)
- Operating wavelength
- Output power level
- Source type
- Optical device temperature controller
- FDA classification (e.g., Class I or Class II)
- Manufacturer product change designation (e.g., issue or revision).

Optical Source Type – The optical source type is characterized by identifying, as a minimum:

- Type of device (e.g., LED, edge emitting LED, SLM laser, MLM laser)
- Material composition of source (e.g., InGaAs)
- Generic device structure [e.g., Distributed Feedback (DFB)]
- Type of modulation used (e.g., direct, Electro-Absorption, Mach-Zehnder).

Transmitter Connector – The transmitter connector is the optical connector provided at the output of the transmitter that attaches to the transmitter pigtail. The transmitter connector description, as a minimum, includes:

- Connector manufacturer
- Connector type (e.g., SC, FC)
- Connector model number
- Connector classification (multimode, single-mode)
- Mating connector model number.

Generic connector requirements appear in GR-326-CORE.

Transmitter Pigtail – The identification of the transmitter pigtail includes the following information (see TIA-4920000B, *Generic Specification for Optical Waveguide Fibers*):

- General fiber type
- Class of fiber
- Mode field diameter.

Maximum Optical Reflection (OR_{max}) – The total reflected optical power (in dB) that a transmitter can accommodate and maintain its stated performance. [Section 4.2.5.1](#) contains the lower bounds to OR_{max} .

Reflection Power Penalty (RP) – The maximum power penalty associated with a worst-case optical reflection at the transmitter (OR_{max}), at a BER of 10^{-10} or 10^{-12} (as appropriate), when operated under standard operating conditions (RP_1) or extended operating conditions (RP_2).

Also specified are the nominal central wavelength λ_{Tnom} or frequency ν_{nom} , transmitter central wavelength range (λ_{Tmin} , λ_{Tmax}) or maximum central frequency deviation $\Delta\nu_{max}$, transmitter power (P_T) and possibly the source frequency chirp factor (α), as defined in [Section 4.2.3](#).

4.3.1.3 Receiver Information

4.3.1.3.1 General Information

General – The unit providing the receiver function is identified by a unique descriptor from which the following information can be determined, using the appropriate documentation:

- Manufacturer
- Terminal equipment association
- System design application (e.g., single-mode long-reach)
- Receiver performance specifications
- Detector type
- Optical device temperature controller
- Manufacturer product change designation (e.g., issue or revision).

Optical Detector Type – The optical detector type is characterized by identifying as a minimum:

- Device type (e.g., PIN, APD)
- Material composition of detector (e.g., Ge, Si).

Receiver Pigtail – The identification of the receiver pigtail includes the following information (see TIA-4920000B):

- General fiber type
- Class of fiber
- Mode field diameter.

Receiver Connector – The receiver connector is the optical connector provided at the input to the receiver that attaches to the receiver pigtail. The receiver connector description, as a minimum, includes:

- Connector manufacturer
- Connector type (e.g., SC, FC)
- Connector model number
- Connector classification (multimode, single-mode)
- Mating connector model number.

Generic connector requirements are in GR-326-CORE.

4.3.1.3.2 Transmission Properties

The transmission parameters to be specified for the receiver unit are as follows.

Receiver Sensitivity (P_R) – The worst-case value of the input optical power to the receiver (on the line side of the receiver module connector), specified as P_{R1} for standard operating conditions or P_{R2} for extended operating conditions, that is necessary to achieve the manufacturer-specified BER as measured using the procedure in TIA-526.3, *OFSTP-3, Fiber Optic Terminal Receiver Sensitivity and Maximum Receiver Input Power*.

The receiver sensitivity value specified includes the following performance degradation factors combined in a worst-case fashion:

- Manufacturing variations with temperature and aging drifts, including any degradation of the receiver connector
- Maximum transmitter power penalty resulting from the use of a transmitter with a worst-case extinction ratio r_e when operated under standard operating conditions (P_{R1}) or extended operating conditions (P_{R2})
- Maximum transmitter power penalty resulting from the use of a transmitter with a worst-case rise/fall time when operated under standard operating conditions (P_{R1}) or extended operating conditions (P_{R2}).

The receiver sensitivity does not include power penalties associated with dispersion (pulse broadening), multiple reflections or jitter.

Receiver Overload Point (R_{max}) – The maximum value of the input optical power to the receiver (on the line side of the receiver module connector, point R), when operated under standard operating conditions R_{max1} , or extended operating conditions R_{max2} , that the receiver will accept and maintain a 10^{-10} or 10^{-12} (as appropriate) BER as measured using the procedure in TIA-526.3. This value does not include the effects of a removable optical attenuator that may be needed to meet the maximum average received power values P_{Rmin} , specified for the applications in Tables 4-6 through 4-25, that could be placed on the receiver side of point R (Figure 4-1) and, consequently, be regarded as part of the receiver.

R_{max} may be greater than the value P_{Rmax} .

Maximum Dispersion (D_{SRmax}) – The maximum dispersion (in ps/nm), due to fiber length between points S and R (see Figure 4-1) that can be accommodated by a transmitter-receiver pair to meet the 10^{-10} or 10^{-12} (as appropriate) BER performance objective, when operated under standard operating conditions (D_{SRmax1}) or extended operating conditions (D_{SRmax2}).

D_{SRmax} may exceed the values that Tables 4-6 through 4-25 specify.

Dispersion Power Penalty (P_D) – The maximum power penalty associated with the worst-case increase in receiver input optical power level to account for the total pulse distortion due to ISI, Mode Partition Noise (MPN) and laser chirp at the specified bit rate, a BER of 10^{-10} or 10^{-12} (as appropriate), and maximum dispersion (D_{SRmax}), when operated under standard operating conditions (P_{D1}) or extended operating conditions (P_{D2}).

TIA-526.10, *OFSTP-10, Measurement of Dispersion Power Penalty in Single-Mode Systems*, contains a procedure for measuring dispersion power penalty. P_D is less than or equal to the value of P_O for each application that Tables 4-6 through 4-25 describe.

The receiver parameters P_R , P_D , D_{SRmax} and R_{max} are provided for a BER of 10^{-10} or 10^{-12} (as appropriate) as Worksheets 2 and 3 show. Optionally, the network provider may request these parameters for other BER values.

4.3.1.4 Attenuators

The supplier provides a full description of the attenuators that are to be used with the system, if needed. The specifications, as a minimum, include the following information:

Manufacturer
Model number.

The supplier also specifies these attenuator parameters:

U_{att} = Insertion loss (dB)
 OR_{att} = Worst-case value of attenuator reflectance (dB).

Attenuator criteria appear in GR-910-CORE.

4.3.1.5 Wavelength Division Multiplex Device

If a WDM device is offered, the supplier specifies:

Manufacturer
Model number
Number of channels.

The supplier also specifies the loss and minimum isolation between channels.

U_{WDM} = Worst-case value (in dB) of the all-inclusive loss associated with WDM equipment (at both ends), including all insertion and additional connector losses as well as other degradations. The allocations must include the effects of temperature, humidity and aging.

$Isol_{min}$ = Worst-case value (in dB) of the isolation associated with the WDM equipment (at both ends). The allocations must include the effects of temperature, humidity and aging.

The loss corresponds to the transmitter wavelength stated in Worksheets 2 and 3.

4.3.1.6 Passive Dispersion Compensation Devices

If a PDC device is offered, the supplier specifies:

Manufacturer
Model number
Manufacturer product change designation (e.g., issue or revision).

The supplier also specifies the following attenuator parameters:

$\lambda_{min} - \lambda_{max}$ = Wavelength range of operation
 D_{max} = The maximum compensation over its operating wavelengths
 D_{min} = The minimum compensation over its operating wavelengths
 U_{PDC} or G_{PDC} = The attenuation or gain of the device over its operating wavelengths.

Generic criteria for PDC devices are in GR-2854-CORE.

4.3.1.7 Optical Fiber Amplifiers

If a unit incorporates an Optical Fiber Amplifier (OFA) device, that device must be identified by a unique descriptor from which the following information can be determined, using the appropriate documentation:

Manufacturer
Model number
Optical device temperature controller
Type of fiber used for amplifier
Manufacturer product change designation (e.g., issue or revision)

The supplier also specifies the following OFA parameters:

$\lambda_{min} - \lambda_{max}$ = Wavelength range of operation
 G_{OFA} = Optical Fiber Amplifier Gain
 $P_{in_{min}}$ and $P_{in_{max}}$ = Minimum and maximum input power levels
 $P_{out_{min}}$ and $P_{out_{max}}$ = Minimum and maximum output power levels

Generic criteria for OFAs appear in GR-1312-CORE.

4.3.1.8 Connectors

The supplier specifies the following connector information:

Connector type (e.g., SC, FC)
Manufacturer
Model number
Connector classification (multimode, single-mode).

The supplier also specifies the following connector parameters:

U_{con} = Worst-case value of connector loss (in dB)
 ΔU_{con} = Maximum value of the difference in insertion loss between mating optical connectors of the same type and model, from the same manufacturer (i.e., connector variation)

OR_{con} = Worst-case value of connector reflectance (in dB).

In addition, the system integrator specifies the following:

N_{con} = Number of single-mode to single-mode connectors. This is the number recommended by a system integrator for a typical point-to-point regenerator section. This should not include the transmitter unit or receiver unit connectors, because they are already accounted for in P_T and P_R , respectively.

Generic connector requirements appear in GR-326-CORE.

4.3.1.9 Station Cable

Station cable represents the optical fiber cable that is used within a building environment to connect the outside plant optical fiber cable to the optical fiber system terminal equipment. The station cable may provide this optical path by means of some form of optical patch panel that allows optical path rearrangement to the outside plant fibers.

The supplier provides the following general station cable information:

Manufacturer
Model number
General fiber type (see TIA-4920000B)
Class of fiber (see TIA-4920000B)

and interconnection-related parameters:

Nominal mode field diameter and tolerance
Nominal cladding diameter and tolerance
Maximum cladding ovality
Maximum core/cladding concentricity error.

The interconnection-related parameters are needed to calculate the connection losses of field-installed splices and connectors.

The supplier also specifies the following transmission parameters:

U_{SM} = Worst-case end-of-life loss of single-mode regenerator station cable
 λ_{cc} = Cable cutoff wavelength, which must be below the minimum value of the transmitter central wavelength
 PMD = The maximum PMD coefficient for the fibers in the cable (applicable for OC-192 and OC-768)
 D_{chrom} = The chromatic dispersion coefficient for the fibers in the cable.

In addition, the system integrator specifies the following:

I_{SM} = Total length, on both ends of a regenerator section, of single-mode regenerator station cable.

GR-20-CORE and GR-409-CORE, *Generic Requirements for Premises Fiber Optic Cable*, contain criteria related to the optical and physical parameters for fiber.

4.3.1.10 Safety Margin

M = Safety margin (in dB) for unexpected losses, to be determined by the system integrator for a specific application.

M does not include penalties for expected losses and degradations (e.g., laser aging or reflections) because these effects are already included in the appropriate transmission parameters, according to the definitions in this document. The specifications in Section 4.2 assume $M = 0$. When a system integrator desires, a non-zero value of M may be used. In this case, the value is to be included in determining the effective attenuation as Tables 4-6 through 4-25 specify. Selecting a non-zero value of M for the system design with transmitters meeting the specifications in Tables 4-6 through 4-25 and receivers meeting the minimum value for receiver overload P_{Rmax} in the tables may require optical attenuators to protect the receiver from damage caused by optical power overload.

4.3.2 Cable Transmission Design Information

For a given cable type, the cable supplier provides the following two categories of cable transmission information.

1. Parameters for specific applications. These are specified at the time of initial installation.
2. Global loss and chromatic dispersion characteristics. Because these may not be guaranteed values, the user should use them for initial feasibility studies only, and should measure the parameters for a specific upgrade. If a known upgrade is planned at the time of initial installation, the required parameters should be specified.

The following sections explain each category.

4.3.2.1 Parameters for a Specific Application

This section discusses the cable transmission parameters to be provided by the system integrator and by the suppliers. These are summarized in Worksheet 4. The parameters discussed in this section are to be given as worst-case values.

The system integrator specifies the following:

- Type of application (e.g., aerial, buried or underground)
- Temperature range (e.g., -7°C to 40°C for underground, -40°C to 77°C for aerial environments)
- The cabled fiber reel length l_R (in km)
- The nominal central wavelength (λ_{Tnom}) and central wavelength range (λ_{Tmin} to λ_{Tmax}) corresponding to the terminal equipment to be used
- The type of splice, if appropriate
- The following splice loss information:


U_S = Maximum allowable splice loss (in dB/splice) at 23°C

U_{ST} = Effect of temperature on splice loss (in dB/splice) at the worst-case temperature conditions, over the specified cable operating temperature range (may be zero if U_S already includes corrections for U_{ST}).

The supplier specifies the following:

- Designation of the cable
- Maximum cable cutoff wavelength, λ_{cc} , (in nm), with cable deployment conditions as shown in Figure 4-4.

Figure 4-4 Cable Configuration for Cabled Fiber Cutoff Wavelength Measurement



Note:

The purpose of the 76-mm (3-inch) loop at each fiber end is to simulate the splice organizer.

The cutoff wavelength of a cabled optical fiber demarcates the wavelength region above which the fiber supports propagation of only a single mode and below which multiple modes are supported. Operation below the cutoff wavelength may result in modal noise, modal distortion (increased pulse broadening), and improper operation of connectors, splices and WDM couplers. For these reasons, the system operating wavelength range, dictated by the transmitter central wavelength range described in Section 4.2.3.1, must be greater than the maximum allowed cutoff wavelength to ensure the system is operating entirely in the fiber’s single-mode regime. In general, the highest value of cabled fiber cutoff wavelength, λ_{ccmax} , will be found in the shortest installation or repair cable length. A criterion that ensures a system is free from high cutoff wavelength problems is:

$$\lambda_{ccmax} < \lambda_{Tmin} \tag{4-11}$$

where λ_{Tmin} is defined in Section 4.2.3.1.


The supplier also specifies the cable loss, dispersion and interconnection parameters discussed below.

- Cable loss parameters:

U_c = Worst-case end-of-life cable loss (in dB/km at 23°C) at the transmitter’s nominal central wavelength λ_{Tnom} . This includes splicing loss caused by the fiber or cable manufacturing process

- U_{λ} = The largest increase in cable loss (in dB/km at 23°C) above U_c that occurs over the transmitter's central wavelength range (λ_{Tmin} to λ_{Tmax}). Figure 4-5 illustrates the determination of this parameter.
- U_{CT} = Effect of temperature on end-of-life cable loss (in dB/km) at the worst-case temperature conditions over the specified cable operating temperature range.

Figure 4-5 Wavelength Dependent Attenuation Characteristics



Note: This figure is only an example. In actuality, U_{λ} may occur at λ_{Tmax} , λ_{Tmin} , or another wavelength between λ_{Tmin} and λ_{Tmax} relative to λ_{Tnom} .

- Dispersion parameters:

λ_{0min} , λ_{0max} = Minimum and maximum values of the zero-dispersion wavelength

S_{0max} = Maximum value [in ps/(nm² × km)] of the zero-dispersion slope (i.e., the dispersion slope at the zero-dispersion wavelength)

D_{max} = Absolute value of the worst case of chromatic dispersion [in ps/(nm × km)] over the transmitter central wavelength range and over the allowed variation in the cable type's dispersion

PMD = The worst case of polarization mode dispersion coefficient [in ps/(km^{0.5})] over the transmitter central wavelength range.

- Interconnection-related parameters:

- Nominal mode field diameter and tolerance
- Nominal cladding diameter and tolerance
- Maximum cladding ovality
- Maximum core/cladding concentricity error

which are needed to calculate the connection losses of field-installed splices and connectors.

If a future system upgrade is planned, the user should specify the performance required for the upgrade. The supplier furnishes the information of Worksheet 4 for the original application and the future upgrade.


4.3.2.2 Global Fiber Parameters

A future application of an installed cable may arise that was not anticipated when the cable was purchased (e.g., an unforeseen upgrading of terminal equipment, rerouting of traffic). In such cases, global fiber characteristics may be helpful for indicating to the user the feasibility of the cable's proposed application.

The fiber global loss and dispersion characteristics are curves that depict these parameters as functions of wavelength. They are to be provided over the fiber's useful wavelength regions (e.g., from 1260 to 1360 nm and from 1430 to 1580 nm). The global characteristics should show typical values because the only worst-case values are those specified when the cable was purchased. Thus, if the global characteristics indicate the new application to be feasible, measurements still must be taken to verify that the cable will perform properly.

For the fiber global loss characteristics, the supplier provides a graph similar to [Figure 4-6](#), which shows the loss at 23°C and the maximum temperature effect over the cable's operating temperature range, U_{CT} , as a function of wavelength. Also, the system integrator describes the maximum allowable splice loss, U_S , at 23°C and the maximum temperature effect on splices, U_{ST} , as a function of wavelength. These characteristics are provided for underground, buried and aerial applications, over the temperature ranges as the user specifies.

Figure 4-6 Example Global Fiber Attenuation Characteristics




For the global dispersion characteristic, the data that [Section 4.3.2.1](#) and [Worksheet 4](#) reference can be used with the following formulas to approximate the fiber's chromatic dispersion coefficient D , in ps/(nm × km), as a function of wavelength. The extreme limits, $D_1(\lambda)$ and $D_2(\lambda)$, are given as:

$$D_1(\lambda) = \frac{S_{o \max}}{4} \left(\lambda - \frac{\lambda_{o \min}^4}{\lambda^3} \right) \quad (4-12)$$

$$D_2(\lambda) = \frac{S_{o \max}}{4} \left(\lambda - \frac{\lambda_{o \max}^4}{\lambda^3} \right) \quad (4-13)$$

[Figure 4-7](#) shows examples of the respective curves.

Figure 4-7 Example Global Fiber Dispersion Characteristics



4.3.3 Fiber Optic System Transmission Design and Analysis

4.3.3.1 Design Approach

This section describes a methodology for engineering the route layout for a single-mode fiber optic transmission system. The methods address the design of a regenerator section and are intended to ensure that all installed fiber paths that meet cable specifications in every regenerator section will be usable at certain proposed transmission rates. A regenerator section, defined in [Figure 4-8](#), is composed of:

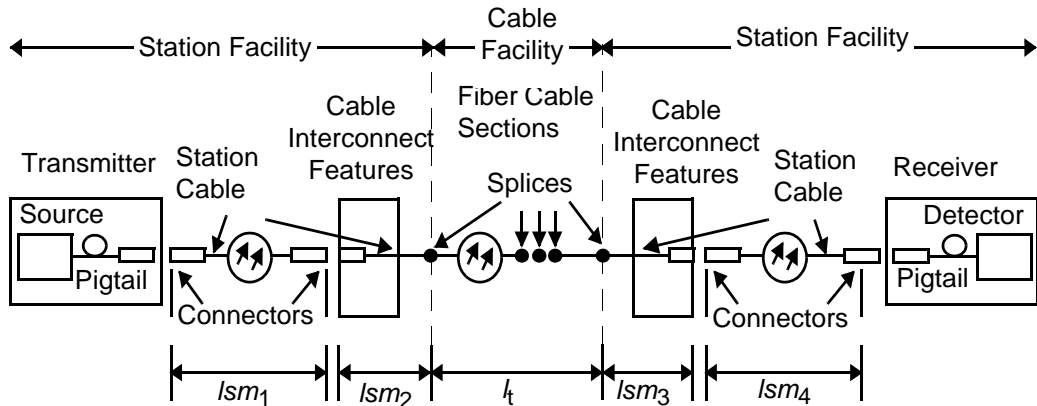
1. Regenerators, connectors, station cables and a fiber interconnection feature at each regenerator station
2. Fiber cable sections, joined by splices, between regenerator stations.

The fiber interconnection feature is the interface between the regenerator station equipment and the fiber cable. It provides connector access to each fiber in the cable for maintenance and restoration purposes.

This section describes how the transmission parameters of Sections 4.3.1 and 4.3.2 are used in the design and analysis of regenerator sections. The computations are based on two constraint relations, one on the system's loss budget and the other on dispersion tolerance (each of which is discussed). Finally, the design and analysis methodology is discussed.

Selecting a non-zero value of M for the system design with transmitters meeting the specifications in Tables 4-6 through 4-25 and receivers meeting the minimum value for receiver overload P_{Rmax} in those tables may require optical attenuators to protect the receiver from damage caused by optical power overload.

Figure 4-8 Engineering of a Fiber Optic Regenerator Section



$$l_{sm} = l_{sm1} + l_{sm2} + l_{sm3} + l_{sm4} = \text{Total Length (in km) of Single-Mode Station Cable}$$

$$l_t = \text{Total Sheath Length (in km) of Spliced Fiber Cable}$$

4.3.3.2 Loss Budget Constraint

Figure 4-9 shows a model for transmission in a regenerator section including transmitter power, receiver sensitivity and various sources of loss. For route layout design, it is convenient to divide the transmission path at the fiber distributing frame into terminal or regenerator station equipment and spliced fiber cable. The system gain, G , of terminal or regenerator station equipment, and the loss, L , introduced by the fiber path between regenerator interfaces, can then be determined. If L is less than or equal to G , then the embodied transmission objectives can be attained. An attenuator may be needed if L is much smaller than G .

In the following equations, the parameters that enter into G and L are specified by their worst-case values.

The equation for the system gain of the terminal or regenerator station equipment is

$$G = P_T - P_R - P_D - M - U_{WDM} - I_{SM} U_{SM} - N_{con} U_{con} \quad (4-14)$$

where the terms are as defined in Section 4.3.

The equation for the loss of the fiber cable within a regenerator section is

$$L = l_t(U_c + U_{cT} + U_\lambda) + N_S(U_S + U_{ST}) \quad (4-15)$$

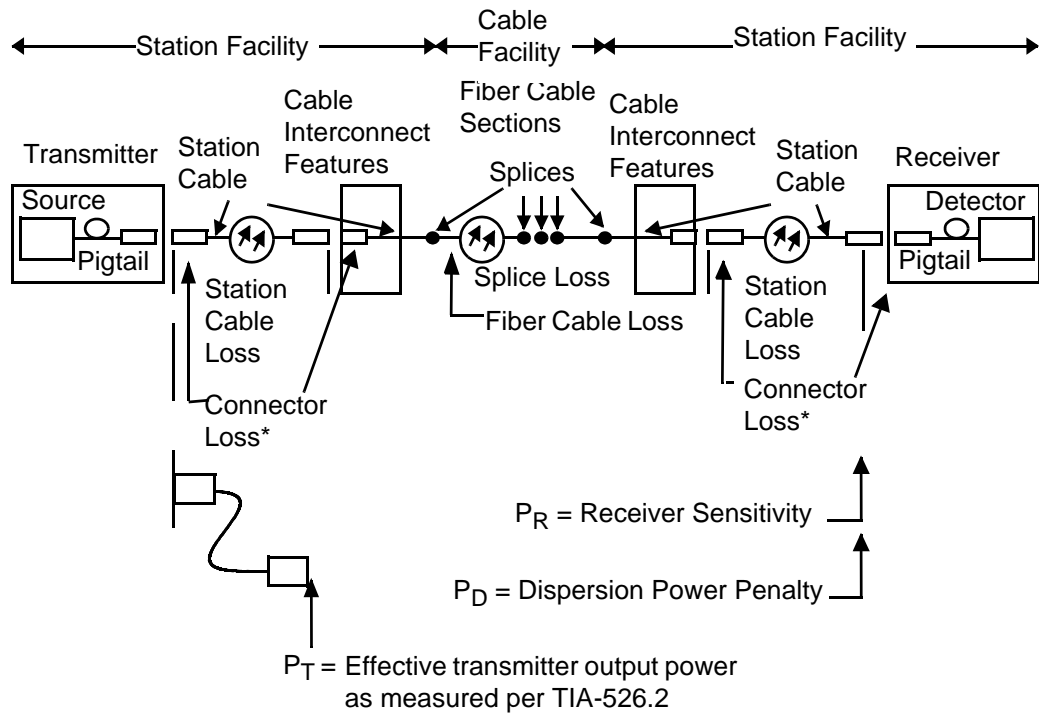
where:

- l_t = Total sheath length of spliced fiber cable in km. This value may include an allowance for cable repair purposes.
- N_S = Number of splices in fiber cable of length l_t (including those at the fiber distributing frame at both ends of the regenerator section), plus any allowance for additional splices for cable repair purposes.

The remaining terms are as defined in Section 4.3.2.1.

Section 4.3.3.4 discusses transmission analysis and design calculations using G and L .

Figure 4-9 Regenerator Section, Sources of Loss



* By convention, the total connector loss is assigned to the connector mate receiving transmitted power

4.3.3.3 Dispersion Limited Length

In general, for single-mode systems operating at relatively low bit rates, regenerator section length may be expected to be limited by loss and not by dispersion. At higher bit rates, however, the length may be limited by dispersion, which arises from the combined effects of chromatic dispersion, mode-partition noise and laser chirp. Thus, it is desirable to check whether a regenerator section's length may be limited by dispersion.

A regenerator section will not be limited by dispersion as long as the following holds:

$$D(\lambda_t) \times l \leq D_{SRmax} \quad (4-16)$$

where:

- $D(\lambda_t)$ = Fiber chromatic dispersion coefficient evaluated at λ_t
[in ps/(nm×km)]
- l = Fiber length (in km)
- D_{SRmax} = Maximum dispersion (in ps/nm) as defined in Section 4.3.1.3.2.
For dispersion-limited systems, Tables 4-6 through 4-25 specify maximum values for D_{SRmax} .

Taking into account the variation in the fiber chromatic dispersion coefficient over the transmitter's central wavelength range, the dispersion limited length is given by

$$l_D = \frac{D_{SRmax}}{D_{max}} \quad (4-17)$$

where:

- D_{max} = Absolute value of the worst case of chromatic dispersion coefficient [in ps/(nm × km)] over the specified range of the transmitter central wavelength and over the specified variation in the cable type's dispersion
- l_D = Dispersion-limited length (in km).

The effect of dispersion is accounted for in the equation for G (Equation 4-14) via the dispersion power penalty P_D (in dB).

4.3.3.4 Design and Analysis Methodology

The transmission design and analysis methodology for fiber optic systems focuses on the individual sections making up the overall system. The design and analysis calculations are performed using specific combinations of terminal equipment and cable parameters.

The design algorithm deals with a proposed fiber optic system route with given distances between terminal and regenerator locations. This document assumes SONET regenerator sections to be designed for a BER of 1×10^{-10} or 1×10^{-12} (depending on the application). However, individual network providers may have different BER requirements. For each regenerator section, one must determine whether a particular combination of terminal equipment and cable can work together (at the desired BER) without exceeding any loss or dispersion limitation. The steps involved in the design algorithm are as follows.

- A. From Worksheet 2 or 3, select the receiver parameters P_R , P_D and D_{SRmax} corresponding to the desired BER.
- B. Calculate G and L (Equations 4-14 and 4-15). If $G - L \geq 0$, then the regenerator section is not limited by loss.

- C. Calculate the end-to-end regenerator section dispersion $D_{max} \times l$. If this is less than or equal to D_{SRmax} , the regenerator section is not limited by dispersion.

If the calculations from steps **B** and **C** are satisfactory, the terminal equipment-cable combination is acceptable.

The analysis algorithm deals with calculating maximum regenerator section lengths. Such calculations may be used in comparing various combinations of terminal equipment and cable. It is suggested that the maximum regenerator section lengths be calculated as the distance between splices is varied, ranging from 0.5 km up to the length of a reel of cable. The steps involved in the analysis algorithm are as follows:

- A. Select a set of receiver parameters P_R , P_D and D_{SRmax} corresponding to the desired BER value from Worksheet 2 or 3
- B. Calculate the length at which $G = L$ (Equations 4-14 and 4-15)
- C. Calculate the dispersion-limited length (Equation 4-17)
- D. Take the maximum regenerator section length to be the smaller of the two lengths calculated in steps **B** and **C** above.

4.4 Electrical Interface Specifications

In some applications, it may be desirable to interconnect SONET NEs using standard electrical interfaces. Central office engineering considerations limit the feasibility of standard electrical interfaces to the STS-1 and STS-3 levels of the SONET hierarchy. These interfaces are based on the following.

- One coaxial line is used for each direction of transmission. Reference cable is 75- Ω coaxial cable with tinned copper shield (AT&T Technologies, Inc. 728A or equivalent).
- Since the SONET signal is scrambled (Section 5.1.3), the electrical interface signal can be adequately modeled as one having a random occurrence of ones and zeros with a mark probability nominally equal to 0.5. This allows specification of the power level at the interface in terms of a wideband power measurement rather than a narrow-band measurement power (as is used for DS1 interfaces). A wideband measurement is typically simpler than a narrow-band measurement because it does not require tightly controlled bandpass filters with specific roll-off characteristics.⁸

8. In the DS_n interface criteria in GR-499-CORE, requirements are given for both power levels and pulse amplitudes. Similarly, for STS-3 electrical signals, a requirement is given for the pulse amplitude at transmitter, and another requirement is given for the power level at the interface. For STS-1 electrical signals, only an interface power level requirement is given. A pulse amplitude requirement is not necessary for conformance testing purposes; however, in some situations (e.g., during trouble-shooting procedures) it may be useful (and simple) to measure the STS-1 pulse amplitude. An STS-1 electrical signal that meets the power level requirement is expected to have pulse amplitudes (at the interface) in the range of approximately 0.45 to 0.9 V. Note that this range (0.45 to 0.9 V) is a conservative estimate calculated for signals having “worst-case” pulse shapes. Therefore, signals with pulse amplitudes somewhat outside this range may still meet the power level requirement.

- The waveform at the interface is specified by an eye diagram mask, which at the STS-1 level is compatible with an isolated pulse template that is also specified. In general, STS-1 electrical signals meeting the pulse template specification (which applies only to isolated pulses within the signal) should also meet the eye diagram mask specification (which applies to the entire signal), and signals meeting the eye diagram mask specification should also meet the pulse template specification. In rare cases, however, a signal with a conforming eye diagram may not meet the pulse template specification, or vice versa. It is therefore recommended that the electrical interfaces be designed and tested against both specifications. In addition, the eye diagram mask specification can be very useful for the purpose of performing a visual check of the waveform at the interface during trouble isolation testing. ANSI T1.102 describes procedures for checking conformance with the pulse and eye diagram masks.

CR4-24 [128v2] A SONET NE may be required to support STS-1 electrical interfaces.

CR4-25 [910] A SONET NE may be required to support STS-3 electrical interfaces.

R4-26 [129v2] If a SONET NE supports STS-1 electrical interfaces, the following apply:


- The transmitter shall generate an interface signal that meets the criteria in [Table 4-26](#) for the entire range of interconnect cable lengths of 0 to 450 feet between the transmitter and the interface
- The receiver shall provide the equivalent of a $75\ \Omega$ ($\pm 5\%$) resistive termination for the incoming signal, and shall accept any interface signal that conforms to the criteria in [Table 4-26](#) (at the interface) and then propagates through a jumper cable (if used, may be up to 27 feet) and/or any additional length of interconnect cable up to 450 feet.

A resistive test load of $75\ \Omega$ ($\pm 5\%$) is used at the interface for the evaluation of the parameters specified above and in [Table 4-26](#). Note that a SONET NE's STS-1 electrical interfaces may switch in Line Build Out (LBO) circuits to maintain the equivalent interconnect distance between 225 and 450 feet in cases where the actual cable length is less than 225 feet.

Table 4-26 STS-1 Electrical Interface Criteria

Parameter	Criteria
Line Rate	51.840 Mb/s (synchronized to the NE clock; see Section 5.4)
Line Code	Bipolar with Three-Zero Substitution (B3ZS, see Section 9.1.1.4 of GR-499-CORE)
Pulse Shape	The shape of every pulse that approximates an isolated pulse (i.e., a pulse preceded by two zeros and followed by one or more zeros) shall fit within the limits of the pulse template in Figure 4-10 . In addition, an eye diagram of the interface signal shall fit within the limits of the eye diagram mask in Figure 4-11 .
Power Level	The wideband power level measured using a low-pass filter with a flat passband and a 3-dB cutoff frequency of 207.36 MHz shall be between -2.7 dBm and +4.7 dBm.
DC Offset	There shall be no Direct Current (DC) power flow across the interface.

Figure 4-10 STS-1 Electrical Interface Pulse Mask




	Time axis range (UI)	Normalized amplitude equation
Upper curve	$-0.85 \leq T \leq -0.68$	0.03
	$-0.68 \leq T \leq 0.26$	$0.5 \{1 + \sin [(\pi/2)(1 + T/0.34)]\} + 0.03$
	$0.26 \leq T \leq 1.40$	$0.1 + 0.61 e^{-2.4 (T - 0.26)}$
Lower curve	$-0.85 \leq T \leq -0.36^a$	-0.03
	$-0.36^a \leq T \leq 0.36$	$0.5 \{1 + \sin [(\pi/2)(1 + T/0.18)]\} - 0.03$
	$0.36 \leq T \leq 1.40$	-0.03

Note:

- a. At least some versions of the STS-1 electrical interface standards have incorrectly listed this value as -0.38.

Figure 4-11 STS-1 Electrical Interface Eye Diagram Mask



Outer region corner points			Inner region corner points		
Point	Time	Amplitude	Point	Time	Amplitude
A	-0.5	0.426	I	-0.245	0.214
B	-0.261	0.904	J	-0.187	0.455
C	-0.136	1.03	K	-0.104	0.67
D	-0.028	1.03	L	-0.017	0.67
E	0.094	0.883	M	0.077	0.581
F	0.187	0.723	N	0.18	0.14
G	0.31	0.566	O	-0.054	0.16
H	0.5	0.426			

Note – Both the inner and outer regions are symmetric about the zero-amplitude axis.

R4-27 [130v2] If a SONET NE supports STS-3 electrical interfaces, the following apply:


- The signal at the output of the transmitter (at the cable terminal) shall have a nominal rectangular pulse shape with a peak amplitude of 0.5 V (+10%) and maximum rise/fall times of 2 ns, as Figures 4-12 and 4-13 show.
- The transmitter shall generate an interface signal that meets the criteria in Table 4-27 for the entire range of interconnect cable lengths of 0 to 225 feet between the transmitter and the interface.
- The receiver shall provide the equivalent of a 75 Ω ($\pm 5\%$) resistive termination for the incoming signal, and shall accept any interface signal that conforms to the criteria in Table 4-27 (at the interface) and then propagates through a jumper cable (if used, may be up to 27 feet) and/or any additional length of interconnect cable up to 225 feet.

A resistive test load of 75 Ω ($\pm 5\%$) is used at the interface for the evaluation of the parameters specified above and in Table 4-27.

Table 4-27 STS-3 Electrical Interface Criteria

Parameter	Criteria
Line Rate	155.520 Mb/s (synchronized to the NE clock; see Section 5.4)
Line Code	Coded Mark Inversion (CMI, see Section 9.1.3 of GR-499-CORE)
Pulse Shape	An eye diagram of the interface signal shall fit within the limits of the eye diagram mask specification in Figure 4-14. The eye diagram is obtained by using a triggering signal at twice the bit rate (i.e., at 311.040 MHz).
Power Level	A wideband power level measured using a low-pass filter with a flat passband and a 3-dB cutoff frequency of 311.04 MHz shall be between -2.5 dBm and +4.3 dBm.
DC Offset	There shall be no DC power flow across the interface.


Figure 4-12 STS-3 Transmitter Pulse Mask Corresponding to a Binary Zero



Notes:

1. The mask does not include the over/undershoot tolerance of 10%.
2. The nominal zero level can be adjusted by $\pm 0.05\text{ V}$ to meet the limits of the mask.


Figure 4-13 STS-3 Transmitter Pulse Mask Corresponding to a Binary One



Notes:

1. The mask does not include the over/undershoot tolerance of 10%.
2. The nominal zero level can be adjusted by ± 0.05 V to meet the limits of the mask.
3. The mask also applies to the inverse pulse.

Figure 4-14 STS-3 Eye Diagram Mask



Corner Points		
Point	Time	Amplitude
A	0.125	0.0
B	0.225	0.25
C	0.275	0.25
D	0.35	0.0
E	0.0	1.1
F	0.5	1.1

Note: The mask is symmetric about zero amplitude. In addition, an identical mask appears in the time interval from -0.5 to 0 UI (and from 0.5 to 1.0 UI).

5 Network Element Architectural Features

This section describes various architectural features required in SONET NEs, particularly:

- Multiplexing procedures
- Overhead function usage
- APS
- Network synchronization
- Framing
- Jitter and wander performance.

5.1 Multiplex Procedures

5.1.1 Interleaving

An STS-N module either can be created directly (e.g., by the equipment that creates an STS-12 and maps ATM traffic into the STS-12c SPE), or it can be formed by byte-interleaving lower-level modules (e.g., STS-1s and STS-Ms). For STS-Ns that are formed by byte-interleaving lower-level modules, the following requirements are applicable:

R5-1 [131] Before byte-interleaving to form an STS-N, the transport overhead byte positions of all the constituent STS-1s and STS-Ms shall be frame aligned.

The alignment of the STS-1s and STS-Ms is accomplished by adjusting the STS payload pointers to reflect the new relative positions of the STS SPEs.¹

Note that for the purposes of these requirements, it is useful to assume that an NE that forms an STS-N where N is greater than 3 will first logically interleave any STS-1 inputs (in sets of three consecutive STS-1s) to form STS-3 modules, and then interleave those STS-3 modules and any other STS-M inputs to form the STS-N. However, there are no requirements concerning the internal architectures of SONET NEs (e.g., an NE could directly interleave STS-1 inputs in the appropriate order to form an STS-N). The important point is that the output byte sequence must be as shown in [Figure 5-1](#).

-
1. In the SONET interface layer model described in [Section 3.3](#), it would be the responsibility of the Line layer to align and interleave the signals that it receives from the STS path layer. These criteria are written based on the assumption that the NE uses STS path and Line layer signals that contain all of the transport overhead byte positions (e.g., that an internal path layer signal carrying an STS-1 SPE has a nominal bit rate of 51.84 Mb/s); however as discussed in [Section 3.3](#), there are no criteria concerning the internal signals used by an NE. If the NE uses some other type of internal signals, then it must perform an equivalent process so that the structure of the interface signals (i.e., the OC-N or STS-N electrical signals transmitted by the NE) is independent of the type of internal signals used.

- R5-2 [132]** To form an STS-3 from STS-1s, three STS-1s shall be interleaved, one byte at a time. The first byte of the STS-3 shall be the A1 byte from the first STS-1, followed sequentially by the A1 byte from the second STS-1, and then the A1 byte from the third STS-1.
- R5-3 [133]** To form a higher-level STS-N ($N > 3$) from lower-level STS-Ms ($3 \leq M < N$), the STS-Ms shall be interleaved $M/3$ bytes at a time. The output byte sequence shall be as shown in [Figure 5-1](#).

The preceding requirements have several implications concerning the transport of STS-Mc SPEs in OC-N ($M < N$) signals. Since the STS-Mc SPE is a single entity (see [Section 3.2.3](#)), the module that contains it is considered an STS-M input to the byte-interleaver (which has an STS-N as its output). Based on these requirements, the STS-M input must occupy the byte positions corresponding to M consecutive STS-1 inputs, and it must start in a byte position corresponding to an STS-1 number $[(X+1),1]$, where $0 \leq X \leq (N-M)/3$ (using the two-level numbering scheme shown in [Figure 5-1](#)). In addition, to simplify the processing capabilities required of NEs that originate, pass and terminate concatenated signals, the possible starting positions for an STS-Mc SPE in an STS-N are further limited by the following requirement.

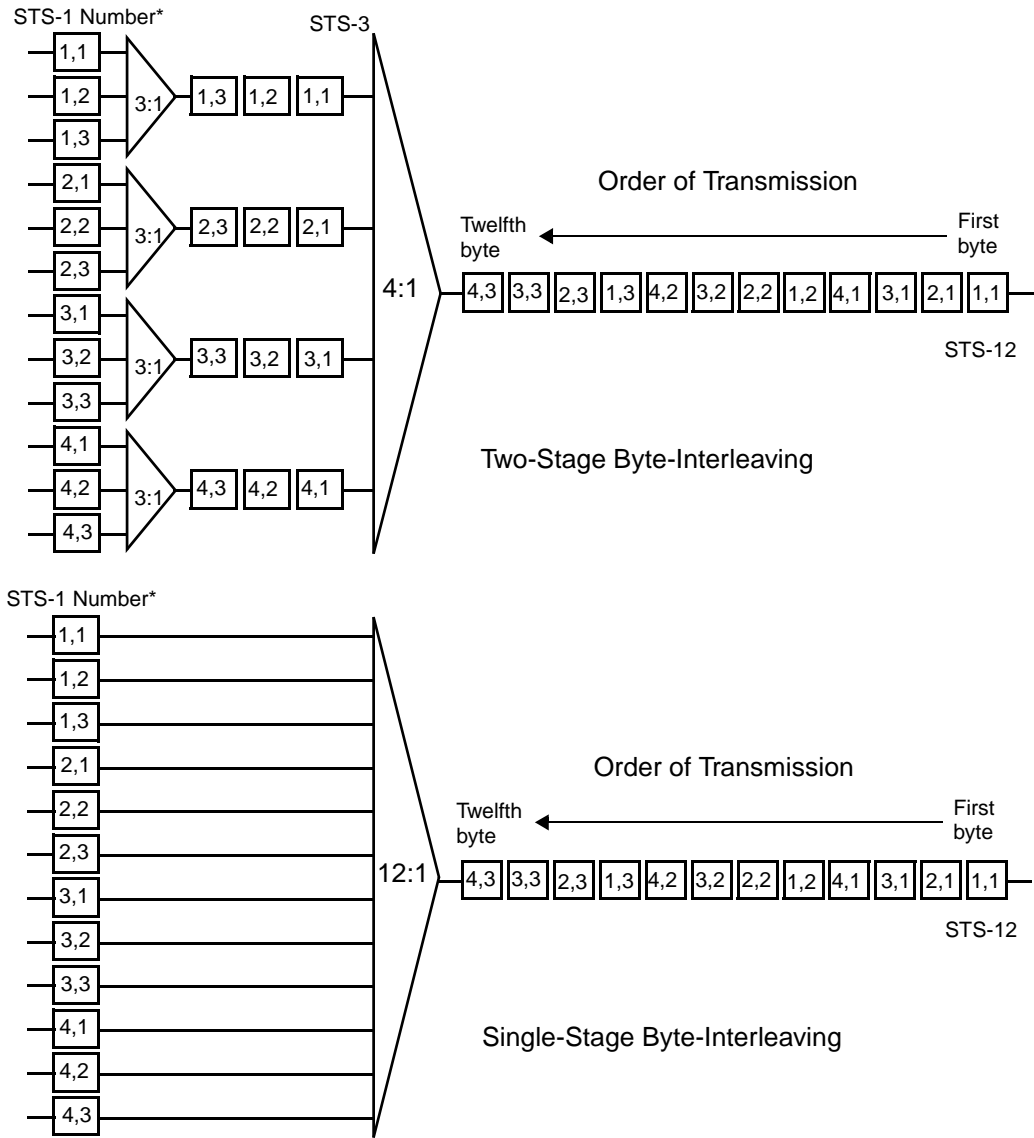
- R5-4 [911v3]** Each STS-Mc SPE in an STS-N shall be completely contained in one of Y groups of M STS-1s, where $Y = (N \div M)$.

The effect of **R5-4 [911v3]** is illustrated in [Table 5-1](#), which lists the possible starting positions for the various size STS-Mc SPEs that could be contained in an OC-768 signal. Note that prior to Issue 4 of this document, all of the possible values of M from 3 to 48 were considered to be valid and were therefore included in previous versions of [Table 5-1](#). However, few (if any) SONET products supported the transport of STS-Mc SPEs other than those that are now considered valid (i.e., other than STS-3c, STS-12c, STS-48c, STS-192c and STS-768c SPEs). Thus, a product's use of such values of M severely restricted its applicability in multi-product configurations. This, along with the definition of the virtual method of concatenating STS-1 or STS-3c SPEs (see [Section 3.2.3.2](#)), led to the removal from this document of 6, 9, 15, 18, 21, ..., 42 and 45 as possible values of M .

- R5-5 [912]** A SONET NE that provides the capability to terminate or pass a particular size STS-Mc SPE shall be capable of performing that function on an STS-Mc SPE that starts in any of the allowed starting positions defined in **R5-4 [911v3]** and shown in [Table 5-1](#).

[Figure 5-2](#) illustrates an interleaving example in which nine STS-1s, one STS-3c (carrying an STS-3c SPE), one STS-12c (carrying an STS-12c SPE), and other STS-1 and STS-Mc modules (which are unspecified in the example) are interleaved to form an STS-48. In the example, the term "first 4 bytes of STS-12c 4,1" means the first four bytes that would be transmitted if that STS-12 were converted to an OC-12 signal instead of being multiplexed into an STS-48. Similarly, "first byte of STS-3c 8,1" means the first byte that would be transmitted if that STS-3 were converted to an STS-3 electrical or OC-3 signal, etc.

Figure 5-1 Example of Byte-Interleaving Sequence, STS-12



* The STS-1s in this figure are numbered using the two-level “STS-3 #/STS-1 #” numbering scheme. This scheme is based on an STS-3 number followed by the number of the STS-1 within the STS-3, and is one of the two schemes that may be used by SONET NEs (see Section 6.1.2). The other scheme that may be used is one in which the STS-1s are numbered “1 to N in order of appearance at the input to the byte-interleaver”. Using that scheme, the order of transmission in this figure would be 1 (first byte), 4, 7, 10, 2, 5, 8, 11, 3, 6, 9, 12 (twelfth byte).

Table 5-1 Possible Starting Positions for an STS-Mc SPE in an OC-768 Signal

STS-1 Number X,1 ^a	STS-1 Number ^b	STS-3c SPE	STS-12c SPE	STS-48c SPE	STS-192c SPE	STS-768c SPE
X = 1; 65; 129; 193	1; 193; 385; 577	Y	Y	Y	Y	Y; NO; NO; NO
X = 2; 66; 130; 194	4; 196; 388; 580	Y	NO	NO	NO	NO
X = 3; 67; 131; 195	7; 199; 391; 583	Y	NO	NO	NO	NO
X = 4; 68; 132; 196	10; 202; 394; 586	Y	NO	NO	NO	NO
X = 5; 69; 133; 197	13; 205; 397; 589	Y	Y	NO	NO	NO
X = 6; 70; 134; 198	16; 208; 400; 592	Y	NO	NO	NO	NO
X = 7; 71; 135; 199	19; 211; 403; 595	Y	NO	NO	NO	NO
X = 8; 72; 136; 200	22; 214; 406; 598	Y	NO	NO	NO	NO
X = 9; 73; 137; 201	25; 217; 409; 601	Y	Y	NO	NO	NO
X = 10; 74; 138; 202	28; 220; 412; 604	Y	NO	NO	NO	NO
X = 11; 75; 139; 203	31; 223; 415; 607	Y	NO	NO	NO	NO
X = 12; 76; 140; 204	34; 226; 418; 610	Y	NO	NO	NO	NO
X = 13; 77; 141; 205	37; 229; 421; 613	Y	Y	NO	NO	NO
X = 14; 78; 142; 206	40; 232; 424; 616	Y	NO	NO	NO	NO
X = 15; 79; 143; 207	43; 235; 427; 619	Y	NO	NO	NO	NO
X = 16; 80; 144; 208	46; 238; 430; 622	Y	NO	NO	NO	NO
X = 17; 81; 145; 209	49; 241; 433; 625	Y	Y	Y	NO	NO
:	:	:	:	:	:	:
:	:	:	:	:	:	:
X = 64; 128; 192; 256	190; 382; 574; 766	Y	NO	NO	NO	NO


Notes:

- The four STS-1 numbers shown in the first column of each row are from the two-level “STS-3 #/STS-1 #” numbering scheme that SONET NEs may use (see [Section 6.1.2](#)). Note that an STS-Mc SPE cannot start in an STS-1 numbered (1,2), (1,3), (2,2), (2,3), ..., or (256,3).
- The four STS-1 numbers shown in the second column of each row are from the single-level “1 to N in order of appearance at the input to the byte-interleaver” numbering scheme that SONET NEs may use (see [Section 6.1.2](#)). Note that an STS-Mc SPE cannot start in an STS-1 numbered 2, 3, 5, 6, ..., or 768.

Y = STS-Mc SPE can start in that STS-1

NO= STS-Mc SPE cannot start in that STS-1

Figure 5-2 Byte-Interleaving Example, Multiple Level Inputs



5.1.2 Concatenation

In all of the uses of STS-Nc SPEs currently defined in SONET, the STS-Nc SPE contains a single payload mapping (e.g., the ATM mapping into an STS-Nc SPE described in Section 3.4.2.2.1), and therefore it is created in a single STS PTE. If

future uses are defined that require an STS-Nc SPE to carry multiple payloads mapped into STS-1 or STS-Mc ($M < N$) SPEs, then the necessary additional criteria will be added to this section.²

Note that [Section 3.2.3.1](#) contains criteria related to the structure of STS-Nc SPEs, [Section 3.4.2](#) describes the mappings of super-rate payloads into STS-Nc SPEs, and [Section 3.5.1.4](#) describes the use of STS payload pointers to identify the STS-1s that make up an STS-Nc.

5.1.3 Scrambling

Currently, all SONET optical interface signals use NRZ line coding, and therefore must be scrambled to assure an adequate number of transitions (zeros to ones, and ones to zeros) for such purposes as line rate clock recovery at the receiver. SONET electrical interface signals use line codes that assure adequate transitions (i.e., B3ZS and CMI, see [Section 4.4](#)); however, they are also scrambled for consistency between the electrical and optical interfaces. In both cases, the scrambler that is used is a frame synchronous scrambler that can be applied identically at the transmitter and the receiver.


R5-6 [134v2] SONET interface signals shall be scrambled (i.e., scrambled at the transmitter and descrambled at the receiver) using a frame synchronous scrambler of sequence length 127, operating at the line rate.

- The generating polynomial for the scrambler shall be $1+x^6+x^7$.
- The scrambler shall be reset to '111111' on the most significant bit of the byte following the Z0 byte in the Nth STS-1 (i.e., the byte following the last Z0 byte). That bit and all subsequent bits to be scrambled shall be added, modulo 2, to the output from the x^7 position of the scrambler, as shown in [Figure 5-3](#).
- For an STS-N where N is less than 768, the scrambler shall run continuously from the bit described in the previous bullet item on throughout the remainder of the STS-N frame.
- For an STS-768, the scrambler shall run continuously through the least significant bit of the last Z0 byte in the next STS-768 frame (i.e., to the point at which it is reset for use in the remainder of that frame). In addition, in this case all bytes except the 64 A1 bytes and the 64 A2 bytes (see [Section 3.3.2.1](#)) shall be scrambled.

Note that for an STS-N where N is less than 768, the $2 \times N$ framing bytes (A1 and A2), the Section trace byte (J0), and the N-1 Section growth (Z0) bytes are not scrambled.

2. Such a use has been defined in SDH, but no equivalent use is defined in SONET. In the SDH case, a single VC-4 (the SDH equivalent to an STS-3c SPE) can be used to transport three independent TUG-3s. Each TUG-3 is equivalent to an STS-1 SPE with the fixed stuff columns removed and a new column containing a payload pointer added. This additional pointer is used to identify the start of the VC-3, which in turn can contain (for example) an asynchronously mapped DS3 signal.

Figure 5-3 Frame Synchronous Scrambler (Functional Diagram)



5.1.4 An Example of STS-1 and OC-N Signal Composition

Figure 5-4 shows one possible set of stages in the formation of an STS-1 and an OC-N signal. The NE in this example has STE, LTE and PTE functionality, and is also capable of passing nonterminated STS SPEs through. It does not support FEC, the extended Line DCC (which is applicable only for N equal to 768), or any tandem connection capabilities. Also, this section merely illustrates the functions, and is not intended to be a description of a required implementation.

The content of a terminating STS-1 SPE is formed from a payload mapping (e.g., a DS3 signal into an STS-1 SPE). Included within the STS-1 SPE are the nine bytes that make up the STS POH. Super-rate payloads can be mapped into an STS-Mc SPE, which also includes nine bytes of STS POH, or M virtually concatenated STS-1 SPEs (each of which includes nine bytes of STS POH). The BIP-8 error check byte is calculated over the entire STS SPE and the result is placed in the B3 byte of the following STS SPE. If the STS path termination is not equipped or provisioned, then an “Unequipped” STS SPE is indicated by inserting an all-zeros pattern.

If the contents of a nonterminating STS SPE are lost (e.g., as the result of a failure at an OC-M input to an OC-M to OC-N multiplex), then an all-ones pattern (for AIS-P) is inserted into the entire STS SPE and its associated H1 through H3 bytes.

The next stage involves frame aligning the STS-1s and STS-Ms, and the addition (or overwriting) of the LOH. The BIP-8 (B2) and payload pointer and pointer action (H1, H2, H3) bytes are present in the LOH of each STS-1 in an STS-N. Thus, Line level BIP-8 error checking and STS payload pointer processing are actually performed on each individual STS-1, irrespective of whether that STS-1 is part of an STS-M carrying an STS-Mc SPE.


The remaining bytes of the LOH for which uses have been defined appear in only one of the STS-1s in an OC-N signal (e.g., the S1 byte in the first STS-1). An all-zeros pattern is sent over the undefined or unused bytes within the other STS-1s. The STS-N is obtained by byte-interleaving the STS-1s and STS-Ms in the proper order.

The final part of the overhead, the SOH, is added (or overwritten) next. The Section BIP-8 (B1), Orderwire (E1), User Channel (F1) and DCC (D1 through D3) bytes are present only in the first STS-1 of an OC-N signal, and are all added before scrambling.

An all-zeros pattern is added in the corresponding undefined bytes in the second through Nth STS-1s. For N less than 768, the STS-N is scrambled, and then the framing (A1, A2) bytes and Section trace (J0) or growth (Z0) bytes are added (or overwritten) for each STS-1. For an STS-768, the undefined, Section trace and growth bytes in the first row of the SOH are added before scrambling, so that only the 128 framing bytes remain unscrambled. The final operation is the calculation of the Section BIP-8 over each 125 μ s frame of the scrambled STS-N signal. The result is placed into the B1 byte in the following frame.

The STS-N signal is then converted into optical pulses for transmission over the fiber.

Figure 5-4 Example of STS-1 and OC-N (N < 768) Signal Composition



5.2 Overhead Function Usage

5.2.1 Generating and Processing Overhead

Table 5-2 summarizes the criteria for generating and processing overhead bytes for SONET NEs, and references the primary sections in this document that contain the applicable criteria.³ Generating overhead refers to the creation of the overhead bits or bytes for the transmitted SONET signal, while processing refers to the interpretation and, if required, the further derivation of the data from the information in the overhead bits and bytes of the received signal. In general, the criteria apply only to the layers that are terminated by a particular NE. For example, the STE, LTE, STS PTE and VT PTE criteria would all be applicable to two ADMs with DS1 tributary interfaces, that generate and terminate an OC-N signal. However, only the STE criteria would apply to an STE regenerator that regenerates that same OC-N signal. Note however, that for some applications [e.g., intermediate path Performance Monitoring (PM), see [Section 6.2.2.9](#)] it may be necessary for an NE to process overhead in a layer that it does not terminate.

The status “R” in **Table 5-2** denotes that overhead generation or processing is required, either for all applications or for certain applications (as noted). The status “CR” denotes that the specified use of the overhead may be required in certain applications. The status “O” denotes that the specified use of the overhead is an objective in certain applications. For other applications, or where the status is unmarked, the overhead is either undefined (generation) or unused (processing). See [Section 3.2](#) for the criteria on undefined and unused overhead.

R5-7 [135] Overhead that is required to be generated shall carry valid data as this document describes. Processing for the required overhead shall adhere to the criteria contained in this document or the appropriate NE-specific GRs and TRs.

Note that the particular signals for which the various overhead bits and bytes listed in **Table 5-2** need to be processed will be affected by the type of APS architecture (if any) that is supported. For example, based on the definition of the linear APS protection switching boundaries in [Section 5.3.1](#), each incoming SONET signal in a linear APS system is separately monitored for several items that are required to be detected on a per-line basis for protection switching and Line PM purposes {e.g., Line BIP (B2 byte) errors, AIS-L [or lower-layer Loss Of Signal (LOS) and Severely Errored Framing/Loss Of Frame (SEF/LOF)] defects, RDI-L defects, REI-L indications}. In addition, the detection of certain of those items on any particular signal (e.g., on the incoming signal on the protection line) is directly reflected in the generation of REI-L and RDI-L indications in the LOH on the corresponding upstream signal (e.g., in the outgoing signal on the protection line). On the other hand, the NE may process various other SOH and LOH bytes such as the orderwire bytes (E1 and E2), the DCC bytes (D1 to D3, D4 to D12 and, in the case of an OC-768 signal, D13 to D156), the STS payload pointer bytes (H1 and H2), and the synchronization status byte (S1) in any of the following:

3. Overhead bits or bytes for which the only currently applicable criteria are those in [Section 3.2](#) are generally not shown in **Table 5-2** (e.g., Z1, Z2, G1 bit 8, Z4, V4, Z6, and Z7 bits 3, 4 and 8).

- All of the incoming signals (but in most cases only use the information from the active line or lines)
- Only the incoming signal or signals from the selected line or lines
- None of the incoming signals (e.g., the orderwire bytes if orderwire is not being used).

Similarly, in a UPSR NE operating at the STS path level, the STS payload pointer bytes and certain STS POH bytes are monitored for protection switching and possibly PM purposes. However, since that STS path is not necessarily terminated within that same NE (in contrast to the linear APS case where the Line layer is always terminated by the NE that performs the switching), the results of that monitoring are not directly reflected in the upstream STS path signals. Instead, the RDI-P and REI-P indications in the upstream STS path signals are based on defects and errors that are detected on the STS path that is eventually terminated.

In addition to processing overhead bits and bytes in nominal-rate SONET signals, it is also important for a SONET NE to be able to accept signals that are off frequency. In the following requirement, to “be capable of receiving and processing” a signal is intended to mean that the NE meets the applicable physical criteria (e.g., receiver sensitivity, jitter tolerance), and can perform the SONET-related functions that it normally performs (e.g., regenerate the signal, process the SONET overhead for the layers that it terminates, process the pointers for various paths, terminate the appropriate paths). It is not intended to mean that non-SONET payload signals dropped by the NE are required to be error-free when the incoming SONET signal is greater than 4.6 ppm off frequency, or that the signal must be accepted as a timing reference signal. Separate criteria (e.g., **O5-263 [349]**, and the GR-1244-CORE, *Clocks for the Synchronized Network: Common Generic Criteria*, criteria referenced in [Section 5.4.6.1](#) of this document) address those issues. Finally, the requirement is applicable when the NE’s internal clock is operating anywhere within its own free-run accuracy specifications. For example, an NE with a SONET Minimum Clock (SMC) that is free-running at -20 ppm must be able to receive and process an incoming SONET signal with a bit rate that is $+20$ ppm off frequency.

R5-8 [1013] A SONET NE shall be capable of receiving and processing incoming SONET signals with bit rates that are, at a minimum, anywhere in the range of ± 20 ppm off frequency from the nominal bit rates for those signals.

Note that based on the criteria in [Section 5.4.4.1](#), SONET NEs that support OC-192 or OC-768 Line terminating functionality are required to provide stratum 3 (or better) clocks, and therefore a traffic-carrying OC-192 or OC-768 signal should always be within ± 4.6 ppm of its nominal bit rate. On the other hand, AIS-L signals generated by OC-192 or OC-768 regenerators could be as much as 20 ppm off frequency, and it is important that downstream NEs be capable of properly receiving and processing those signals.

Table 5-2 SONET Overhead Generating and Processing Criteria

	OH Bytes	Function or Definition	Generation Criteria		Processing Criteria	
			Status	Section	Status	Section
S T E	A1, A2	Framing ¹	R	3.3.2.1	R ²	5.5
	J0	Section Trace/Fixed Pattern	O ³ (Trace)/ R (Fixed)	3.3.2.1	CR	6.2.3.2.3
	Z0	Section Growth (Fixed Pattern)	R	3.3.2.1		
	B1	Section BIP-8	R ³	3.3.2.1	R ^{2,3,4}	6.2.2.3
	E1	Local Orderwire	CR ^{3,5}	5.2.2	CR ^{3,5}	5.2.2
	F1	Section User Channel	CR ^{3,5}	3.3.2.1	CR ^{3,5}	5.2.3
	D1, D2, D3	Section DCC	R ^{3,5,6}	8.3.1.3	R ^{3,5,6}	8.3.1.3
L T E	H1, H2, H3	STS Payload Pointer and Action	R ⁷	3.5.1	R ⁷	3.5.1
	B2	Line BIP-8	R	3.3.2.3	R	6.2.2.4
	K1, K2	APS Channel	R ⁸	5.3.5	R ⁸	5.3.5
	K2 bits 6-8	RDI-L AIS-L Monitoring	R	6.2.1.3.1	R R	6.2.1.3.1 6.2.1.2.1
	D4-D12	Line DCC	CR ^{3,5,6}	8.3.1.3	CR ^{3,5,6}	8.3.1.3
	D13-D156	Extended Line DCC (OC-768 only)	CR ^{3,5,6}	8.3.1.3	CR ^{3,5,6}	8.3.1.3
	S1	Synchronization Status	R ⁵	5.4.2	R ⁵	5.4.2
	M0/M1	REI-L	R	3.3.2.3	R	6.2.2.4
	E2	Express Orderwire	CR ^{3,5}	5.2.2	CR ^{3,5}	5.2.2
S T S P T E	H1, H2, H3	STS Payload Pointer and Action AIS-P Monitoring (H1 & H2 only)	R ⁷	3.5.1	R ⁷ R	3.5.1 6.2.1.2.2
	J1	STS Path Trace and TIM-P	R	6.2.1.1.9 6.2.3.2.3	R	6.2.1.1.9 6.2.3.2.3
	B3	STS Path BIP-8	R	3.3.2.4	R	6.2.2.5
	C2	STS Path Signal Label PDI-P Monitoring	R	3.3.2.4	R CR	6.2.1.1.8 6.2.1.4.1
	G1 bits 1-4	REI-P	R	3.3.2.4	R	6.2.2.5
	G1 bit 5	One-bit RDI-P	R	6.2.1.3.2	R	6.2.1.3.2
	G1 bits 5-7	STS Path Enhanced RDI (ERDI-P)	O	6.2.1.3.2	O	6.2.1.3.2
	F2	Path User Channel ⁹	CR	3.3.2.4	CR	5.2.3
	H4	Indicator	R ¹⁰	3.4 5.2.5	R ¹⁰	3.4 5.2.5
	Z3	Growth ⁹				
	N1	STS Tandem Connection and Path Data		3.3.2.4		
	V1, V2, V3	VT Payload Pointer and Action	R ¹¹	3.5.2	R ¹¹	3.5.2

Table 5-2 SONET Overhead Generating and Processing Criteria (Continued)

	OH Bytes	Function or Definition	Generation Criteria		Processing Criteria	
			Status	Section	Status	Section
V T	V1, V2, V3	VT Payload Pointer and Action AIS-V Monitoring (V1 and V2 only)	R ¹¹	3.5.2	R ¹¹ R	3.5.2 6.2.1.2.3
	V5 bits 1, 2	VT Path BIP-2	R	3.3.3	R	6.2.2.6
	V5 bit 3	REI-V	R	3.3.3	R	6.2.2.6
P T E	V5 bit 4	RFI-V	R ¹²	6.2.1.3.3	R ¹²	6.2.1.3.3
	V5 bits 5-7	VT Path Signal Label	R	3.3.3	R	6.2.1.1.8
	V5 bit 8	RDI-V	R	6.2.1.3.3	R	6.2.1.3.3
	J2	VT Path Trace and TIM-V	CR	6.2.1.1.9 6.2.3.2.3	CR	6.2.1.1.9 6.2.3.2.3
	Z7 bit 1	Extended VT Signal Label	R ¹³	3.3.3	R ¹³	6.2.1.1.8
	Z7 bit 2	VT Virtual Concatenation	R ¹⁴	5.2.5	R ¹⁴	5.2.5
	Z7 bits 5-7	VT Path Enhanced RDI (ERDI-V)	O	6.2.1.3.3	O	6.2.1.3.3

Notes:

1. The criteria related to the generation of the undefined bytes in certain of the A1 and A2 byte positions in an STS-768 also appear in [Section 3.3.2.1](#). Processing of those bytes is not required.
2. Also required for physical layer regenerators (see TR-NWT-000917).
3. Not required, or not applicable, for drop-side signals (see [Section 3.3.2](#) for the definitions of line-side and drop-side signals).
4. Required for regenerators. Conditionally required for the NEs with both STE and LTE functionality in applications with regenerators. Not required in applications with no regenerators.
5. Not required, or not applicable, for lines 2 through n of a 1:n protected system.
6. Not required, or not applicable, for B-ISDN User Network Interface (UNI) applications.
7. Not required for nonterminated STSs that are dropped from one SONET signal to a lower-speed SONET signal, if all of the STSs in the lower-speed signal are from the same high-speed signal and the lower-speed signal uses through-timing such that there are no changes in the offsets between the pointers and the first bytes of the STS SPEs (i.e., the incoming pointers are passed through unchanged from the high-speed signal to the lower-speed signal). STS payload pointers for other nonterminated STSs may be processed at the LTE where the incoming SONET signal is terminated, at the LTE where the outgoing SONET signal is generated, or both. STS payload pointers for terminated STSs may be processed at the LTE, the STS PTE, or both.
8. Required only for the protection line.
9. The F2 and Z3 bytes are used for payload-specific functions in the DQDB mapping. See [Section 3.4.2.1.4](#).

10. Currently required only for STS virtual concatenation, VT-structured STS-1 SPEs and the DQDB mapping.
11. VT payload pointers for nonterminated VTs may be processed at the STS PTE where the incoming SONET signal is terminated, at the STS PTE where the outgoing SONET signal is generated, or both. The VT payload pointers for terminated VTs may be processed at the STS PTE, the VT PTE, or both. [Note that the frequency justification method previously discussed in TR-NWT-000496 (which has since been replaced by GR-496-CORE), of using the STS pointer to justify through VTs when an STS-1 is terminated but only some of the VTs are dropped, is not recommended for most applications. That method can result in unnecessary “hits” on added VTs during a recovery from an incoming signal failure.]
12. Required only for the byte-synchronous DS1 mapping.
13. Required only if a mapping that utilizes an extended VT signal label is supported.
14. Required only if a mapping that utilizes VT virtual concatenation is supported.

5.2.2 Orderwire

Craftspeople use the orderwire channels for voice communications between different sites during failures in the network, or during coordinated maintenance activities such as provisioning and testing of new services. Two channels are allocated for orderwire in the SONET overhead. These channels are:

- Local Orderwire (LOW), which is carried in the E1 byte of the SOH
- Express Orderwire (EOW), which is carried in the E2 byte of the LOH.

Both of these channels have bit rates of 64 kb/s and are located in the overhead of the first STS-1 in an OC-N signal. This section contains the orderwire criteria applicable to a SONET NE.

CR5-9 [136] SONET NEs with line-side interfaces may be required to provide orderwire functionality.

Support of orderwire at drop-side interfaces is not required or expected.

The primary distinction between the LOW and the EOW is that the LOW can be accessed at STE (i.e., it can be generated and processed at STE regenerators, and at all NEs that also have LTE functionality), while the EOW can be accessed only at LTE. Therefore, for systems with no regenerators, support of both the LOW and the EOW is not necessary. To promote orderwire interoperability between NEs that only support one of the orderwire channels, the following objective is applicable:

O5-10 [137] If only a single orderwire channel is supported, it should be the LOW.

If an NE does not support either the LOW or the EOW, then either the E1 byte or the E2 byte (respectively) is considered undefined, and the criteria in [Section 3.2](#) are applicable.⁴

5.2.2.1 Orderwire Access

The following orderwire access criteria apply if orderwire functionality is supported.

O5-11 [138] Access to the orderwire circuit should be through a 4-wire analog interface at 0 dBm. The input impedance should be $600\ \Omega$ ($\pm 5\%$), and the speech encoding should be μ -law PCM.

R5-12 [139] If a 4-wire analog interface is not provided, either a 2-wire analog interface [0 dBm, $900\ \Omega$ ($\pm 5\%$), μ -law PCM], or a digital interface shall be provided.

Whether the digital encoding is performed by the SONET NE (as is the case for 4-wire or 2-wire analog interfaces), or by external equipment (digital interfaces), the 8-bit PCM sample must be placed in the OC-N signal as follows.

R5-13 [140] The 8-bit PCM sample shall be synchronized to the STS-N frame, and the PCM bits shall be assigned to the corresponding bits of the appropriate E1 or E2 byte (see [Figure 3-2](#)).

The following requirement is applicable for the orderwire channel or channels that are supported.

R5-14 [141] The NE shall have the capability to generate the “quiet” PCM code (i.e., ‘01111111’) on its supported orderwire channels.

5.2.2.2 Orderwire System Communication

The LOW is intended to provide voice communications between any two NEs with STE functionality along a SONET line (including the end NEs that have both STE and LTE functionality). This means that all STE must be able to access the E1 bytes. It also means that the NEs that provide only STE functionality (e.g., STE regenerators) must be able to pass⁵ the LOW channel, to allow voice communications between non-consecutive sites on the line.

-
4. It is also acceptable for the NE to insert the “quiet” Pulse Code Modulation (PCM) code (see [Section 5.2.2.1](#)) on the unsupported E1 or E2 bytes.
 5. Typically when the word “pass” is used with respect to a digital signal, it is intended to mean that the incoming and outgoing bit-streams in the same direction of transmission are identical. However, when “pass” is used in this section, it would also be acceptable for the NE to perform back-to-back digital-to-analog and analog-to-digital conversions, and to provide circuitry that would allow multi-point communications. For example, a regenerator could provide party-line access in which the local input is added (in the analog domain) to the “passed” orderwire channel in each direction, and the received orderwire channels from both directions are added to become the local output.

-
- R5-15 [142]** An NE with STE functionality but no LTE functionality shall provide the capability to pass the incoming LOW channel through to the outgoing LOW channel on the same line.
-

The EOW is intended to provide voice communications between the LTE on a SONET line without any dependency on the LOW functionality of intermediate STE. This means the LTE at each end of the SONET line must be able to access the E2 bytes to allow end-to-end communication. In addition, in some applications it may be useful for an NE that terminates multiple SONET optical lines to provide the capability to interconnect the EOW channels on different lines.

-
- CR5-16 [144]** A SONET NE that terminates multiple SONET optical lines may be required to provide the capability to pass the EOW circuit between any two of those lines.
-

In some applications, it may be useful for the orderwire channel to be protected.

-
- CR5-17 [145]** A SONET NE may be required to support orderwire channel protection.

- R5-18 [146]** If orderwire channel protection is supported, then the orderwire protection scheme shall be the overhead protection scheme in which overhead channels are protected along with the traffic (see [Section 8.3.1.3](#)).
-

Note that using this protection scheme, in a configuration where two NEs supporting linear APS are connected via diversely-routed working and protection lines that include one or more regenerators, communications between an end NE and a regenerator site may not be possible. Specifically, if the end NE is selecting traffic from the working line then communications with a regenerator site that supports only the protection line will not be possible, while if the end NE is selecting traffic from the protection line then communications with a site that supports only the working line will not be possible.

5.2.2.3 Orderwire Operations

Monitoring and control capabilities (e.g., selective signaling, ringing capabilities, call-out functions, audible and visual call indications, orderwire channel selection) to assist users in establishing orderwire connections are FFS.

5.2.3 User Channels

As discussed in [Section 3.3.2](#), overhead bytes are allocated in the SONET signal for use by the network provider. The following criteria apply to the use of those user channels.

-
- CR5-19 [147]** An NE with STE functionality may be required to allow the user to access the Section user channel (i.e., the F1 byte) in line-side signals.
-

- CR5-20 [148]** An NE that provides only STE functionality (e.g., an STE regenerator) may be required to be capable of passing the incoming F1 byte through to the outgoing signal on the same line.

Support of the Section user channel is not required for drop-side signals.

- CR5-21 [149]** An NE with STS PTE functionality may be required to allow the user to access the path user channel (i.e., the F2 byte).
-

5.2.4 Remote Error Indications

Prior to Issue 4 of this document, there were no explicit criteria related to when and how the REI-L, REI-P and REI-V codes are supposed to be generated. Instead, it was simply assumed that LTE, STS PTE and VT PTE would be designed to insert those codes on a frame-by-frame or superframe-by-superframe basis at the next available opportunity. [For example, if two Line BIP errors were detected in one frame of an incoming OC-N signal, the M1 byte in the next frame of the upstream (outgoing) OC-N signal would be set to '0000010'.] That assumption has turned out to be incorrect. That is, some products have exhibited significant delays between the detection of a BIP error and the indication of that error in the upstream signal, and some of those same products have accumulated the errors detected over a number of frames of the incoming signal before inserting a single REI-L or REI-P code corresponding to the total number of detected errors. In addition, the assumption did not account for the possibility of a timing offset between the incoming and upstream frame or VT superframe rates (in which case the number of frames or VT superframes of BIP error results that need to be indicated over some period of time will be different than the number of REI code generation opportunities available during that period). To address both of these issues, the following criteria were added as of Issue 4.

-
- R5-22 [1160]** With the exception of the cases addressed in **O5-23 [1161]** and **R5-25 [1163]**, all SONET REI codes shall be generated on a frame-by-frame (REI-L and REI-P) or superframe-by-superframe (REI-V) basis. In addition, the maximum delay between the detection of BIP errors in a frame or VT superframe of the incoming signal and the insertion of the appropriate REI code in the corresponding outgoing signal shall be 1.0 ms at the Line and STS path levels, or 4.0 ms at the VT path level.
- O5-23 [1161]** If one frame or VT superframe of BIP error results have not yet been reflected in the appropriate REI channel when it is time to indicate the BIP error results for the next frame or VT superframe, the results of those two frames or VT superframes should be added.

- R5-24 [1162]** If **O5-23 [1161]** is met and the sum is less than or equal to the maximum value that can (according to **R3-30 [21v2]**, **R3-31 [22v2]**, **R3-38 [27v2]** or **R3-40 [29v2]**) be transmitted and correctly interpreted by the receiving NE, then that sum shall be indicated. On the other hand, if the sum is greater than that maximum value, then the maximum value shall be indicated.
- R5-25 [1163]** If no “new” BIP error results (i.e., BIP error results that have not yet been indicated in the appropriate REI channel) are available when it is time to indicate such results in an REI channel, the code for “zero” errors shall be inserted.
- R5-26 [1164]** The BIP error detection and REI generation processes shall be such that the levels of jitter and wander that an NE is required to tolerate (based on the criteria in [Sections 5.4.4](#) and [5.6.2.2.2](#)) shall not cause multiple cases where the results of two or zero frames or VT superframes of BIP error results need or are available to be indicated.
- O5-27 [1165]** The maximum delay between the detection of BIP errors in a frame or VT superframe of the incoming signal and the insertion of the appropriate REI code in the corresponding outgoing signal should be 139 μ s at the Line and STS path levels, or 514 μ s at the VT path level.

Note that the maximum delays listed in **R5-22 [1160]** are consistent with the specifications provided in ITU-T Rec. G.783, *Characteristics of synchronous digital hierarchy (SDH) equipment functional blocks*, while those listed in **O5-27 [1165]** are basically consistent with the “next available opportunity” assumption (i.e., one frame or VT superframe plus a processing delay corresponding to approximately one row of a SONET frame).

5.2.5 Virtual Concatenation

As indicated in [Sections 3.2.3.2](#) and [3.2.5](#), the H4 byte and bit 2 of the Z7 byte are used for multiframe alignment and sequence indication purposes by PTE supporting STS and VT virtual concatenation, respectively. This section contains criteria related to those functions, as well as to capacity adjustment.

5.2.5.1 Capacity Adjustment

In some applications, PTE supporting the virtual concatenation function may need to be able to utilize different numbers of STS or VT SPEs. This may be necessary (for example) as a result of changes in the needed bandwidth or failures affecting one or more of the constituent STS or VT path signals. While this capability can be controlled by the user via provisioning commands to “permanently” change the value of *X*, it can also be provided dynamically via LCAS. That function utilizes control messages that are also carried in the H4 byte or bit 2 of the Z7 byte, and is specified in ITU-T Rec. G.7042, *Link capacity adjustment scheme (LCAS) for virtual concatenated signals*.

- CR5-28 [1166]** STS or VT PTE that supports virtual concatenation and a corresponding payload mapping that does not require a fixed amount of bandwidth may be required to be capable of operating using different values of X .
- R5-29 [1167]** If STS or VT PTE that supports virtual concatenation is capable of operating using different values of X , it shall also support the capability for the user to provision the particular value of X to be used.
- CR5-30 [1168]** STS or VT PTE that supports virtual concatenation and a corresponding payload mapping that does not require a fixed amount of bandwidth may be required to also support the capability to dynamically adjust the number of SPEs that it is using.
- R5-31 [1169]** If STS or VT PTE that supports virtual concatenation also supports the capability to dynamically adjust the number of SPEs that it is using, that capability shall be provided using LCAS and according to the specifications in ITU-T Rec. G.7042.
- R5-32 [1170]** LCAS-enabled STS or VT PTE shall detect if it is connected to non-LCAS equipment, and shall automatically adjust its receiver's mode of operation to interoperate with that equipment (as described in ANSI T1.105 and ITU-T Rec. G.7042). In addition, the LCAS-enabled PTE shall continue to transmit as if it were connected to other LCAS-enabled (rather than non-LCAS) equipment.

Note that LCAS-enabled PTE would never be expected to transmit all zeros in the LCAS Control (CTRL) word and CRC bits, while that is the code that is required to be transmitted by non-LCAS equipment (see **R5-35 [1173]**). Thus, the reception of that code at LCAS-enabled PTE can be used by that equipment to determine that it is connected to non-LCAS equipment and needs to ignore the codes received in the remaining LCAS-specific fields. (See [Section 5.2.5.2](#), ANSI T1.105 and ITU-T Rec. G.7042 for more information related to the CTRL bits and other LCAS-specific fields.)

5.2.5.2 Overhead Bit and Byte Usage

In the case of STS virtual concatenation, the H4 byte in each of the constituent STS-1 or STS-3c SPEs consists of four bits that are used for several different functions (bits 1 through 4) and four bits that are used to carry the 1st Multiframe Indication (MFI1, carried in bits 5 through 8). MFI1 is a 4-bit counter that is incremented every frame (counting from 0 to 15), and is used for SPE-alignment purposes (see below) and to indicate the particular information that appears in bits 1 through 4 of the same byte. This latter information includes the 2nd Multiframe Indication (MFI2), which is an 8-bit counter that is incremented in each 16-frame MFI1 multiframe. As shown in [Table 5-3](#), MFI2 is split into two 4-bit nibbles that are transmitted in the first and second frames of the MFI1 multiframe (i.e., in MFI1 frames 0 and 1). Taken together, MFI1 and MFI2 provide a two-stage, 12-bit counter that has $2^{12} = 4096$ values, and therefore has a cycle time of 512 ms (i.e., $4096 \times 125 \mu\text{s}$). By aligning the

counters carried in each of the constituent STS SPEs at the transmitter, differential delays of up to ± 256 ms can be detected and (depending on the buffering capabilities of the receiving PTE) possibly accommodated at the receiver.

In cases where LCAS is not supported, the only other information carried in bits 1 through 4 of the H4 byte is the STS-1 or STS-3c SPE's Sequence indicator (SQ). This is an 8-bit value that is used to uniquely identify the location of each SPE within the STS-1- X_v or STS-3c- X_v SPE [from 0 for the first SPE to $(X - 1)$ for the last SPE]. Similar to MFI2, the SQ value is split into two 4-bit nibbles and those nibbles are transmitted in MFI1 frames 14 and 15. In addition, the 8 bits allocated for this function mean that the maximum possible value of X is 256 (i.e., 2^8).

In cases where LCAS is supported, bits 1 through 4 of the H4 byte are also used to carry CTRL words, a Group Identification (GID), member status information, a Re-Sequence Acknowledgement (RS-ACK), and a CRC-8. Although these will not be discussed further in this document (see **R5-31 [1169]**), they are shown in [Table 5-3](#).

R5-33 [1171] STS PTE that supports virtual concatenation shall utilize the H4 byte carried in each of the constituent STS-1 or STS-3c SPEs as described above and shown in [Table 5-3](#).

Table 5-3 H4 Byte Usage in STS Virtual Concatenation Applications

Use of Bits 1 to 4	H4 Bits								MFI1 Frame Number
	Bit 1	Bit 2	Bit 3	Bit 4	Bit 5	Bit 6	Bit 7	Bit 8	
MF12 MSB ^a	b1	b2	b3	b4	0	0	0	0	0
MF12 LSB ^a	b5	b6	b7	b8	0	0	0	1	1
CTRL ^{b,c}	x	x	x	x	0	0	1	0	2
GID ^{b,c,d}	0	0	0	x	0	0	1	1	3
Reserved	0	0	0	0	0	1	0	0	4
Reserved	0	0	0	0	0	1	0	1	5
CRC-8 MSB ^{b,e}	C ₁	C ₂	C ₃	C ₄	0	1	1	0	6
CRC-8 LSB ^{b,e}	C ₅	C ₆	C ₇	C ₈	0	1	1	1	7
Member Status ^{b,c}	x	x	x	x	1	0	0	0	8
Member Status ^{b,c}	x	x	x	x	1	0	0	1	9
RS-ACK ^{b,c}	0	0	0	x	1	0	1	0	10
Reserved	0	0	0	0	1	0	1	1	11
Reserved	0	0	0	0	1	1	0	0	12
Reserved	0	0	0	0	1	1	0	1	13
SQ MSB ^a	b1	b2	b3	b4	1	1	1	0	14
SQ LSB ^a	b5	b6	b7	b8	1	1	1	1	15

Notes:

- a b1 is the most significant bit (and depending on the particular value being transmitted may be either '0' or '1'), b2 is the second most significant bit, ..., and b8 is the least significant bit.
- b This entry is applicable only in applications where LCAS is supported. In other applications these bits are Reserved and set to all zeros.
- c Depending on the situation, "x" may be either '0' or '1'.
- d The GID bit (bit 4) is cycled through a 2¹⁵-1 Pseudo Random Bit Sequence (PRBS) pattern, and is set identically in each of the STS-1-Xv's or STS-3c-Xv's constituent SPEs.
- e C₁ is the most significant bit of the remainder of a CRC calculation performed on the first 14 nibbles of the 16-nibble LCAS control packet (which starts with the first Member Status nibble in the preceding MFI1 multiframe and ends with these CRC-8 nibbles), C₂ is the second most significant bit, ..., and C₈ is the least significant bit. Note that the MFI1 bits are not in the LCAS control packet, and therefore are not included in the CRC calculation.

In VT virtual concatenation, bit 2 of the Z7 byte is used for most of the same functions as are supported by the H4 byte in STS virtual concatenation (i.e., multiframe alignment, SQ, LCAS-related functions). On the other hand, in the VT case only one bit (per VT) is available every 500 μs, as opposed to 8 bits (per STS) every 125 μs in the STS case. To accommodate this difference, several functions utilize fewer bits in the VT case, and the extended VT path signal label discussed in

[Section 3.3.2.4](#) and carried in Z7 bit 1 is used for the alignment and field identification functions supported by MFI1 in STS virtual concatenation.⁶ This is illustrated in [Table 5-4](#), which shows that five bits are assigned to a Frame Indicator, six bits are assigned to SQ (thus limiting the number of VTs in a VTn-Xv to $2^6 = 64$), and three bits are used to carry the results of a CRC-3 calculation (rather than a CRC-8 calculation) in applications where LCAS is supported.

Similar to MFI2 in STS virtual concatenation, the Frame Indicator in the VT virtual concatenation case is a counter that is incremented in each multiframe defined by the MFAS in Z7 bit 1. In this case the count is from 0 to 31, and this, combined with the MFAS multiframe length of 32 VT superframes and the VT superframe time of 500 μ s, results in a total cycle time of 512 ms (i.e., the same total cycle time as in the STS virtual concatenation case). Thus, by aligning the counters carried in each of the constituent VT SPEs at the transmitter, differential delays of up to ± 256 ms can be detected and (depending on the buffering capabilities of the receiving PTE) possibly accommodated at the receiver.

-
- R5-34 [1172]** VT PTE that supports virtual concatenation shall utilize bits 1 and 2 of the Z7 byte carried in each of the constituent VTn SPEs as described above and shown in [Table 5-4](#).
-

6. No mappings that have been (or will be) assigned a non-extended VT path signal label will utilize a VTn-Xv SPE. Therefore, in all applications in which Z7 bit 2 is needed to carry information specific to VT virtual concatenation, bit 1 of that same byte will be required to be carrying an extended VT path signal label.

Table 5-4 Z7 Bits 1 and 2 Usage in VT Virtual Concatenation

VT Superframe Number	Z7 Bit 1 Use	Z7 Bit 1 Value	Z7 Bit 2 Use	Z7 Bit 2 Value		
1	MFAS (used instead of a VT virtual concatenation specific MFI1)	0	Frame Indicator (equivalent to MFI2, but with a shorter counting cycle of 0 to 31)	b1		
2		1		b2		
3		1		b3		
4		1		b4		
5		1		b5		
6		1		SQ	b1	
7		1			b2	
8		1			b3	
9		1			b4	
10		1			b5	
11		0			b6	
12	Signal Label	b1	CTRL	x		
13		b2		x		
14		b3		x		
15		b4		x		
16		b5		GID	x	
17		b6			Reserved	0
18		b7				0
19		b8				0
20	–	0	0			
21	Reserved	0	RS-ACK	x		
22		0	Member Status	x		
23		0		x		
24		0		x		
25		0		x		
26		0		x		
27		0		x		
28		0		x		
29		0		x		
30		0		CRC-3	C ₁	
31		0	C ₂			
32	0	C ₃				

Notes:

- a b1 is the most significant bit (and depending on the particular value being transmitted may be either ‘0’ or ‘1’), b2 is the second most significant bit, etc.
- b This entry is applicable only in applications where LCAS is supported. In other applications these bits are Reserved and set to all zeros.
- c Depending on the situation, “x” may be either ‘0’ or ‘1’.
- d The GID bit (bit 4) is cycled through a 2¹⁵–1 PRBS pattern, and is set identically in each of the VTn-Xv’s constituent SPEs.
- e C₁ is the most significant bit of the remainder of a CRC calculation performed on the first 29 bits of the 32-bit LCAS control packet (which starts with b1 of the frame indicator and ends with these CRC-3 bits), C₂ is the second most significant bit, and C₃ is the least significant bit.

Finally, in order to facilitate interworking between LCAS-enabled and non-LCAS equipment, **O3-2 [5]** is upgraded to a requirement as follows.

-
- R5-35 [1173]** STS or VT PTE that supports virtual concatenation but does not support (or is not provisioned to utilize) LCAS shall transmit all zeros in the Reserved bits that have been defined to carry the CTRL word in LCAS applications.
-

5.2.5.3 Differential Delay Accommodation

As indicated in [Section 5.2.5.2](#), in both the STS and VT virtual concatenation cases, differential delays of up to ± 256 ms can be detected and possibly accommodated at the receiver. While the ability to accommodate such large delays could be useful in some applications, in most cases any potential benefit would be more than offset by the negative impacts that this would have in such areas as equipment costs (e.g., additional buffer capacity) and end-to-end transmission delays for any delay-sensitive payloads. Therefore, only a relatively small amount of differential delay is required to be accommodated (see **R5-36 [1174]**). Note that for the purposes of the following criteria, the “total” differential delay is defined to be the difference in the delays experienced by the constituent SPE with the largest individual delay and the SPE with the smallest individual delay. Also note that “accommodate” in this context is defined to mean that the equipment can properly process SPEs having any combination of differential delays (amongst themselves) that falls within the boundary condition of the particular specified maximum total differential delay.

-
- R5-36 [1174]** STS or VT PTE that supports virtual concatenation shall be capable of accommodating a total differential delay of 125 μ s.
-

In general, equipment that conforms to **R5-36 [1174]** can be expected to operate properly in any configuration in which all of the constituent SPEs traverse the same (or nearly the same) route through the network. However, in order to be used in configurations in which the SPEs can traverse substantially different routes, the receiving PTE would need to be able to accommodate larger delays. For example, in a 1200-km ring with ten equally spaced nodes, path signals that travel in opposite directions around the ring between two adjacent nodes would be subject to a fiber transmission delay difference of approximately 4.8 ms.

-
- O5-37 [1175]** STS or VT PTE that supports virtual concatenation should be capable of accommodating a total differential delay of Z μ s, where Z is greater than 125.

- R5-38 [1176]** The maximum total differential delay that can be accommodated by STS or VT PTE that supports virtual concatenation shall be clearly documented.
-

5.3 Automatic Protection Switching

Several APS schemes have been defined for SONET signals. This section contains criteria related to the scheme known as linear APS. Other Telcordia criteria documents [e.g., GR-1230-CORE, *SONET Bidirectional Line-Switched Ring Equipment Generic Criteria*, and GR-1400-CORE, *SONET Dual-Fed Unidirectional Path Switched Ring (UPSR) Equipment Generic Criteria*] contain criteria applicable to systems that use other protection switching schemes. In addition, Section 5.6 of GR-499-CORE contains circuit pack protection switching criteria that are applicable to SONET NEs that provide redundant or protection circuit packs.

CR5-39 [150] SONET LTE that terminates optical lines may be required to provide linear APS.

Support of linear APS at STS-N electrical interfaces is not expected or required.

R5-40 [151] If linear APS is provided, the SONET NE shall provide the capability for the user to disable the feature on a per-interface basis.

The capability to disable the linear APS feature is needed for interworking with SONET NEs that do not support it.

Linear APS, and in particular, the protocol for the APS channel, is standardized to allow interworking between SONET LTE from different suppliers. The criteria in this section apply to all SONET LTE that provides linear APS.

5.3.1 Protection Switching Boundaries

Linear APS is defined to provide protection at the line layer (see [Section 3.3](#)). Therefore, all of the STS SPEs carried in an OC-N signal are protected together (i.e., if a switch occurs, all of the STS SPEs are switched simultaneously).

R5-41 [152v2] Protection shall cover the multiplexer/optics units from a point at or before where the LOH is inserted and after the STS POH is inserted, to a point at or beyond where the LOH is terminated and before the STS POH is terminated.

5.3.2 Linear APS Architectures

Two linear APS architectures are defined in the following sections. These are the 1+1 architecture and the 1:n architecture. In addition, the special case of the 1:n architecture where n is equal to 1 is also discussed. Several criteria are given that apply only to the 1:1 case, to allow interworking with LTE using the 1+1 architecture.

5.3.2.1 1+1 Architecture

The 1+1 architecture is defined as follows:

An architecture in which the protected signals are continuously bridged (at the electrical level) to working and protection equipment so that the same STS path signals (and thus the same payloads) are transmitted identically to the selector-end working and protection equipment. At the selector end, the working and protection OC-N signals are monitored independently and identically for failures. The receiving equipment chooses either the working or the protection signal as the one from which to select the traffic, based on the switch initiation criteria contained in [Section 5.3.3](#). Because of the continuous bridge, the 1+1 architecture does not allow an unprotected extra traffic channel to be provided.

The 1+1 architecture could be supported as either a user-provisionable option (e.g., with the 1:1 case of the 1:n architecture as the other option) or as the only supported architecture. If it is the only supported architecture, it may still be upgradable to the 1:n architecture.

A system using the 1+1 architecture operates, as a default, in a unidirectional mode. In this mode, the switching is complete when a channel in the failed direction is switched to the protection line. Although end-to-end signaling is not needed for the purpose of completing and releasing switches, the APS channel (which is carried in the K1 and K2 bytes of the signal on the protection line) is still used to indicate the local switch actions and the mode of operation. A bidirectional switching mode may also be provided as a user-provisionable option. In this mode, a channel is switched to the protection line in both directions. Switching of only one direction is not allowed. End-to-end signaling is accomplished using the APS channel.

A 1+1 system also uses, as a default, nonrevertive switching. In nonrevertive switching, a switch to the protection line is maintained even after the working line has recovered from the failure that caused the switch, or the manual switch command is cleared. Revertive switching may also be provided as a user-provisionable option. In revertive switching, the traffic is switched back to the working line when the working line has recovered from the failure or the manual command is cleared.

CR5-42 [153] The LTE may be required to support the 1+1 architecture.

R5-43 [154] If the 1+1 architecture is provided, the following apply.

- The unidirectional mode shall be provided.
- If bidirectional switching is provided (see **CR5-44 [155]**), the mode shall be user provisionable, with a default mode of unidirectional.
- If bidirectional switching is provided, the switching operation shall be bidirectional only if the LTE at both ends is provisioned to operate in the bidirectional mode. Otherwise, the LTE shall operate in the unidirectional mode. (The K2 byte is used to determine the mode of the far-end LTE as described in [Section 5.3.5.2](#).)
- Nonrevertive switching shall be provided.

- If revertive switching is provided (see **CR5-45 [156]**), the choice of nonrevertive or revertive shall be user provisionable, with a default of nonrevertive.

CR5-44 [155] If the 1+1 architecture is provided, the bidirectional mode may be required to be provided.

CR5-45 [156] If the 1+1 architecture is provided, revertive switching may be required to be provided.

CR5-46 [157] 1+1 LTE may be required to be upgradable to the 1:n architecture.

5.3.2.2 1:n Architecture

The 1:n architecture is defined as follows:

An architecture in which any of the n working channels can be bridged to a single protection line. Permissible values of n are from 1 to 14. End-to-end signaling is accomplished by using the APS channel. Because the bridge end is switchable, the protection line can be used to carry an extra traffic channel.

The 1:n architecture could be supported as either a user-provisionable option (i.e., with the 1+1 architecture as the other user-provisionable option) or as the only supported architecture (although see [Section 5.3.2.3](#) for criteria related to the 1:1 case).

In a 1:n system, all switching is revertive, and both unidirectional and bidirectional switching modes are provided. The mode of operation is user provisionable, with bidirectional as the default.

CR5-47 [158] LTE may be required to support the 1:n architecture.

CR5-48 [159] LTE supporting the 1:n architecture may be required to support values of n greater than 1.

R5-49 [160] If the 1:n architecture is provided, the following apply.

- All switching shall be revertive.
- The unidirectional mode shall be provided.
- The bidirectional mode shall be provided.
- The mode shall be user provisionable, with a default mode of bidirectional.
- The switching operation shall be unidirectional only if the LTE at both ends is provisioned to operate in the unidirectional mode. Otherwise, the LTE shall operate in the bidirectional mode. (The K2 byte is used to determine the provisioned mode of the far-end LTE as described in [Section 5.3.5.2](#).)

O5-50 [161] If the 1:n architecture is provided, the LTE should provide the capability to transport extra traffic on the protection line when it is not being used for protection.

- R5-51 [162]** If the capability to transport extra traffic is provided, the SONET NE shall provide the capability for the user to disable that feature for interworking with other SONET NEs that do not support it.
-

5.3.2.3 1:1 Case of the 1:n Architecture

The 1:1 case is a subset of the 1:n architecture, with n equal to 1. In addition to the criteria for the 1:n architecture, the following criteria are specific to the 1:1 case.

- CR5-52 [163]** LTE that supports only the 1:1 case of the 1:n architecture may be required to be upgradable to support values of n greater than 1.
-

The following requirement allows 1:1 LTE to interwork with LTE using the 1+1 architecture.

- R5-53 [164]** 1:1 LTE (which indicates the 1:n architecture on bit 5 of the transmitted K2 byte) shall operate as 1+1 LTE if the far-end LTE indicates that it is 1+1 LTE, as detected on the received K2 byte. The 1:1 LTE shall continue to indicate the 1:n architecture on the transmitted K2 byte unless it is reprovisioned by the user to 1+1. It shall also continue to indicate its provisioned unidirectional/bidirectional switching mode; however, it shall meet the criteria in [Section 5.3.2.1](#) (instead of the criteria in [Section 5.3.2.2](#)) concerning which of those modes to actually operate in.
-

Although nonrevertive switching is the default for the 1+1 architecture, it is not essential to the operation of a system for both ends to use nonrevertive switching (or for both to use revertive), and therefore the type of switching used locally is not indicated to the far end. Therefore, 1:1 LTE can use either revertive or nonrevertive switching when it is operating as 1+1 LTE.

There are currently no criteria concerning the time at which 1:1 LTE must change its operation to 1+1. The time at which it declares an APS Mode Mismatch failure (see [Section 6.2.1.1.6.C](#)) would be an appropriate time.

5.3.3 Switch Initiation and Completion Criteria

5.3.3.1 Switch Initiation Criteria

The following two automatic switch initiation criteria are defined for linear APS.

1. Signal Fail (SF): A “hard failure” condition detected on the incoming OC-N signal.

- R5-54 [165]** Loss of Signal, Loss of Frame and AIS-L defects (see [Section 6.2.1](#)), and a Line BER exceeding 10^{-3} on an incoming OC-N shall be detected as SF conditions on that line.
-

Other protectable hard failures (e.g., a stuck bit) may also be detected as SF conditions.

CR5-55 [166] The BER threshold for an SF condition may be required to be user provisionable⁷ over the range of 10^{-3} to 10^{-5} .

2. Signal Degrade (SD): A “soft failure” condition resulting from the Line BER exceeding a pre-selected threshold.

R5-56 [167] A BER exceeding the SD threshold on an incoming OC-N shall be detected as an SD condition on that line.

R5-57 [168] The BER threshold for an SD condition shall be user provisionable over the range of 10^{-5} to 10^{-9} .

5.3.3.2 Switch Initiation Time

The switch initiation time is the time that it takes LTE to detect an SF or SD condition and initiate a switch (if appropriate).

R5-58 [169] For SF conditions caused by LOS, LOF or AIS-L defects, the switch initiation time shall be 10 ms or less.

O5-59 [170] For SF conditions caused by LOS, LOF or AIS-L defects, the switch initiation time should be 8 ms or less.

SF and SD conditions based on the BER criteria are detected when the BER of a line exceeds a threshold. That BER is derived from the sum of Line BIP-8 violation counts of the individual STS-1s in the STS-N, and the time allowed to detect that the BER is greater than or equal to a threshold (and to initiate a switch, if appropriate) depends on the actual BER, not the threshold.⁸

In general, the occurrence of errors on an optical signal is expected to be a random process. Therefore, for any particular actual BER, there will be some probability that the number of errors detected in a particular time period will be less than the threshold for detecting an SD or SF condition, and a corresponding probability that the number of errors will be greater than or equal to the threshold. The intent of the

7. For both SF and SD thresholds, providing thresholds at 1×10^{-n} , where n is an integer, is sufficient.

8. On the other hand, the time that an NE actually takes (as opposed to being allowed to take) may depend on both the actual BER and the threshold. In addition, it is possible (depending on the algorithms that are used) that decreasing the threshold will reduce the actual detection time. For example, the time needed to determine that the BER of a signal with an actual BER of 2×10^{-7} is worse than a threshold of 1×10^{-8} could easily be less than the time needed to determine that the BER of that same signal is worse than a threshold of 1×10^{-7} .

following criteria (which were written assuming a Poisson distribution of errors) is to promote the development of BER detection algorithms that meet the goals listed below:

- When the actual BER is greater than or equal to an SD or SF threshold, it is quickly detected as such, and a switch (if appropriate) is initiated
- The probability of detecting a threshold crossing when the actual BER is below the threshold is very small and tolerant to burst errors.

R5-60 [171] For SF and SD conditions based on BER, the switch initiation time shall be below the “maximum” curve in [Figure 5-5](#) (assuming the actual BER is greater than or equal to the threshold).⁹

O5-61 [172] For SF and SD conditions based on BER, the probability that the switch initiation time will be less than the “objective” curve in [Figure 5-5](#) (for the particular rate OC-N signal) should be greater than 0.95 (assuming the actual BER is greater than or equal to the threshold).

Whereas the “maximum” switch initiation time curve shown in [Figure 5-5](#) applies to all OC-N rates, the objective switch initiation time curves depend on the level N. The curves in [Figure 5-5](#) take into account the error detection saturation effect at high BERs.


R5-62 [173v2] For an SF or SD detection threshold of 10^{-n} and an actual BER of $1 \times 10^{-(n+1)}$ or less, the probability that the SD or SF condition will be detected in the “maximum” switch initiation time from [Figure 5-5](#) (for that particular threshold) shall be less than or equal to 10^{-6} .

For example, if the SD threshold is 10^{-5} and the actual BER is 1×10^{-6} , the probability that an SD condition will be detected during any 1-second interval must be less than or equal to 10^{-6} .

In an analysis situation, it is likely that the BER detection algorithm details needed to calculate LTE’s conformance to these criteria (using the assumption of a Poisson distribution of errors) will not be available. Therefore, tests will need to be performed where errors are inserted onto an actual signal. Several methods of generating errors are possible, including attenuation of the optical signal and insertion of errors using a SONET test set. If errors are inserted using a SONET test set, their distribution may be either periodic or “random” (e.g., generated using some type of random number generation process).

9. For a random distribution of errors, it is impossible to specify a true “maximum” switch initiation time (i.e., one where the probability of detecting a threshold crossing and initiating a switch within the specified time is 1.0). However, the probability that the switch initiation time will be within the “maximum” time from [Figure 5-5](#) must be very close to 1.0.

Figure 5-5 Switch Initiation Time Criteria



5.3.3.3 Switch Completion Time

R5-63 [174] The time to complete a switch, once it is initiated, shall be 50 ms or less.

R5-64 [175] For bidirectional switching, the LTE at both ends shall complete the switch within the same 50-ms switch completion time, from the time the request is initiated.

For LTE operating in the 1:n architecture or the 1+1 bidirectional mode, the LTE at both ends of the line must perform various actions before a switch can be completed. No criteria exist concerning the times for the individual actions. Therefore, in configurations where the LTE is made by different suppliers, the overall switch completion time may be greater than 50 ms. The need for additional criteria in this area is FFS.

In addition, note that **R5-63 [174]** and **R5-64 [175]** are the applicable requirements for automatically initiated switches caused by degradations or failures, and also for manually initiated switches (e.g., switches initiated by the user from the craft interface). This means that traffic could be interrupted for as long as 50 ms during a manually initiated switch. Some users may find an interruption of that length to be unacceptable, and therefore the following conditional requirement applies for manually initiated switches.

CR5-65 [913] Manually initiated facility protection switches may be required to be error-free.

Note that in many applications, differences in the transmission delays for the working and protection lines may preclude error-free switching. However, it should still be possible for the traffic interruptions caused by a manually initiated switch to be much shorter than the 50 ms switch completion time requirement (i.e., the switch initiation time criteria in [Section 5.3.3.2](#) are not applicable for manually initiated switches, so it should be possible for the NE to take additional time to prepare for the switch, minimizing the time it takes to complete the switch after it is initiated).¹⁰

Finally, it should be noted that the switch initiation time criteria in [Section 5.3.3.2](#) and the switch completion time requirement in this section are separately applicable. They are not intended to be combined into (for example) a 60 ms total switching time requirement for hard failures. Therefore, if a particular system (e.g., a system that uses a unidirectional switching architecture and semiconductor switching mechanisms) is capable of completing switches much faster than 50 ms, that system is still required to meet the switch initiation time requirements in [Section 5.3.3.2](#) and therefore will have a total switching time for hard failures that is much shorter than 60 ms.

5.3.4 Restoral and Clearing of SD and SF Conditions

For LTE using either revertive or nonrevertive switching, a hysteresis method of clearing SD and SF conditions based on the BER is used. This method specifies an SD or SF clearing threshold ten times lower than the SD or SF threshold.

R5-66 [176] The clearing threshold for an SD or SF condition based on the BER shall be one-tenth the threshold for declaring the SD or SF.

10. There currently are no criteria in this document concerning switch initiation times for manually initiated switches. However, other Telcordia documents (e.g., TR-TSY-000824, *OTGR Section 10.1: User System Interface – User System Access*) indicate that an NE must respond to user input within approximately 2 seconds. If a user enters one of the switch commands discussed in [Section 5.3.6.1](#) and that switch is not successfully completed within 2 seconds, it is assumed that the NE will indicate to the user that the command could not be completed. Therefore, the switch initiation time would need to be somewhat less than 2 seconds. The need for explicit criteria in this area is under study.

For 1:n systems, the process of determining that the working line is operating at a BER lower than the clearing threshold needs to be reasonably fast so that the system can restore and free up the protection line for use by other working channels (or the extra traffic channel). In 1+1 systems, the protection line is not needed to carry other channels; however, SD and SF conditions still need to be cleared in a reasonable amount of time to allow the use of the better line when the other line also has an SD or SF condition. On the other hand, it is also important that an existing SD or SF condition not be cleared unless the BER is actually better than the corresponding clearing threshold.

- R5-67 [1046]** If an SD or SF condition has been detected and the incoming signal’s BER is greater than or equal to that SD or SF threshold, the probability that the LTE will detect that the BER is less than the SD or SF clearing threshold within the “maximum clearing” time listed in [Table 5-5](#) shall be less than or equal to 10^{-6} .
- R5-68 [1047]** If an SD or SF condition has been detected and the incoming signal’s BER is less than or equal to the SD or SF clearing threshold, the probability that the LTE will detect that the BER is less than that threshold within the “maximum clearing” time shown in [Table 5-5](#) shall be greater than or equal to 0.99.
- O5-69 [177v3]** If an SD or SF condition has been detected and the incoming signal’s BER is less than or equal to the SD or SF clearing threshold, the probability that the LTE will detect that the BER is less than that threshold within the “objective clearing” time shown in [Table 5-5](#) should be greater than 0.95.

Table 5-5 Clearing Time Criteria for BER-based SF and SD Conditions

SD or SF Threshold	Clearing Threshold	Maximum Clearing Time	Objective Clearing Time / SONET Rate ^a					
			OC-1	OC-3	OC-12	OC-48	OC-192	OC-768
10^{-3}	10^{-4}	10 ms	None	None	None	None	None	None
10^{-4}	10^{-5}	100 ms	None	None	25 ms	10 ms	10 ms	10 ms
10^{-5}	10^{-6}	1 s	None	None	250 ms	63 ms	16 ms	10 ms
10^{-6}	10^{-7}	10 s	None	None	2.5 s	625 ms	160 ms	40 ms
10^{-7}	10^{-8}	100 s	None	83 s	21 s	5.2 s	1.3 s	325 ms
10^{-8}	10^{-9}	1,000 s	None	667 s	167 s	42 s	11 s	2.6 s
10^{-9}	10^{-10}	10,000 s	None	5360 s	1330 s	340 s	85 s	21 s

Note:

- a. Where no objective time is listed for a particular SONET rate, **O5-69 [177v3]** is not applicable.

For consistency with the switch initiation time criteria, the clearing detection times shown above are based on the curves in [Figure 5-5](#). However, in contrast to the switch initiation time criteria, which are a function of the OC-N rate and the actual BER, the clearing time criteria are a function of the OC-N rate and the clearing

threshold. For example, for an OC-12 signal and an SD clearing threshold of 1×10^{-7} , the clearing detection time requirement is 10 seconds and the objective is 2.5 seconds whether the actual BER is 9×10^{-8} or 1×10^{-10} . In this example the SD threshold would be 1×10^{-6} , and the switch initiation time criteria would be 10 seconds (requirement) and 0.25 seconds (objective) for an actual BER of 1×10^{-6} and 0.1 seconds (requirement) and 0.008 seconds (objective) for an actual BER of 1×10^{-4} .

-
- R5-70 [914]** Once LTE has detected that the BER is less than the SD or SF clearing threshold, it shall clear the SD or SF condition within an additional 10.5 seconds (although see below for possible exceptions in cases of intermittent SD and SF conditions) assuming the LTE does not detect that the BER is greater than or equal to the SD or SF threshold before that condition is cleared.
-

Depending on the SONET rate and the particular SD and SF thresholds that LTE is using, **O5-69 [177v3]** could result in extremely short detection times for the purposes of clearing SD and SF conditions (e.g., for an OC-48 SD threshold of 10^{-5} , the required SD clearing threshold is 10^{-6} and the detection time objective is 62.5 ms). While it is important to clear SD and SF conditions fairly quickly (as discussed above), very short clearing times could result in rapid switches between channels in some situations. For example, oscillations could occur if an SD condition was detected on one line and an intermittent SF condition was detected and cleared repeatedly on another line. Therefore, the following criteria are applicable.

-
- CR5-71 [915]** LTE may be required to be designed to reduce the chance or number of rapid protection switching oscillations that could occur in multiple failure or degradation situations where one or more of the failures or degradations is intermittent.
- R5-72 [916]** If LTE is designed to reduce the chance or number of rapid protection switching oscillations, the method used shall be clearly documented.
-

Two methods that LTE could use to meet **CR5-71 [915]** are:

- Delay clearing SD and SF conditions for approximately 10 seconds (as allowed in **R5-70 [914]**) after the LTE has detected that the BER of the incoming signal is less than the clearing threshold. The LTE would then clear the condition only if it determines that the SD or SF threshold has not been crossed during the delay period. If the SD or SF threshold is crossed, then the process for clearing the condition is restarted.
- Monitor how often an SF or SD condition is detected and cleared on a line, and “lock on” to that condition if it is detected more than x times in y seconds. After the LTE has locked on to the condition, it would consider that line to be in an SF or SD condition until it determines that the intermittent failure or degradation is gone (e.g., it would clear the SF condition when it has detected that the BER is less than the SF clearing threshold and no hard failures have occurred for z seconds).

The first method discussed above would reduce the chance of rapid oscillations in some situations; however, in other cases it could extend the time that traffic is interrupted.¹¹ The second method would not reduce the chance of rapid oscillations; however, it would limit the number of oscillations that could occur.

If a system uses revertive switching, then frequent automatically initiated switches could also (in addition to the multiple failure situations discussed above) occur as the result of an intermittent failure or degradation on a single working line. To prevent this, a Wait To Restore (WTR) period is defined for LTE using revertive switching.¹² After the BER of the working line meets the clearing threshold and the SD or SF condition is cleared, a WTR period is allowed to elapse before the revertive switch is performed. Note that the WTR period is not used after manually initiated switches, or after an SD or SF condition on the protection line clears.

-
- R5-73 [178v2]** For LTE using revertive switching, a WTR period of 5 to 12 minutes shall be provided after the condition that caused an automatically initiated switch to the protection line clears. The length of the WTR period shall be user provisionable on a per-protection-line (or per-protection-group) basis.¹³
-

5.3.5 APS Channel Protocol

This section presents the protocol for the APS channel, which is carried in the K1 and K2 bytes in the signal on the protection line. The APS controllers at the LTE terminating SONET lines use the channel to exchange requests and acknowledgments for protection switch actions. Even if the protocol is not needed to complete the protection switch actions (as in the case of 1+1 unidirectional

-
11. For example, suppose SF conditions have been detected on both working line 1 and the protection line. If the protection line is repaired first and a delay period is used, the NE would not remove the “SF – null channel” request (see [Section 5.3.5](#)) for 10 additional seconds. Therefore, it would not be able to insert an “SF – working channel 1” request for those 10 seconds, and the traffic would remain interrupted during the delay period. (Note that if working line 1 was repaired first, the traffic would be restored as soon as the LOS defect on that line was terminated, which would not be affected by whether there is a delay period for clearing SF conditions.)
12. Note that the term “restore” is commonly used to refer to several different processes related to protection switching, including performing a protection switch in response to a failure (e.g., restore the traffic), “fixing” a failed entity and determining that the entity is once again capable of transporting traffic (e.g., restore an optical fiber or circuit pack), and switching from a protection facility or circuit pack back to the working facility or pack (e.g., restore traffic to the working line or pack). The WTR period referred to here is concerned with the delay between the last two of these processes (i.e., between the time at which it is determined that the working line is capable of satisfactorily transporting traffic and the alarm associated with the failure is cleared, and the time at which the traffic is actually switched back to that line). Such a period only applies in cases where revertive switching is used and, therefore, it could be more accurately referred to as a “Wait-To-Revert” period. However, “Wait To Restore” is the phrase that has traditionally been used, and it will continue to be used here (at least in this issue of the document).
13. It is sufficient to provide provisionable WTR times from 5 to 12 minutes in 1-minute increments.

operation), it can still be used for other purposes¹⁴ and therefore the appropriate codes must be transmitted. The bit assignments for these bytes and the bit-oriented protocol are defined in this section.

5.3.5.1 K1 Byte

This section contains the bit assignments and the APS channel generation rules for the K1 byte carried on the protection line. This byte is used to indicate a request by a channel for a switch action.

5.3.5.1.1 Bit Assignments for the K1 Byte

The K1 byte is divided into two parts. Bits 1 through 4 are used to indicate a request, while bits 5 through 8 are used to indicate which channel is making the request.

Three categories of requests can be indicated using K1 bits 1 through 4:

1. Automatically initiated requests (i.e., SF and SD), which are used to indicate that a failure or degradation condition has been detected on the line associated with the requesting channel
2. An external request (i.e., Lockout of Protection, Forced Switch, Manual Switch, Exercise)
3. A state request (i.e., Wait To Restore, Do Not Revert, No Request, Reverse Request), which is used to indicate the state of the APS controller when no other request is indicated.

R5-74 [179] Bits 1 through 4 of the K1 byte shall indicate the current request using the codes listed in [Table 5-6](#).

R5-75 [180] An NE using the 1:n architecture shall provide the capability to provision each working channel and the null channel (for conditions detected on the Protection line) as high or low priority, with low priority as the default. These priorities shall determine which of the listed codes is used for SF and SD requests.

14. For example, LTE could determine (from the received K1 byte) which of its transmitted OC-N signals is being selected at the far-end LTE, and could provide that information to the user.

Table 5-6 K1 Byte, Bits 1 through 4: Type of Request

Bits 1234	Automatically Initiated, External, or State Request ¹
1111	Lockout of Protection
1110	Forced Switch
1101	SF - High Priority ²
1100	SF - Low Priority
1011	SD - High Priority ²
1010	SD - Low Priority
1001	(not used)
1000	Manual Switch
0111	(not used)
0110	Wait To Restore ³
0101	(not used)
0100	Exercise ⁴
0011	(not used)
0010	Reverse Request ⁵
0001	Do Not Revert ⁶
0000	No Request

Notes:

1. Request priority is in descending order with the following exceptions.

In architectures and modes other than the 1+1 unidirectional mode, an SF request by the null channel (for an SF condition detected on the protection line) has a higher priority than a Forced Switch of Working (to Protection). Note that for systems operating in the 1+1 bidirectional mode this creates a situation in which a request of one type [Forced Switch of Protection (to Working)] is higher priority than a second type of request (SF – Null channel), which in turn is higher priority than a third type of request [Forced Switch of Working (to Protection)], but in the absence of the second type of request the first and third request types are of equal priority.

For LTE operating in the 1+1 unidirectional mode, an SF request by the null channel may be considered higher priority than a Forced Switch consistent with the criteria that appeared in this document prior to Issue 4, lower priority than a Forced Switch as shown in this table and consistent with the specifications in ANSI T1.105.01, *Synchronous Optical Network (SONET) – Automatic Protection Switching*, or intermediate between Forced Switch of Protection (to Working) and Forced Switch of Working (to Protection) consistent with the 1+1 bidirectional case described above and the specifications in ITU-T Rec. G.841, *Types and characteristics of SDH network protection architectures*.

2. High Priority codes apply only to the 1:n architecture.
3. 1+1 LTE provisioned for nonrevertive switching does not transmit Wait To Restore.
4. Exercise may not be applicable in some linear APS systems.
5. Reverse Request applies only to bidirectional systems.
6. Only 1+1 LTE provisioned for nonrevertive switching transmits Do Not Revert.

R5-76 [181] Bits 5 through 8 of the K1 byte shall indicate the number of the channel for which the request is issued, using the codes shown in [Table 5-7](#).

Table 5-7 Channel Number Code Assignments, K1 Bits 5 to 8 (and K2 Bits 1 to 4)

Code	Channel and Notes
0 (i.e., 0000)	Null channel SD and SF requests apply to conditions detected on the protection line. For 1+1 systems, Forced and Manual Switch requests apply to the protection line. Only code 0 is used with the Lockout of Protection request.
1 through 14 (i.e., 0001 through 1110)	Working channels Codes 1 through n apply in a 1:n architecture. Only code 1 applies in a 1+1 architecture. Conditions SD and SF with the provisioned priority (high/low) apply to the corresponding working lines.
15 (i.e., 1111)	Extra traffic channel May exist only when provisioned in a 1:n architecture Only No Request is used with code 15 ¹

Note:

1. In contrast to this document and ANSI T1.105.01, ITU-T Rec. G.841 contains specifications that allow several other requests to be used with code 15. If those specifications are included in a future revision of T1.105.01, then it is expected that they will also be included in any subsequent issue of this document. Note however, that based on the responses that were received regarding GR-253-ILR Issue ID 253-9, no changes are currently expected to be made in this area.

5.3.5.1.2 K1 Byte Generation

This section contains the K1 byte generation criteria, and the criteria to determine which of the possible state requests to indicate after an automatically initiated request or external request is cleared, or after a WTR state times out.

The K1 byte generation criteria are divided into three parts. These are an evaluation to determine the highest priority local request, a comparison of the highest priority local request with the current local request, and a comparison of the current local request with the remote request. The parts that apply to a particular system depend on the mode of operation.

R5-77 [182] All local requests [i.e., any locally detected SF or SD conditions, local WTR, Do Not Revert (DNR) or No Request state, or external request from a received switch command] shall be evaluated to determine the highest priority local request based on the order of request priorities in [Table 5-6](#). If local SF or SD conditions of the same priority have been detected and are still present on different lines at the same time, then the condition with the lowest channel number shall take precedence in the evaluation.

The highest priority local request shall be compared to the current local request. If the highest priority local request is of higher priority than the current local request (based on the order of priorities in [Table 5-6](#)) or if the current local request is no longer a valid request (e.g., the condition or external request that caused it has been cleared) then the highest priority local request shall replace the current local request (i.e., it becomes the new current local request). In all other cases the current local request shall not change.¹⁵

The LTE's use of the current local request in generating the transmitted K1 byte depends on its mode of operation, as follows.

Bidirectional operation:

R5-78 [183] In the bidirectional mode, the priorities of the current local request and the remote request on the received K1 byte shall be compared according to the order of priorities in [Table 5-6](#). A received Reverse Request shall not be considered in the comparison, because it assumes the priority of the request to which it is responding. The transmitted K1 byte shall be set to indicate a Reverse Request if any of the following are true:

- The remote request is of higher priority
- The requests are of the same priority, they are higher priority than a No Request, and the transmitted K1 byte is already set to indicate Reverse Request
- The requests are of the same priority, they are higher priority than a No Request, the transmitted K1 byte is not set to indicate Reverse Request, and the remote request indicates a lower channel number.

The transmitted K1 byte shall be set to indicate the local request in all other cases.

¹⁵In the bidirectional mode of operation, if the current local request is an SD or SF request that is not being indicated on the K1 byte (i.e., Reverse Request is being indicated instead), then it is acceptable for the highest priority local request to replace the current local request if they have the same priority and the highest priority local request has a lower channel number. In this situation, whether or not the current local request is changed is not critical to the operation of the linear APS system, since the switch requested by the far-end LTE will not be dropped in either case (see **R5-78 [183]**). What is critical is that if an SD or SF condition on a particular line is still present and is being indicated on the transmitted K1 byte, then it must not be replaced by another SD or SF request with the same priority but a different channel number (i.e., an existing switch that is still needed must not be dropped for another switch with the same priority but a different channel number).

One result of the above requirement is that if LTE is transmitting a particular local request, and then it receives a remote request that has the same priority and the same channel number, it will continue to transmit the local request (and consider the remote request to be a valid response). Another result is that WTR and DNR requests are normally acknowledged by a Reverse Request, while a No Request is acknowledged by a No Request.

Unidirectional operation:

-
- R5-79 [184]** In the unidirectional mode, the transmitted K1 byte shall be set to indicate the current local request (i.e., Reverse Request is never indicated).
-

When a switch action is no longer requested (i.e., an SD or SF condition clears, an external request is cleared, or a WTR state times out), a new state request is activated and is evaluated according to the rules described above. The new state request depends on the type of switching (i.e., revertive or nonrevertive), the request that is being replaced, and the channel for which that request was issued, as follows.

Working channels at LTE using revertive switching:

-
- R5-80 [185]** For working channels at LTE using revertive switching, when a local condition that caused an automatically initiated switch clears, a local WTR state shall be activated.
-

If the WTR state is the highest priority request, it is indicated on the transmitted K1 byte and maintains the switch of that channel.

-
- R5-81 [186]** A WTR state shall normally time out and become a No Request – null channel (or No Request – Channel 15, if applicable). The WTR timer shall deactivate earlier if the transmitted K1 byte no longer indicates WTR (i.e., when any request of higher priority preempts this state). When the higher priority request is cleared, the preempted WTR state shall not be reactivated. (Note however, that a new WTR state would be required to be activated if the higher priority request was for an automatically initiated switch of a working channel.)

- R5-82 [187]** When an external request is cleared, the No Request – null channel (or No Request – Channel 15, if applicable) state shall be activated (i.e., the WTR state is not activated).
-

The working channel at LTE using nonrevertive switching:

-
- R5-83 [188]** For the working channel at LTE using nonrevertive switching, the selection of the working channel from the protection line shall be maintained by activating a DNR state (instead of a WTR or No Request state). The DNR state shall be deactivated if the transmitted K1 byte no longer indicates DNR (i.e., when any request of higher priority preempts this state).
-

Note that if the most recent request was an Exercise request for the working channel, then the selector will be in the released position (see [Section 5.3.5.4](#)) and there will be no selection to maintain. Thus, the new request is expected to be No Request – null channel rather than DNR – Channel 1.

Null channel (nonrevertive and revertive):

R5-84 [189] After any request for the null channel is cleared, the No Request – null channel (or No Request – Channel 15, if applicable) state shall be activated.

5.3.5.2 K2 Byte

This section contains the bit assignments and the APS channel generation rules for the K2 byte carried on the protection line. This byte is used to indicate the bridging actions performed at the LTE, and the provisioned architecture and mode of operation.

5.3.5.2.1 Bit Assignments for the K2 Byte

R5-85 [190v2] The bit assignments for the K2 byte shall be as follows (see [Section 5.3.5.2.2](#) for the corresponding K2 byte generation rules):

- Bits 1 through 4 – A channel number using the codes shown in [Table 5-7](#).
 - Bit 5 – Indication of architecture (1+1 or 1:n), as provisioned.
 - Bits 6 through 8 – Indication of the mode of operation, as provisioned, or non-APS channel uses (i.e., AIS-L, RDI-L).
-

5.3.5.2.2 K2 Byte Generation Rules

R5-86 [191] For all architectures and modes of operation, bits 1 through 4 of the K2 byte shall be set to indicate:

- The null channel (0) if the received K1 byte indicates the null channel
 - The number of the channel bridged onto the protection line in all other cases.
-

Note that in some situations (e.g., in the bidirectional mode, when the received K1 byte contains a request for a channel that has been locked out by a “Lockout a Working Channel” command), the null channel may still be the channel bridged onto the protection line even though the received K1 byte does not indicate the null channel.

Also note that although LTE operating in the 1+1 architecture continuously transmits the working channel on the protection line, the working channel is only considered to be “bridged” for the purpose of generating the K2 byte if the received K1 byte indicates a request for a valid channel number (i.e., it is not considered bridged for the purpose of generating K2 if the received K1 byte contains any of the

invalid channel numbers, 2 through 15). Similarly, even though LTE operating in the 1:n architecture is allowed to transmit a working channel on the protection line when it is detecting a request for the null channel on the received K1 byte (see [Section 5.3.7.1.1](#)), that working channel is also not considered bridged for the purpose of generating the K2 byte. For both the 1+1 and 1:n architectures, non-null channels that are transmitted on the protection line are only considered bridged for the purpose of generating the K2 byte if they are being transmitted in response to a request.¹⁶ For example, assume 1:n LTE is receiving the No Request – null channel code on K1, and is transmitting working channel 1 on the protection line (and indicating the null channel on K2, as required). If that LTE subsequently receives a request for working channel 2, it should continue to indicate the null channel on the transmitted K2 byte until it completes the requested bridge and changes the code on K2 to ‘0010.’

R5-87 [192] Bit 5 of the K2 byte shall be set to indicate:

- Code ‘0’ if the provisioned (or only supported) architecture is 1+1
- Code ‘1’ if the provisioned (or only supported) architecture is 1:n.

R5-88 [193] Bits 6 through 8 of the K2 byte shall be set to indicate:

- Code ‘101’ if the provisioned mode is bidirectional
 - Code ‘100’ if the provisioned (or only supported) mode is unidirectional.
-

The codes ‘011,’ ‘010,’ ‘001,’ and ‘000’ in K2 bits 6 through 8 are reserved, possibly for future use in 1:n drop and insert (nested) protection switching, and the code ‘111’ is used in detecting an incoming AIS-L. Also, **R5-88 [193]** is overridden by the requirements in [Section 6.2.1.3.1](#) when a defect that causes RDI-L (i.e., ‘110’) to be generated is detected on the incoming signal on the protection line.

5.3.5.2.3 Mode of Operation

The K2 byte can be used to determine the protection configuration of the far-end LTE. As discussed in [Section 6.2.1.1.6.C](#), LTE provisioned to operate in any architecture and mode other than the 1+1 unidirectional mode is required to monitor K2 bit 5 and K2 bits 6 through 8 (on the protection line) for APS Mode Mismatch defects and failures.

R5-89 [194] An APS Mode Mismatch indication resulting from a mismatch of K2 bit 5 shall be used to modify the operation of the 1:1 LTE to interwork with the 1+1 LTE (see [Section 5.3.2.3](#)).

¹⁶This interpretation is not considered essential to the operation of the linear APS system, and therefore is not listed as a criterion. It is included here for clarification purposes only.

- R5-90 [195]** An APS Mode Mismatch indication resulting from a mismatch of K2 bits 6 through 8 shall be used by 1+1 LTE provisioned for bidirectional switching to operate unidirectionally, or by 1:n LTE provisioned for unidirectional switching to operate bidirectionally (see [Sections 5.3.2.1](#) and [5.3.2.2](#)).

In cases where both **R5-89 [194]** and **R5-90 [195]** apply (i.e., 1+1 LTE provisioned for bidirectional switching connected to 1:1 LTE provisioned for unidirectional switching, or 1+1 LTE provisioned for unidirectional switching connected to 1:1 LTE provisioned for bidirectional switching), **R5-89 [194]** applies first, and therefore the result is that the LTE at both ends will operate in the 1+1 unidirectional switching mode.

- R5-91 [196v2]** If LTE stops receiving an indication of the provisioned mode of operation from the far-end LTE, it shall maintain its current mode of operation. Similarly, if LTE has never been able to determine the architecture or mode of operation of the far-end LTE (e.g., because it has received only RDI-L, invalid codes in K2 bits 6 through 8, or AIS-L since it was last initialized), then it shall operate in its provisioned architecture and switching mode.

The near-end LTE will normally stop receiving (or will not have access to) the far-end LTE's indication of its provisioned mode when the far-end LTE inserts RDI-L on the protection line, or when an LOS or LOF defect is detected on the protection line by a regenerator or STE within the same NE (either of which, in the functional model used throughout this document, will insert AIS-L downstream, thus resulting in the reception of AIS-L by the near-end LTE on the protection line). In general, these situations are unrelated to the far-end LTE's provisioned mode of operation, and therefore there is no reason for the near-end LTE to alter its own mode of operation in response.

Unlike the cases discussed above, if the near-end LTE receives '000' (for longer than 200 ms, see [Section 6.2.1.3](#)), '001', '010' or '011' in K2 bits 6 through 8, it is likely that there is a serious provisioning problem (e.g., the far-end LTE has not been provisioned to support linear APS) or connection problem (e.g., the transmitter supporting the working line at the far end of the system is connected to the protection line receiver at the near end). In addition, in some such cases the "best possible response" might be for the near-end LTE to alter its mode of operation. However, whether or not the best possible response involves a mode change (and in cases where it does, the appropriate final mode) is highly dependent upon the particular situation. In addition, a response that might be appropriate in one situation may be inappropriate in a slightly different situation, and therefore **R5-91 [196v2]** also indicates that the near-end LTE must not alter its mode of operation in these cases.

5.3.5.3 Control of the Bridge

The criteria on the bridging of a channel to the protection line depend on the architecture and mode of operation being used, as follows. Also see [Figures 5-6](#) and [5-7](#).

R5-92 [197v3] In the 1:n architecture, the channel whose number is indicated on the received K1 byte shall be bridged to the protection line unless the request is invalid.

See [Section 5.3.5.5](#) for a discussion of various types of invalid requests and the timing of actions performed in response to the receipt of invalid codes.

R5-93 [198] If a local SF condition is detected on the protection line or a Protection Switching Byte failure is declared, the current bridge shall be maintained if the mode of operation is unidirectional, or shall be released if the mode is bidirectional.

R5-94 [199] In the 1+1 architecture, the working channel shall be continuously bridged to the protection line.

Figure 5-6 Linear APS Switch – 1:n Architecture (in released position)

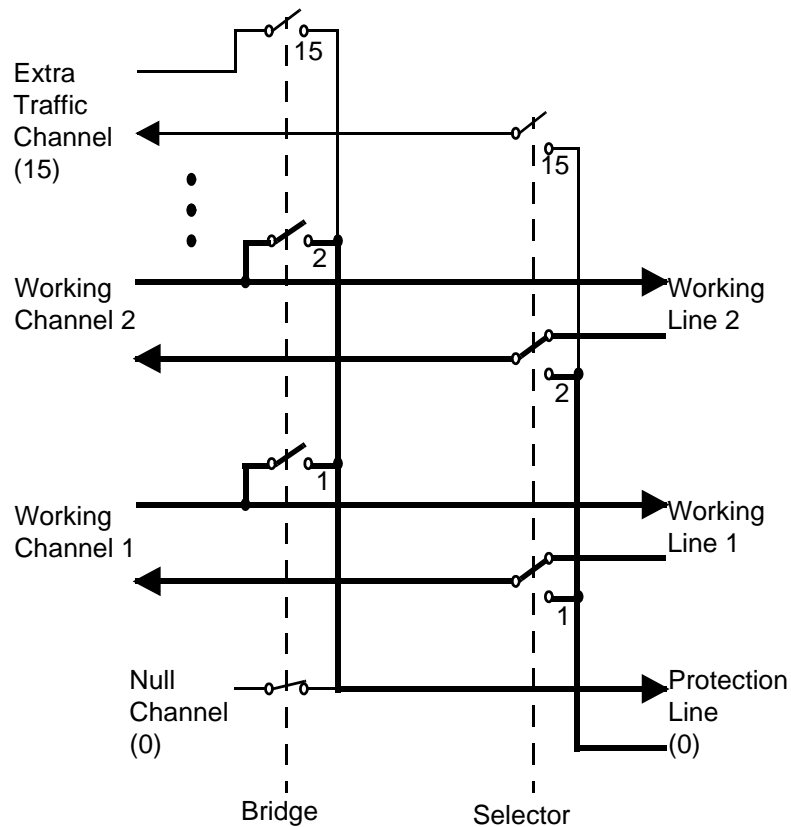



Figure 5-7 Linear APS Switch – 1+1 Architecture (in released position)



5.3.5.4 Control of the Selector

The criteria on the control of the selector depend on the architecture and mode of operation being used. In all architectures and modes except the 1+1 unidirectional mode, the selector is controlled by comparing the channel numbers indicated on the transmitted K1 and received K2 bytes, and the following criteria apply.

R5-95 [200] In all architectures and modes except the 1+1 unidirectional mode, if there is a match of the transmitted K1 and received K2 bytes, then the indicated channel shall be selected from the protection line unless one of the following is true (in which case the selector shall be in the released position, see Figures 5-6 and 5-7):

- The match is for the null channel
- An Exercise request is indicated on the transmitted K1 byte (unidirectional and bidirectional), or the received and acknowledged K1 byte (bidirectional only).

R5-96 [201v2] The selector shall also be in the released position when there is a mismatch of the channel numbers¹⁷ at LTE operating in the 1:n architecture or (with the possible exception described below) the 1+1 bidirectional mode.

17. The phrase “when there is a mismatch of channel numbers” was intended to mean “after the third of three consecutive received frames containing a channel number (or numbers) in K2 that is different than the transmitted channel number in K1.” If an NE meets the Channel Mismatch defect detection objective in Section 6.2.1.1.6.B, then “when there is a mismatch of channel numbers” can be replaced by “when a Channel Mismatch defect is detected”. However, if the NE only meets the Channel Mismatch defect requirement, then delaying the release of the selector until the Channel Mismatch defect is detected could result in traffic from one channel being misrouted to a different channel for as long as 50 ms. Therefore, the LTE must release the selector after three frames.

In the 1+1 bidirectional mode, the working channel is continuously bridged to the protection line, and therefore the selection of the working channel from the protection line can be made before a match is obtained during the process of completing a switch to the protection line (i.e., the selector does not have to remain in the released position when there is a mismatch during that process). However, if the selection is made before a match is obtained, the match must still be obtained to avoid the declaration of a Channel Mismatch failure (see [Section 6.2.1.1.6.B](#)).

In the 1+1 unidirectional mode, the highest priority local request controls the selector, and the following requirement applies.

-
- R5-97 [202]** In the 1+1 unidirectional mode, the working channel shall be selected from the protection line if channel number 1 is indicated on the transmitted K1 byte.
-

5.3.5.5 Transmission and Acceptance of Bytes K1 and K2

-
- R5-98 [203]** The linear APS protocol shall be carried between LTE in the APS channel (i.e., the K1 and K2 bytes) transmitted on the protection line.
-

The values transmitted by LTE in the K1 and K2 bytes on the working lines are undefined (except for K2 bits 6, 7 and 8 when RDI-L is being transmitted), and therefore the criteria in [Section 3.2](#) are applicable.¹⁸

-
- R5-99 [204]** A new code on the received K1 and K2 bytes shall replace the current received code if it is received identically in three consecutive frames.
-

An invalid code is defined as an unused code (e.g., see [Table 5-6](#)) or a code irrelevant for the specific operation or mode of operation (e.g., a switch request issued for a nonexistent channel). Note that as a protection switch is being completed or released, a code that was valid may become invalid. In most cases this is only a temporary situation, as the LTE that is generating that code will insert a different code that is appropriate for the new situation. However in other cases where the transmitting LTE is nonconforming to certain criteria or the system has not been correctly provisioned (e.g., where a Lockout a Working Channel command has been entered at only one end of a system operating in the 1:n bidirectional mode), it may persist. In any case, criteria such as **R5-93 [198]**, **R5-96 [201v2]** and **R5-105 [209]** apply regarding the actions (and the timing of the actions) that need to be taken when the received codes are invalid.

-
- R5-100 [205v2]** LTE shall not change its transmitted codes so that they are invalid. In addition, LTE shall not continue to transmit (previously valid) codes that become

18. It is also acceptable for the K1 and K2 bytes in the signals carried on working lines to carry the APS channel information. However, LTE cannot assume that the APS channel information will be present in the received signal, and therefore it must be capable of ignoring any APS codes received in those bytes.

invalid due to changes in the system, beyond the time required to determine and insert the appropriate new codes (see **R5-63 [174]** and **R5-64 [175]**).

- R5-101 [206]** If the capability to transport extra traffic on the protection line is provided, the No Request code shall be used to keep the extra traffic channel on the protection line (i.e., the No Request code is the only valid code to transmit with a channel number of '1111').

The preceding requirement implies that the extra traffic channel will be preempted when any request with priority higher than No Request is indicated (i.e., a Lockout of Protection, any request by a working channel to use the protection line, or an SD or SF condition detected on the protection line). One or more of the working channels can be precluded from preempting the extra traffic channel with the "Lockout a Working Channel" control command (see [Section 5.3.6.2](#)). However, there is no mechanism defined such that the extra traffic can continue to be transported when an SD condition is detected on the protection line. This issue was discussed at length in GR-253-ILR, Issue ID 253-9. Based on the comments that were received in response to that issue, it was determined that no additional criteria are needed in this area (although see Note 1 in [Table 5-7](#)).

It is expected that the LTE at both ends of a linear APS system will normally have consistent views of the priorities (for SD and SF requests) of each of the working and null channels that they jointly support. For example if one LTE in a 1:2 system is provisioned to consider working channel 1 as high priority, and working channel 2 and the null channel as low priority, then the other LTE would normally be expected to be provisioned the same way. In addition, this document and ANSI T1.105.01 both indicate that only the low priority SD and SF request code are applicable in 1+1 systems. However, the LTE in a 1:n system could easily be provisioned (or misprovisioned) with different views of the channel priorities, and ITU-T Recommendations G.841 and G.783 indicate that only the high priority codes are applicable in 1+1 systems. To promote compatibility between LTE with different views of the priorities of the various channels, the following objective is applicable.

- O5-102 [917]** LTE should not consider an SD or SF request detected on the incoming K1 bits 1 through 4 to be invalid based (solely) on the high/low priority of that request.

For example, LTE that detects a high priority SD request for a channel that it views as low priority should not consider the incoming request to be invalid based solely on the fact that its view of the requesting channel's priority does not match the priority indicated by the incoming request. Note however, that such a request could be considered invalid for other reasons (e.g., in a 1:n system operating in the bidirectional mode, it would be considered invalid if the requesting channel has been locked out using the "Lockout a Working Channel" control command).

In addition, the following criteria are applicable to LTE operating in the 1+1 bidirectional mode.

- R5-103 [207]** LTE operating in the 1+1 bidirectional mode and using nonrevertive switching shall consider the WTR code for the working channel to be valid.

O5-104 [208] LTE operating in the 1+1 bidirectional mode should consider the Do Not Revert code for either the null channel or the working channel as valid.

R5-103 [207] and **O5-104 [208]** are needed to facilitate interworking between LTE that uses revertive switching and LTE that uses nonrevertive switching. If the NEs in a 1+1 bidirectional system with one revertive NE and one nonrevertive NE do not conform to **R5-103 [207]** and **O5-104 [208]**, then they will detect and declare unnecessary Protection Switching Byte defects and failures (see [Section 6.2.1.1.6.A](#)) after many completed switch requests are cleared. As discussed in **R5-106 [210]**, the protection line is considered to be in an SF condition when a Protection Switching Byte failure is declared. That SF condition will cause an immediate switch back to the working line, defeating the purpose of both the WTR and DNR states. Conversely, if the NEs conform to the criteria, then no unnecessary failures will be declared and the type of switching used by the system will simply depend on which of the NEs is “in control” (i.e., which NE originated the most recent request that resulted in a completed switch) as described below.

- If the revertive NE is in control, then when the request is cleared that NE will normally send a WTR request (if the original request was an automatically initiated request for a working channel) or a No Request with a channel number of ‘0’ (if it was an external request or a request for the null channel). The nonrevertive NE will then normally send a Reverse Request (if it receives a WTR) or a No Request (if it receives a No Request). This is exactly the same response as would be expected if the second NE was provisioned for revertive switching rather than nonrevertive switching. Thus, the system uses revertive switching if the revertive NE is in control.
- If the nonrevertive NE is in control, then when the request is cleared that NE will normally send a DNR request (if the original request was for a working channel) or a No Request (if it was for the null channel). The revertive NE will then normally send a Reverse Request (if it receives a DNR) or a No Request (if it receives a No Request). This is exactly the same response as would be expected if the second NE was provisioned for nonrevertive switching rather than revertive switching. Thus, the system uses nonrevertive switching if the nonrevertive NE is in control.

O5-104 [208] is also needed to facilitate interworking between LTE designed to different issues of the Telcordia SONET criteria documents. Some earlier issues of those documents indicated that a nonrevertive NE was supposed to send the DNR code with a channel number of ‘0’ when it was selecting the traffic from the working line and had no higher priority requests. Other issues, including this document, indicate that such an NE is required to send the No Request code (see **R5-84 [189]**).

Receipt of an invalid code or persistently unacceptable codes in the K1 byte results in a Protection Switching Byte defect, and if that defect persists for an extended period, a Protection Switching Byte failure is declared (see [Section 6.2.1.1.6.A](#)). Similarly, receipt of an invalid code or persistently unacceptable codes in bits 1 through 4 of the K2 byte results in a Channel Mismatch defect, and if that defect persists for an extended period, a Channel Mismatch failure is declared (see [Section 6.2.1.1.6.B](#)).

R5-105 [209] Even when accepted as the current code, an invalid code in K1 shall not result in any immediate protection switching action.

R5-106 [210] For LTE operating in the bidirectional mode, the protection line shall be considered to be in the SF condition when a Protection Switching Byte failure is declared. An SF condition resulting from a Protection Switching Byte failure shall be cleared when the Protection Switching Byte failure is cleared.

For LTE operating in the unidirectional mode, the received request does not affect the transmitted request, so the preceding requirement does not apply.

5.3.6 Linear APS Commands

This section describes the linear APS commands defined to allow the user to perform protection switch actions or to provision the linear APS controller.

5.3.6.1 Switch Commands

A switch command issued at the APS controller interface initiates one external request for evaluation as described in [Section 5.3.5.1.2](#).

R5-107 [211] The following switch commands shall be provided, as described.

Clear – Clears all of the switch commands listed below, for the channel or channels specified in the command.

Lockout of Protection – Prevents any of the working channels from switching to the protection line by issuing a Lockout of Protection request [unless a request of equal priority (i.e., a Lockout of Protection) is already in effect].

Forced Switch of Working (to Protection) – Switches the specified working channel to the protection line unless a request of equal or higher priority is in effect, by issuing a Forced Switch request.

Forced Switch of Protection (to Working) – Switches the working channel back from the protection line to the working line unless a request of equal or higher priority is in effect, by issuing a Forced Switch request for the null channel. This command applies only in the 1+1 architecture.

Manual Switch of Working (to Protection) – Switches the working channel to the protection line unless a request of equal or higher priority is in effect, by issuing a Manual Switch request.

Manual Switch of Protection (to Working) – Switches the working channel back from the protection line to the working line unless a request of equal or higher priority is in effect, by issuing a Manual Switch request for the null channel. This command applies only in the 1+1 architecture.

In addition to being used to clear Lockout, Forced and Manual switch commands, the Clear command defined above may (but is not required to) be implemented such that it clears any active WTR state related to the specified line or lines.

-
- R5-108 [212]** When a higher priority local or remote request preempts an external request, the preempted request shall not be retained (i.e., when the higher priority request is cleared, the preempted switch request shall not be reinitiated).
-

In some situations it could be useful for a previously completed external request to be reinitiated after a higher priority automatically initiated request has preempted it and then been cleared (i.e., reinitiate a Manual Switch after an SD or SF request). However, allowing external requests to be retained and automatically reinitiated would also result in protection switch oscillations in some situations. Therefore **R5-108 [212]** is not expected to be changed in any future issue of this document.

In addition to the required switch commands discussed above, in some applications it may be necessary to support the Exercise command.

- CR5-109 [213]** LTE capable of operating in the 1+1 bidirectional mode or the 1:n architecture may be required to support the Exercise command.

- R5-110 [214]** If the Exercise command is supported, it shall cause the LTE to perform as described below.

Exercise – Exercises the protocol for a protection switch of the specified channel, unless a request of equal or higher priority is in effect, by issuing an Exercise request for that channel and checking the response on the APS channel.

As discussed in [Section 5.3.5.4](#), the switch is not actually completed during the Exercise routine (i.e., the selector remains released). Therefore, all actions are the same as those taken for a Manual Switch command, except that the selector remains released.

- R5-111 [215]** If the Exercise command is supported, it shall be cleared automatically at the end of the Exercise routine or the required switch completion time, whichever is sooner.
-

Support of the Exercise command is not required or expected for LTE that supports only the 1+1 unidirectional mode. For LTE operating in any of the other possible modes, the following requirement is applicable to facilitate interworking between LTE that supports the Exercise command and LTE that does not support it.

- R5-112 [216]** LTE that does not support the Exercise command shall consider an incoming Exercise request with a valid channel number to be valid, and shall respond as requested (per [Section 5.3.5](#)).
-

5.3.6.2 Control Commands

Control commands set and modify the linear APS operation. The commands that are currently defined apply only for LTE that supports the 1:n architecture.

R5-113 [217v2] The following control commands shall be supported by 1:n LTE, and shall cause the LTE to perform as described below.

Lockout a Working Channel – Prevents the specified working channel (or channels) from switching to the protection line.

Clear Lockout-a-Working-Channel – Clears the Lockout a Working Channel command for the channel or channels specified in the clear command.

The functionality of the Lockout a Working Channel command depends on the switching mode (i.e., unidirectional or bidirectional) being used. If the switching mode is bidirectional, then the lockout is also bidirectional. Specifically, local requests for the locked out channel are not considered in the K1 byte generation process, and remote requests for that channel are not acknowledged (i.e., they also are not considered in the K1 byte generation process, and the requested bridge is not performed or indicated in K2 bits 1 through 4). Similarly, if the switching mode is unidirectional, then the lockout is also unidirectional. That is, local requests for the locked out channel are not considered in the K1 byte generation process, but remote requests for that channel are acknowledged (i.e., the requested bridge is performed and indicated in K2 bits 1 through 4).

Note that in the bidirectional mode, a Lockout a Working Channel command for a particular channel must be applied at both ends of the SONET Line for proper operation.

5.3.7 Switch Operation

This section describes the operation of the linear APS system and provides an example of the operation in a tabular format.

5.3.7.1 1:n Architecture

5.3.7.1.1 1:n Bidirectional Mode

Table 5-8 gives an example of the protection switching actions between two multiplexer sites (denoted by A and C) of the 1:n bidirectional linear APS system shown in Figure 5-8. A description of the operation follows the table.

Figure 5-8 1:n Linear APS Architecture Example

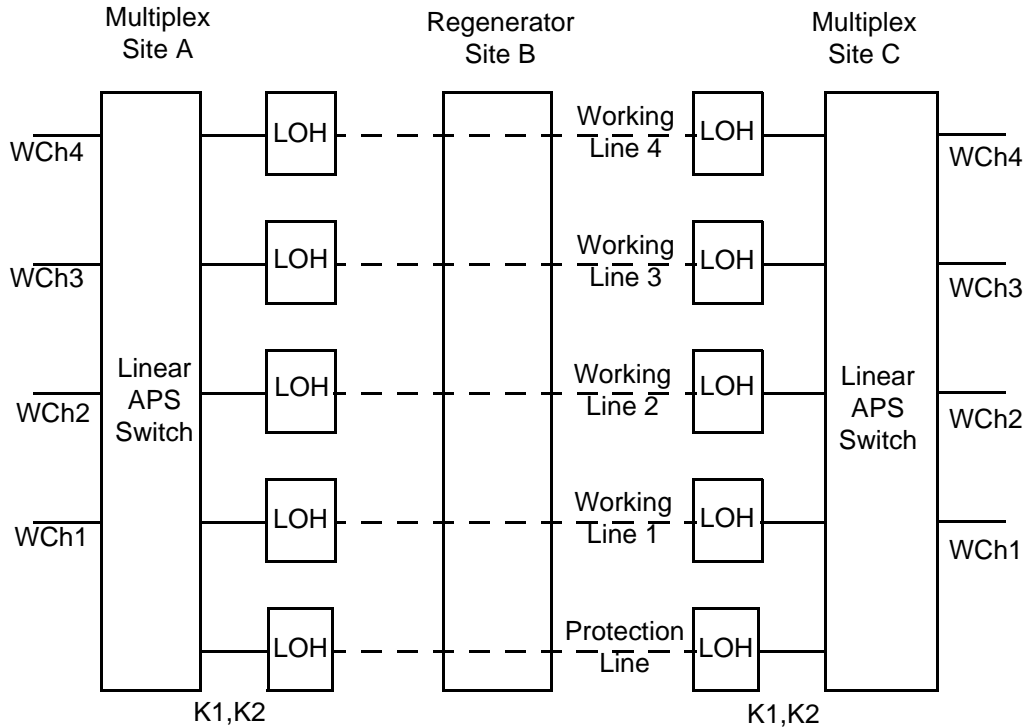


Table 5-8 1:n Bidirectional Switching Example

Failure Condition or Controller State	APS Bytes				Action	
	C → A		A → C		At Site C	At Site A
	K1 Byte	K2 Byte	K1 Byte	K2 Byte		
No Failures (Protection Line is not in use)	00000000	00001101	00000000	00001101	WCh 3 is transmitted on the protection line to provide a valid signal. Selector is released.	WCh 4 is transmitted on the protection line to provide a valid signal. Selector is released.
Working Line 2 Degraded in Direction A → C	10100010	00001101	00000000	00001101	Failure detected. Request WCh 2 bridge - SD.	
	10100010	00001101	00100010	00101101		Bridge WCh 2. Send Reverse Request for WCh 2 bridge.
	10100010	00101101	00100010	00101101	Select WCh 2. Bridge WCh 2.	
	10100010	00101101	00100010	00101101		Select WCh2. Bidirectional switch completed.

Table 5-8 1:n Bidirectional Switching Example (Continued)

Failure Condition or Controller State	APS Bytes				Action	
	C → A		A → C		At Site C	At Site A
	K1 Byte	K2 Byte	K1 Byte	K2 Byte		
Working Line 1 Failed in Direction C -> A (This preempts the WCh 2 Switch)	10100010	00101101	11000001	00101101 (Note a)		Failure detected. Request WCh 1 bridge - SF. Release WCh 2 selection.
	00100001	00011101	11000001	00101101 (Note a)	Bridge WCh 1. Send Reverse Request for WCh 1 bridge. Release WCh 2 selection.	
	00100001	00011101	11000001	00011101		Select WCh 1. Bridge WCh 1.
	00100001	00011101	11000001	00011101	Select WCh 1. Bidirectional switch completed.	
Working Line 1 Repaired (Working Line 2 still Degraded)	00100001	00011101	01100001	00011101		Send Wait To Restore for WCh 1.
	10100010	00011101 (Note a)	01100001	00011101	Request WCh 2 bridge - SD. Release WCh 1 selection.	
	10100010	00011101 (Note a)	00100010	00101101		Bridge WCh 2. Send Reverse Request for WCh 2 bridge. Release WCh 1 selection.
	10100010	00101101	00100010	00101101	Bridge WCh 2. Select WCh 2.	
	10100010	00101101	00100010	00101101		Select WCh 2. Bidirectional switch completed.
Working Line 2 Repaired	01100010	00101101	00100010	00101101	Send Wait To Restore for WCh 2.	

Table 5-8 1:n Bidirectional Switching Example (Continued)

Failure Condition or Controller State	APS Bytes				Action	
	C → A		A → C		At Site C	At Site A
	K1 Byte	K2 Byte	K1 Byte	K2 Byte		
Wait to Restore Expired (No Failures)	00000000	00101101	00100010	00101101	Drop WCh 2 request.	Release WCh 2 selection.
	00000000	00101101	00000000	00001101		Release WCh 2 bridge. (Transmit WCh 4 on the protection line.) Drop WCh 2 request. Release WCh 2 selection.
	00000000	00001101	00000000	00001101	Release WCh 2 bridge. (Transmit WCh 3 on the protection line.)	
No Failures (Protection Line is not in use)	00000000	00001101	00000000	00001101	WCh 3 is transmitted on the protection line to provide a valid signal. Selector is released.	WCh 4 is transmitted on the protection line to provide a valid signal. Selector is released.

Note:

- Note that Corrigendum 1 to ITU Rec. G.841 indicates that K2 bits 1 through 4 are supposed to be set to '0000' at these points in the sequence. While the transmission of that code would not be expected to adversely impact system performance or cause interoperability issues with LTE that transmits the codes shown here, it would be inconsistent with the intent of the criteria in this document. For example, in the scenario described in this table, when LTE A detects the failure of working line 1, it is still receiving a request for working channel 2 in the K1 byte (on the protection line). According to **R5-92 [197v3]**, it is therefore supposed to bridge working channel 2 to the protection line (and signal accordingly in its outgoing K2 byte, see **R5-86 [191]**) "unless the request is invalid". While it could be argued that this last phrase was meant to indicate that the bridge is supposed to be released immediately (since the incoming SD request is lower priority than the outgoing SF request, it is invalid for the particular switching operation), other criteria (e.g., **R5-93 [198]**, **R5-105 [209]**) indicate that any actions that are taken as a result of an invalid code on the received K1 byte are required to be delayed until a Protection Switching Byte failure is declared. Thus according to the criteria in this document, LTE A would need to release its bridge and send '0000' in K2 bits 1 through 4 only if the incoming SD request were to persist for approximately 2.5 seconds.

When the protection line is not in use, the null channel is indicated on the K1 bytes. The null channel is also indicated on the K2 bytes, according to K2 byte generation rules. However, any working channel may actually be transmitted on the protection line. The receiving equipment does not assume or require any specific channel. In the example in **Table 5-8**, Working Channel (WCh) 3 is transmitted at site C, and WCh 4 is transmitted at site A.

When an SF or SD condition is detected or a switch command is received, the APS controller at that end (the “tail end”) compares the priority of this new condition with the request priority of the channel (if any) on the protection line. The comparison includes the priority of any remote request on the received K1 byte. If the new request is of higher priority, the K1 byte is loaded with the request and the number of the channel asking for use of the protection line. In the example, an SD condition is detected at C on working line 2, and an SD request is sent on the K1 byte to A.

At the far end (the “head end”), when this new incoming K1 byte has been verified (i.e., received identically for three consecutive frames) and evaluated by the APS controller, the outgoing K1 byte is sent with a Reverse Request and the appropriate channel number as a confirmation for that channel to use the protection line, and to request a bridge at the tail end for the same channel. A Reverse Request is returned for all requests except a No Request. This clearly identifies which end originated the switch request. If the head end also originates an identical request (not yet confirmed by a Reverse Request) for the same channel, then both ends continue transmitting the identical K1 byte and perform the requested switch action.

Also, at the head end, the indicated channel is bridged to the protection line. When the channel is bridged, the K2 byte is set to indicate the number of the channel on the protection line.

At the tail end, when the channel number on the received K2 byte matches the number of the channel requesting the switch, that channel is selected from the protection line. This completes the switch to the protection line for one direction. The tail end also performs the bridge as requested by the K1 byte and indicates the bridged channel on the K2 byte. The head end completes the bidirectional switch by selecting the channel from the protection line when it receives a matching K2 byte.

Normally, a switch is completed within 50 ms. If the switch cannot be completed because an appropriate code is not returned in the K1 byte sent by the head end to the tail end, or because one end does not perform and indicate the appropriate bridge, then a Protection Switching Byte defect and/or a Channel Mismatch defect is detected. If a one of these defects persists for 2.5 seconds, then the appropriate failure is declared. In the case of a Protection Switching Byte failure, the end declaring the failure changes its transmitted K1 byte to indicate an SF on the protection line.

Table 5-8 further illustrates how an existing switch is preempted by a higher priority request. In the example, an SF condition on WCh 1 preempts the WCh 2 switch. The selectors are temporarily released before selecting WCh 1, because of the temporary channel number mismatches on the transmitted K1 and received K2 bytes. The example also illustrates the switch back to WCh 2 after the failure on working line 1 is repaired.

When the switch is no longer needed (e.g., the failure on working line 2 is repaired) and the WTR time has expired, the tail end indicates a No Request – null channel on the K1 byte. This releases the selector because of a channel number mismatch. The head end then releases the bridge and replies with the same K1 byte and the null channel number on the K2 byte. The selector at the head end is also released because of a channel number mismatch. Receiving the number for the null channel on the K1 byte causes the tail end to release the bridge. Because the K2 bytes now

indicate the null channel, and that matches the null channel numbers on the K1 bytes, the selectors remain released, any Channel Mismatch defects are terminated, and restoral is completed.

5.3.7.1.2 1:n Unidirectional Mode

All actions are as described in [Section 5.3.7.1.1](#) for the bidirectional mode, except that the unidirectional switch is completed when the tail end selects the channel for which it issued a request (on the K1 byte) from the protection line. This difference in operation is obtained by not considering remote requests in the priority logic and, therefore, by not issuing Reverse Requests.

5.3.7.2 1+1 Architecture

5.3.7.2.1 1+1 Bidirectional Mode

For 1+1 bidirectional systems, the K1 and K2 bytes are exchanged as [Section 5.3.7.1.1](#) describes to complete a switch, except that the head end maintains a continuous bridge of the working channel to the protection line, and therefore separate bridging actions are not performed for each request.

In addition, each end is allowed to switch immediately, before receiving a bridge confirmation from the other end. However, the channel matches must still be obtained to avoid Channel Mismatch failures.

Finally, for revertive switching the restoral takes place as described in [Section 5.3.7.1.1](#). For nonrevertive switching (assuming the working channel is being selected from the protection line) when the working line is repaired or a switch command is cleared, the tail end maintains the selection and indicates Do Not Revert for WCh 1. The head end also maintains the selection and continues indicating Reverse Request. The Do Not Revert is removed only when preempted by a signal degrade or failure condition, or an external request.

5.3.7.2.2 1+1 Unidirectional Mode


For 1+1 unidirectional switching, the working channel is continuously bridged to the protection line, and the channel selection is based on only the local conditions and requests. Therefore, each end operates independently of the other end, and the K1 and K2 bytes are not needed to coordinate switch actions. However, the K1 byte is still used to inform the other end of the local action, and the K2 byte is set to indicate that the K1 byte is being received (i.e., by indicating the same channel number as the received K1), and to inform the other end of the provisioned architecture and mode of operation.

5.4 Network Synchronization

SONET uses the existing synchronization network as described in GR-436-CORE and ANSI T1.101. This section discusses applications with respect to SONET NEs and service provider synchronization networks, timing modes for SONET NEs, criteria for SONET NE internal clocks, criteria for SONET-based timing distribution, timing reference switching criteria, and criteria regarding the use of synchronization status messages. This section also indicates which of the criteria in GR-1244-CORE are applicable to SONET NEs.

Note that in this section “OC-N” and “STS-N electrical” are generally used to identify the high-speed SONET “interfaces”¹⁹ to a particular NE, and “OC-M” and “STS-M electrical” are used to identify any tributary SONET “interfaces”. For example, in [Figure 5-9](#) the ADM’s OC-12 “interfaces” are OC-N “interfaces” and its OC-3 “interfaces” are OC-M “interfaces”.

Figure 5-9 OC-N and OC-M Example



5.4.1 SONET NE Clock Applications

In general, SONET NEs are required to have internal clocks of ± 20 -ppm minimum free-run accuracy. Based on the work done in T1X1 on ANSI T1.105.09, *Synchronous Optical Network (SONET): Network Element Timing and Synchronization*, clocks that are contained in NEs that support SONET Line terminating functions, and that meet only this minimum accuracy requirement are called SONET Minimum Clocks (SMCs).

19. Throughout Section 5.4, an OC-N (or OC-M) “interface” is defined to include all of the OC-N (or OC-M) signals that are part of a “line APS” system (if supported). In turn, line APS is used to refer to protection switching architectures where an NE receives two or more OC-N (or OC-M) signals from a peer NE, and one of those signals is used to provide SONET Line layer protection for the working channel(s) carried on the other signal(s). This includes systems that support linear APS (see [Section 5.3](#)), and span-switching in 4-fiber bidirectional line switched rings (see GR-1230-CORE). If no line APS is supported for a particular pair of OC-N or OC-M signals (e.g., for the OC-N signals that connect two adjacent nodes in a UPSR), then the “interface” includes only those two (i.e., one incoming and one outgoing) signals. In this section, quotation marks will be used to distinguish between an “interface” as defined here and other types of interfaces that are commonly defined (e.g., a point in the transmission medium at which the physical layer characteristics of a signal are specified and measured).

Many synchronization-related criteria are application-specific, particularly for SONET ADMs. Either external or line timing may be appropriate for an ADM, and protection switching schemes and dropped OC-M and STS-M electrical signals further complicate the selection of synchronization options. In general, most of the information and criteria applicable to a SONET ADM appear in this document; however, GR-496-CORE also contains several criteria related to the timing options that should be supported. Also note that other types of NEs, such as DCSs and SONET regenerators, have specific synchronization criteria that are covered in GRs and TRs specific to those types of NEs.

The following general statements describe conditions necessary for synchronization and SONET networks to be compatible:

- Where Building Integrated Timing Supply (BITS) timing is available (see GR-436-CORE), SONET NEs are externally timed from the BITS clock
- Where no BITS timing is available, SONET NEs are timed from a received OC-N (or OC-M) signal
- External-timing references to a SONET NE are from a BITS clock of stratum 3 or better quality
- Timing signals delivered to the synchronization network from a SONET NE are derived directly from a terminating OC-N (or OC-M)
- Stratum 3E or 2 clocks are used between non-SONET timing distribution networks²⁰ and any SONET NEs containing SMC or stratum 3 clocks that require timing that is traceable from or through those networks.

5.4.1.1 Physical Interface to Synchronization Network

The physical interface for synchronization signals is important so that SONET NEs can be easily integrated into the BITS plan. BITS clocks provide two types of timing outputs. These are DS1 and Composite Clock (CC) outputs, both of which are balanced signals. As described in [Section 5.4.3.1](#), DS1 signals are the usual timing reference signals for most SONET NEs, but CC signals are required for SONET NEs that support DS0 interconnections (see [Section 6](#) of GR-1244-CORE). In addition, the derived DS1 signals from SONET NEs may be used as references for the BITS clock.

[Section 3.2.1.2](#) of GR-1244-CORE describes various criteria for the physical interface to the synchronization network, such as termination requirements and wire-wrap terminals.

R5-114 [918v2] The physical interface between the SONET NE and the synchronization network shall meet the criteria²¹ in [Section 3.2.1.2](#) of GR-1244-CORE.

²⁰ “Non-SONET synchronization distribution network” is used here to refer to a synchronization network where a timing reference signal may contain a significantly higher amplitude of wander than allowed by [Figures 5-14](#) and [5-15](#) of this document [i.e., wander corresponding to the wander tolerance Time Deviation (TDEV) mask in GR-1244-CORE].

5.4.1.2 TDEV and MTIE Measurements

The discussion of the TDEV and Maximum Time Interval Error (MTIE) parameters that appeared in this section in Issue 1 of this document now appears in GR-1244-CORE.

5.4.2 Synchronization Status Messages

Synchronization status messages have been defined as a nibble (bits 5 to 8) in the S1 byte of the SONET LOH and as a bit-oriented message in the data link of ESF DS1 signals. These messages contain clock quality labels that allow (for example) a SONET NE to select the most suitable synchronization reference from the set of available references. The purpose of these messages is to allow SONET NEs and Timing Signal Generators (TSGs) to reconfigure their synchronization references autonomously while avoiding the creation of timing loops. However, it is critical to realize that the use of synchronization status messages alone will not preclude the creation of timing loops. Synchronization engineering following the guidelines in GR-436-CORE is still required.

Table 5-9 lists the synchronization status messages that have been defined at this time for the S1 byte and the ESF DS1 format. This table is based on the information contained in ANSI T1.101, as are the criteria in this section related to the interoperability of equipment supporting the set of messages that appear here (the “Generation 2” message set) and older equipment that supports the set of messages that appeared in Issue 2 of this document (the “Generation 1” message set).

21. Note that in cases where a requirement in this section indicates that the criteria in a section of another document “shall” be met, the intent is not to upgrade any objectives or conditional requirements contained in the other document to requirements. Rather, requirements in the other document need to be treated as requirements, objectives as objectives, and conditional requirements as conditional requirements.

Table 5-9 Synchronization Status Message Definitions

Description	Acronym	Quality Level	DS1 ESF Data Link Code Word ^a	S1 bits 5678 ^b
Stratum 1 Traceable	PRS	1	00000100 11111111	0001
Synchronized - Traceability Unknown	STU	2	00001000 11111111	0000 ^c
Stratum 2 Traceable	ST2	3	00001100 11111111	0111
Transit Node Clock Traceable ^d	TNC	4	01111000 11111111	0100
Stratum 3E Traceable ^d	ST3E	5	01111100 11111111	1101
Stratum 3 Traceable	ST3	6	00010000 11111111	1010
SONET Minimum Clock Traceable	SMC	7	00100010 11111111	1100
Stratum 4 Traceable	ST4	8	00101000 11111111	N/A
DON'T USE for Synchronization	DUS	9	00110000 11111111	1111
Provisionable by the Network Operator ^e	PNO	user assignable	01000000 11111111	1110

Notes:

- a. The ESF synchronization status messages are transmitted rightmost bit first. Also, in all of the applications for ESF messages described in this document (i.e., external timing references from a BITS clock and the derived DS1 outputs from a SONET NE) the DS1 is dedicated to synchronization, so the ESF messages are sent continuously (see **R5-202 [294]**).
- b. The S1 bits are transmitted leftmost bit first.
- c. The S1 code for the “Synchronized – Traceability Unknown” message was intentionally selected to be an all-zeros pattern because that is the preferred way to populate SONET overhead bits and bytes that are unused (see **Section 3.2**). Thus, the expected output of an NE when synchronization status messaging is not supported matches the message that indicates “Synchronized – Traceability Unknown.” (Note that the same strategy could not be used in defining that message on the ESF data link.) In general, the “Synchronized – Traceability Unknown” message was prevalent throughout the network when synchronization status messaging was first implemented. However as the embedded base of SONET NEs and BITS clocks has been upgraded to support synchronization status messaging, the prevalence of that message has diminished.
- d. Except for the purpose of defining synchronization status messages, TNC clocks are not considered in this document. In addition, these codes (i.e., the codes for TNC and ST3E) were not included in the Generation 1 set of synchronization status messages.
- e. This message was previously referred to as the Reserved for Network Synchronization Use (RES) message.

The use of synchronization status messages in the S1 byte can provide the following benefits:

- Automatic reconfiguration of line-timed rings
- Improved reliability of interoffice timing distribution
- Trouble-shooting of synchronization-related problems.

Therefore, full support for synchronization status messages is required for almost all line-side SONET signals (and is allowed for the remaining such signals).

R5-115 [224v2] SONET LTE shall generate and provide the capability to process the synchronization status messages listed in [Table 5-9](#) on bits 5 through 8 of the S1 bytes of all signals at SONET line-side “interfaces” (except for lines 2 through n at an OC-N “interface” where 1:n linear APS is being used).

In addition, synchronization status messages can be useful as a troubleshooting tool for drop-side signals at network interfaces (such as at a customer premise location), even if those signals are not required to be used as timing references.

R5-116 [225v4] SONET LTE shall generate the synchronization status messages listed in [Table 5-9](#) on bits 5 through 8 of the S1 bytes of all signals at SONET drop-side “interfaces” (except for lines 2 through n at an OC-M “interface” where 1:n linear APS is being used).

CR5-117 [1048v2] SONET LTE may be required to provide the capability to process the synchronization status messages listed in [Table 5-9](#) on bits 5 through 8 of the S1 bytes of all signals at SONET drop-side “interfaces” (except for lines 2 through n at an OC-M “interface” where 1:n linear APS is being used).

Messages in the DS1 ESF data link are not needed to realize the benefits listed above. Therefore, a SONET NE is only conditionally required to support DS1 ESF messages (see **CR5-197 [287v3]** and **CR1244-8**). However, some service providers believe that the use of DS1 ESF synchronization status messages will add robustness to their interoffice synchronization distribution networks. Therefore, support for DS1 ESF synchronization status messages will be required by some service providers. In addition, it is important to note that the benefit of preventing timing loops for interoffice synchronization distribution (particularly for SONET rings) can only be realized if synchronization status messaging is supported by all TSGs and SONET NEs. However, even if synchronization status messaging is implemented for all SONET signals and ESF DS1s, careful synchronization planning is still necessary.

In the criteria shown and referenced above, the use of the words “generate” and “provide the capability to process” is intended to indicate that synchronization status messages must always be generated on outgoing signals, but may not need to be processed on incoming signals. This is similar to the criteria for various other overhead bits and bytes, such as those used for REI-L, REI-P and REI-V (e.g., the M1 byte). Those bits and bytes are required to always be generated on originating Line and path signals, but are only required to be processed if a far-end PM accumulation feature is active for the incoming Line or path. [Note however that in the case of

synchronization status messages, a user can effectively override the required generation of those messages at any desired “interfaces” by provisioning them to be set to the DUS message (see **R5-232 [324]**).

As for the processing of the S1 byte contained in incoming SONET signals, the following criteria apply.

CR5-118 [1049] A SONET NE that contains LTE and supports line-timing (or through-timing), or that is capable of being provisioned to derive DS1s from the incoming signals at one or more of its SONET “interfaces”, may be required to be able to be provisioned by the user to ignore the incoming S1 byte at its provisioned reference and/or derived DS1 source “interfaces”.

R5-119 [1050] If a SONET NE allows the user to disable the processing of the S1 byte at its SONET “interfaces” that are provisioned as timing references or derived DS1 sources, the default shall be that the S1 byte is processed.

In general it is assumed that in any particular application, an NE will either use the messages received on all of its provisioned reference and derived DS1 source “interfaces”, or not use the messages on any of them. Therefore it would be sufficient to provide the capability to disable the use of the incoming synchronization status messages on a per-NE basis. However, it is also acceptable for that capability to be provided on a per-reference basis, or a per-derived-DS1 or per-derived-DS1-source basis. In addition, note that if the use of the incoming messages has been disabled, then various other criteria (or parts of criteria) in this document will not be applicable (e.g., **R5-215 [302v2]** on considering a reference failed or unavailable based on its synchronization status message). Finally, note that for interfaces that are not or cannot be provisioned as timing references or derived DS1 sources, the only defined use of the incoming synchronization status messages is that the NE should allow the user to retrieve those messages. That capability is supposed to be provided independent of whether processing has been enabled or disabled (see **O5-226 [320v4]**), and therefore it is not necessary for an NE to support the capability to disable processing at those interfaces.

It is important to note that interoperability issues exist in environments where S1 synchronization status messaging is not uniformly supported (e.g., because of an embedded base that was deployed before that capability was developed). For example, the creation of timing loops cannot be precluded in certain scenarios (e.g., in a line-timed ring where some NEs support synchronization status messaging and others do not, a timing loop may be created if the NEs are allowed to reconfigure their timing sources autonomously). This is mainly due to the fact that equipment that does not support S1 synchronization status messages will generate an all-zeros code in the S1 byte, which corresponds to the “Synchronized – Traceability Unknown” message. Careful engineering is required in mixed environments to ensure that timing loops are not created.

In addition to environments where some NEs support synchronization status messages and others do not, there may be environments in which some NEs support the Generation 2 message set and other (old) NEs support the Generation 1 set. To promote interoperability in these cases, the following criteria apply to the Generation-2 NEs.

CR5-120 [1177] A SONET NE that supports the synchronization status message set shown in [Table 5-9](#) (the Generation 2 message set), may be required to be able to interoperate with equipment that supports the message set that appeared in Issue 2 of this document (the Generation 1 message set).

R5-121 [1178] If the capability to interoperate with Generation 1 equipment is claimed, that function shall be user provisionable on a per-“interface” or derived-DS1 basis.

R5-122 [1179] If the capability to interoperate with Generation 1 equipment has been activated by the user for a particular “interface” or derived DS1, the ST3 message shall be generated on the outgoing signals at that “interface” or derived DS1 port in any situation in which the TNC or ST3E message would otherwise be sent.

Finally, Telcordia has not assigned a use for the “Provisionable by the Network Operator” message.

R5-123 [222v2] The “Provisionable by the Network Operator” message shall be treated as a “DON’T USE for Synchronization” message at intercarrier interfaces.

R5-124 [223v2] Any proprietary use of the “Provisionable by the Network Operator” message shall be clearly documented as per [Section 3.2](#).

Synchronization status messages impact criteria in many subsections of [Section 5.4](#). Specifically, external-timing mode criteria, derived DS1 criteria, and reference switching criteria are all affected. Also, a section specific to synchronization status message validation and generation is included.

5.4.3 SONET Timing Modes

This section describes the four timing modes for SONET NEs: external, line, loop and through. The timing mode determines the source of timing for OC-N/OC-M, STS-N/STS-M electrical and synchronous DS1 signals transmitted by the NE, along with the STS and VT SPEs created by that NE (with the possible exception of VT1.5 SPEs containing byte-synchronously mapped DS1s, as discussed in [Section 3.4.1.1.1](#)). Specific documents that describe the application of an NE identify which timing mode or modes the NE is required to support.

R5-125 [226] For NEs that support the external-timing mode, that mode shall be the default timing mode.

R5-126 [227] For NEs that support automatic switching between timing modes, the switching shall be revertive.

Note that the above switching requirement applies to *timing mode* switching and not to *reference* switching or *hardware* switching. In general, NEs are provisioned to use only a single timing mode and this requirement would not be pertinent. For example, it is *not* expected that an NE would be provisioned to enter line-timing if

all of its external references fail. As discussed in GR-436-CORE, maintaining and administering a network (especially avoiding the creation of timing loops) becomes difficult if this type of switching is allowed. The one case where timing mode switching would apply is for NEs that switch from through-timing to line-timing when an OC-N line fails, as per Section 5.4.3.4.

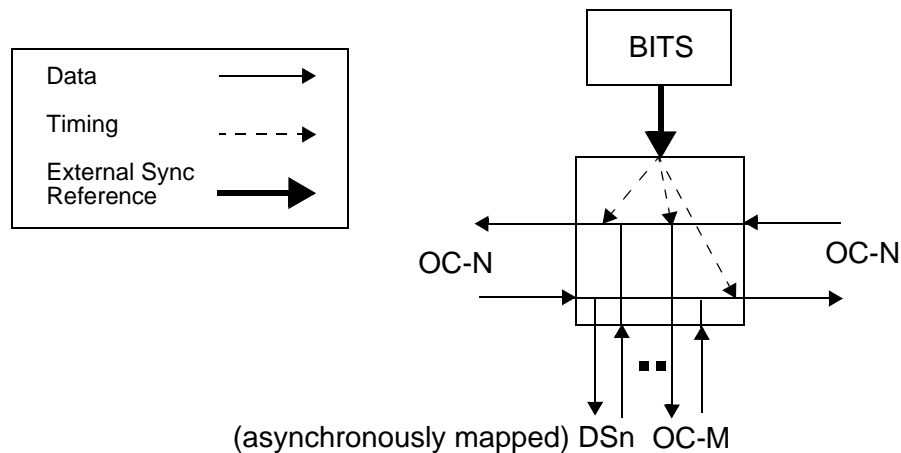
Sections 5.4.3.2, 5.4.3.3 and 5.4.3.4 describe methods of recovering timing from a terminating OC-N. Use of one of these timing modes is necessary when no external timing reference is available. The ability to recover timing from terminating STS-N electrical signals is not required.

5.4.3.1 External Timing

In accordance with the BITS concept, external timing from a BITS clock is the preferred mode of synchronizing SONET NEs that contain LTE. Figure 5-10 illustrates external timing.

R5-127 [919v2] SONET NEs with external-timing interfaces shall meet the external timing criteria in Section 3.2.1.1 of GR-1244-CORE.

Figure 5-10 External-Timing Mode Example



5.4.3.2 Line Timing

Figure 5-11 illustrates the line-timing mode for an ADM. Some suppliers have suggested that an alternative to line timing would be to use a derived DS1 (defined in Section 5.4.5.1) as an external reference. This may be acceptable in some applications, but cannot be called line timing.

R5-128 [920v2] A SONET NE with line-timing “interfaces” shall meet the line-timing criteria in Section 3.2.2 of GR-1244-CORE, with the exception that OC-M/STS-N and

synchronous DS1 interfaces are not necessarily required to be able to be provisioned as synchronization sources.

Note that the criteria in GR-1244-CORE allow (for example) an ADM to switch between its “east” and “west” “interfaces” if the “interface” currently used as the timing reference becomes unavailable as defined in Section 5.4.6. However, provisioning multiple “interfaces” as synchronization references must be done with care to avoid the creation of timing loops. Refer to GR-436-CORE for more information about synchronization planning and administration. In general, NEs without S1 synchronization status messaging capabilities should have only one OC-N “interface” provisioned as a reference.

While it is not required (or recommended by Telcordia) that an NE that supports line APS also support the provisioning of working line 1 and the protection line at a single SONET “interface” as separate references, such provisioning is allowed and may be supported by some NEs. In addition, some of those NEs may not allow the user to provision a line APS “interface” as a single reference. In that case, the following requirement applies.

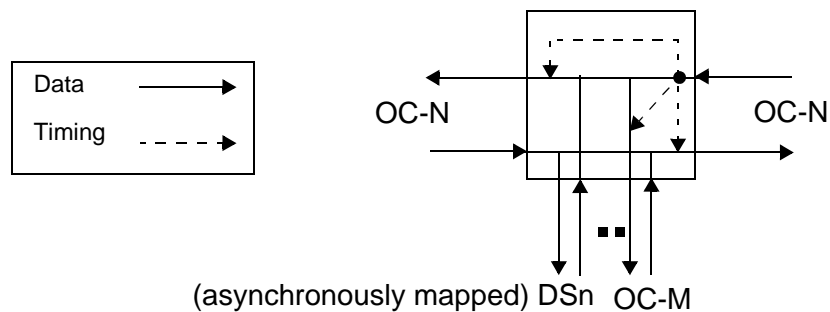
R5-129 [1051] If an NE that supports line APS does not support provisioning of an OC-N/M “interface” as a single reference, that fact shall be clearly documented.

Note that even if an NE supports the provisioning of working line 1 and the protection line as separate references (or as separate derived DS1 sources), it still must meet the applicable criteria related to the generation of the DUS synchronization status message (e.g., **R5-242 [330]**, **R5-234 [326v3]**).

CR5-130 [239] In some applications the NE may be required to provide the user the capability to provision a line-side OC-M “interface” as a synchronization source.

A typical application where this would apply is one where the timing distribution path is through a tributary “interface” on an ADM. For example, a ring on a customer’s campus environment that has a low-speed SONET connection to a central office (to carry the traffic going on and off of the campus) may need to receive timing from the central office via that SONET “interface”.


Figure 5-11 Line-Timing Mode Example



5.4.3.3 Loop Timing

Loop timing is a special case of line timing. It applies to NEs that have only one OC-N “interface” (see [Figure 5-12](#)).

Figure 5-12 Loop-Timing Mode



5.4.3.4 Through Timing


Through timing is not a recommended timing mode for SONET NEs that contain LTE (e.g., ADMs). For such NEs, the through-timing mode can be a very complex timing scheme. For example, the user may find it unclear what the timing source is for a tributary “interface”, or for the protection line at an NE that supports line APS (particularly if a 1:n APS architecture is being used). On the other hand, through timing is the required timing mode for regenerators (see TR-NWT-000917), and is supported by some other types of deployed SONET NEs. Therefore, aspects of the timing mode must be specified.

[Figure 5-13](#) illustrates the through-timing mode for an ADM. The following criteria apply to NEs configured for through timing.

-
- R5-131 [240v2]** When an NE is through timed, the transmitted signals at the “west” OC-N or STS-N electrical “interface” shall be timed by the terminating signals at the “east” OC-N or STS-N electrical “interface”, and the transmitted signals at the “east” “interface” shall be timed by the terminating signals at the “west” “interface”.
- R5-132 [241]** An ADM that supports through timing shall provide the user with the capability to provision for automatic switching to line timing using the non-failed OC-N “interface” when one of the OC-N “interfaces” becomes unavailable as a reference, as defined in [Section 5.4.6](#).
- R5-133 [242]** If a through-timed ADM has OC-M, STS-M electrical, or synchronous DS1 interfaces, the timing source for these outputs, as a group or individually, shall be user provisionable from either of the OC-N “interfaces”.
-

Note that in a chain of through-timed ADMs provisioned to perform the type of switching described in **R5-132 [241]**, the use of synchronization status messages would be essential in order to prevent the formation of timing loops.

Figure 5-13 Through-Timing Mode Example



5.4.4 SONET Internal Clock

This section provides criteria applicable to an NE's internal clocks. All clock performance characteristics are measured at the synchronous outputs (e.g., OC-N outputs) of the NE.

Note that the criteria in this section are intended for SONET NEs that have stratified clocks or SMCs, and that some SONET NEs are not expected to provide such clocks. In particular, SONET NEs that do not contain LTE (e.g., regenerators) need only provide a free-running clock with ± 20 ppm accuracy (see TR-NWT-000917). That clock is used to generate AIS-L when the NE detects an incoming signal defect (i.e., an LOS or LOF), and therefore is not receiving a valid timing reference.²² For such NEs, the Category II jitter transfer requirement in Section 5.6.2.1.2 is applicable, rather than the wander transfer and generation criteria in this section.

5.4.4.1 Stratum Clocks for SONET Applications

Stratum clock criteria are contained in GR-1244-CORE. Some applications require that SONET NEs provide clocks that meet, as a minimum, stratum 3 criteria.

R5-134 [1109v2] As a minimum, SONET NEs that support Line termination functionality for OC-192 and/or OC-768 signals shall provide stratum 3 clocks.

CR5-135 [243v2] Some service providers may require stratum 3 clocks for SONET NEs used in applications that do not explicitly require stratum 3 clocks.

In addition to stratum 3 clocks (and SMCs), some SONET NEs may support stratum 3E or 2 clocks so that they can be used, for example, as timing distribution hubs. In general, that type of application is not consistent with the BITS concept or recommended by Telcordia, and therefore the conditional requirements that appeared in early issues of this document to indicate that SONET NEs might be

²² Although the accuracy requirements for an SMC and a regenerator's clock are identical (± 20 ppm), the regenerator's clock is not required to meet the other performance criteria applicable to an SMC and therefore it is not considered to be an SMC.

required to support stratum 3E or 2 clocks have been removed. In addition, to avoid performance issues that can occur when multiple stratum 3E or 2 clocks are deployed in a single office (e.g., in the BITS clock and one or more SONET NEs), the following objective applies to NEs that support those types of clocks.

O5-136 [1110] A SONET NE that can be equipped with stratum 3E or 2 internal clocks should also be able to be equipped with SMCs or stratum 3 internal clocks.

R5-137 [246v2] A stratum 3, 3E or 2 clock in a SONET NE shall meet the following criteria in GR-1244-CORE:

- Minimum free-run accuracy (GR-1244-CORE, Section 5.1)
- Holdover stability (GR-1244-CORE, Section 5.2)
- Pull-in/hold-in range and settling times (GR-1244-CORE, Section 3.5)
- Phase changes during pull-in (GR-1244-CORE, Section 5.8).

R5-138 [247v2] A stratum 3E or 2 clock in a SONET NE shall meet the following requirements in GR-1244-CORE:

- Wander tolerance (GR-1244-CORE, Section 4.3)
- Wander transfer (GR-1244-CORE, Section 5.4)
- Phase transients during both the various synchronization rearrangement operations listed in GR-1244-CORE and the additional, SONET-specific rearrangement listed in [Section 5.4.4.3.3](#) of this document (GR-1244-CORE, Section 5.6).

R5-139 [921v2] A stratum 3 clock in a SONET NE shall meet the SMC wander tolerance and transfer requirements in [Section 5.4.4.2.4](#) (of this document).

Note that the “output phase transient in response to an input phase transient” requirement in this document (i.e., **R5-158 [1182]**) is more stringent than the corresponding requirement in GR-1244-CORE, and therefore overrides that requirement for SONET NEs containing stratum 3E or 2 clocks. Also, see [Section 5.4.4.3.7](#) regarding the applicability of the phase build-out criteria in Section 5.7 of GR-1244-CORE.

5.4.4.2 SONET Minimum Clock Applications

In general, the following criteria apply to the clocks in all SONET NEs that contain LTE and do not contain a stratum clock (i.e., to SMCs as defined in ANSI T1.105.09). Note however, that [Section 5.4.4.2.4](#) also applies to stratum 3 clocks, as per **R5-139 [921v2]**.

5.4.4.2.1 Free-Run Accuracy for SMCs

Free-run accuracy is defined in Section 5.1 of GR-1244-CORE.

R5-140 [248] The minimum free-run accuracy of an SMC shall be ± 20 ppm.

5.4.4.2.2 Holdover Stability for SMCs

Holdover is defined in GR-1244-CORE (e.g., see Sections 3.6 and 5.2 of that document). In addition to the holdover stability requirements listed below, holdover-related criteria also exist in such areas as triggers for entry into holdover (see [Section 5.4.6](#)) and phase transients during entry into holdover (see [Section 5.4.4.3.3](#)).

R5-141 [923v2] The initial fractional frequency offset of an SMC in holdover shall be less than 0.05 ppm.

R5-142 [924v2] The frequency drift rate of an SMC in holdover shall be less than 5.8×10^{-6} ppm/second.

R5-143 [925v2] The (additional) fractional frequency offset of an SMC in holdover due to varying temperature conditions shall not exceed 4.1 ppm.

The equations to be used in calculating a clock's fractional frequency offset and drift are given in Section 2.7.5 of GR-1244-CORE. For an SMC, the applicable values of the regression time or measurement period ($N\tau_0$) are 60 seconds for calculating the fractional frequency offset, and 2000 seconds for calculating the drift. In addition, a phase transient may be generated at any time during the first 64 seconds after an SMC's entry into holdover (see [Section 5.4.4.3.3](#)), and therefore that time period is excluded from the holdover stability calculations.

Note that the first two requirements listed above are applicable independent of any temperature changes. However, an NE is tested against them at constant temperature so that the components of the fractional frequency offset that they are intended to limit can be separated from any component caused by changing temperatures. The cumulative result of all three of these requirements is to limit the maximum holdover frequency offset to less than 4.6 ppm for the first 24-hour period of holdover under all operating conditions.

5.4.4.2.3 Pull-in/Hold-in for SMCs

Pull-in and hold-in are defined in Section 3.5 of GR-1244-CORE, and the pull-in/hold-in criteria are critical in maintaining the stratum hierarchy because they assure that any clock of a given stratum level will be able to take timing from any other clock of the same or higher stratum level.

R5-144 [253] If a SONET NE with an SMC is timed from an external reference, the NE clock shall pull-in and hold-in to an external reference that is off frequency by ± 4.6 ppm (i.e., from a free-running stratum 3 clock).

- R5-145 [254]** If a SONET NE with an SMC is timed from an OC-N reference, the NE clock shall pull-in and hold-in to an OC-N that is off frequency by ± 20 ppm.
- R5-146 [1180]** A SONET NE with an SMC shall conform to the stratum 3 clock criteria on phase changes during pull-in in Section 5.8 of GR-1244-CORE.

As discussed in GR-1244-CORE, pull-in/hold-in range requirements are applicable when the free-run frequency of the clock is anywhere within its specified accuracy limits. For example, a line-timed SONET NE containing an SMC must be able to pull in and hold to an OC-N reference that is $+20$ ppm off frequency (from nominal), independent of whether the SMC's current free-run frequency is $+20$, 0 or -20 ppm off frequency (from nominal). In addition, it is important to realize that payload integrity is not necessarily guaranteed past frequency offsets of 4.6 ppm (from nominal).

In environments where some NEs with SMCs and some NEs with stratum 3 internal clocks are deployed, possible interoperability issues exist. Specifically, an NE with a stratum 3 clock will not be able to pull-in to a signal from an upstream SMC in free-run. Careful engineering, following the hierarchical rules in GR-436-CORE, is necessary in architectures that mix NEs with SMCs and stratified clocks.

Conformance with pull-in/hold-in requirements is partially verified by testing for wander generation as per [Section 5.4.4.3.2](#) after some "settling time." Note that during this settling time the NE will not necessarily conform with all of the wander generation and transfer criteria. Also note that this requirement is intended to apply in cases where the NE and clock are "warmed up" (e.g., after the frequency of the active reference changes). It is not intended to apply to clock start-up situations such as recovery after a power outage, or after the clock circuit pack is (re)inserted into the NE. That issue (i.e., clock warm-up times) may be discussed in a future issue of GR-253-ILR.

-
- R5-147 [926]** The maximum settling time for an SMC shall be 100 seconds.
-

5.4.4.2.4 Wander Tolerance and Transfer for SMCs and Stratum 3 Clocks

The following criteria define the amount of phase-noise that SMCs and stratum 3 clocks in SONET NEs must tolerate, as well as the minimum amount of phase noise filtering that those clocks need to provide. These requirements are based on work done in T1X1.3 that was intended to (for example) control the amount of jitter on transported DS3 payloads during worst-case external DS1 timing reference phase transients, and limit the severity of pointer adjustment bursts.


-
- R5-148 [1111v2]** An SMC or stratum 3 clock in a SONET NE shall tolerate (i.e., give no indication of improper operation, in addition to meeting **R5-149 [255v3]**) any arbitrary input reference signal having wander TDEV characteristics less than or equal to the input mask in [Figure 5-14](#) (for an OC-N reference signal) or [Figure 5-15](#) (for an external DS1 reference).

- R5-149 [255v3]** OC-N/OC-M and STS-N/STS-M electrical outputs, when referenced to an external timing signal that meets the (input) wander TDEV mask in [Figure 5-15](#), shall be less than or equal to the output wander TDEV mask given in [Figure 5-14](#).
- R5-150 [256v3]** OC-N/OC-M and STS-N/STS-M electrical outputs, when referenced to an OC-N timing signal that meets the input wander TDEV mask in [Figure 5-14](#), shall be less than or equal to the output wander TDEV mask in [Figure 5-14](#).

Note that the masks in [Figures 5-14](#) and [5-15](#) allow for 2% gain (or peaking) in the passband of the clock. While this is slightly different than the 0.2 dB gain that was allowed by the criteria in Issue 3 of this document, it is consistent with the criteria in GR-1244-CORE that apply to other types of clocks, and also with the specifications in ANSI T1.101. Also note that if an NE does not appear to meet **R5-149 [255v3]** or **R5-150 [256v3]** when it is tested using an input reference signal whose TDEV characteristics are equal to the appropriate input mask, the results could indicate a nonconformance to either the wander tolerance or wander transfer requirements. In such cases it may be possible to perform additional tests using input signals with TDEV characteristics that are well below the input mask (i.e., that the NE should not have any difficulty tolerating) to determine the phase-noise filtering characteristics of the clock, and thus whether the problem is a wander tolerance problem or a wander transfer problem.


One result of the requirements in this section is that a stratum 3 clock contained in a SONET NE is not required (but is certainly allowed) to tolerate the same level in wander on its input reference signals as a stratum 3 clock contained in non-SONET equipment covered by the wander tolerance criteria in GR-1244-CORE. The reason for this is that it is assumed that SONET NEs that contain stratum 3 clocks (or SMCs) will *not* be used in applications where they would need to receive timing from a stratum 3 BITS clock that in turn receives timing directly from the non-SONET synchronization distribution network. That is, it is assumed that a stratum 3E or 2 clock (either of which would provide significantly more filtering than a stratum 3 clock) will be present between the non-SONET synchronization distribution network and any SONET NEs that require timing that is traceable from or through that network.

Figure 5-14 OC-N Input and Output Wander Time Deviation



Integration Time, τ (seconds)	Input TDEV (nanoseconds)	Output TDEV (nanoseconds)
$\tau < 0.1$	N/A	N/A
$0.1 \leq \tau \leq 1.73$	10	10.2
$1.73 < \tau \leq 30$	$5.77 \times \tau$	$5.88 \times \tau$
$30 < \tau \leq 1000$	$31.623 \times \tau^{0.5}$	$32.26 \times \tau^{0.5}$
$1000 < \tau$	N/A	N/A

Figure 5-15 Time Deviation of Filtered Network Input to SONET NEs



Integration Time, τ (seconds)	TDEV (nanoseconds)
$\tau < 0.1$	N/A
$0.1 \leq \tau \leq 3$	17
$3 < \tau \leq 30$	$5.67 \times \tau$
$30 < \tau \leq 1000$	$31.62 \times \tau^{0.5}$
$1000 < \tau$	N/A

5.4.4.3 All SONET Stratum Clocks and SMCs

The following criteria apply to the clocks in all SONET NEs that contain LTE, independent of whether the clocks are stratum clocks or SMCs.

5.4.4.3.1 Clock Hardware

In general, large amounts of traffic are dependant upon the availability and quality of the clock. Therefore the clock in a SONET NE should be very reliable. To provide this level of reliability, the clock is required to be effectively duplicated by methods which allow corrective maintenance actions to be performed on failed elements without affecting the in-service elements.

-
- R5-151 [264v3]** A stratum clock or SMC in a SONET NE shall meet the duplication of equipment criteria in Section 3.3 of GR-1244-CORE (including the referenced criteria in Section 5.6 of GR-499-CORE).
-

5.4.4.3.2 Wander Generation

Wander generation is the process whereby wander appears at the output of a clock in the absence of input wander. Additional information about the intent of wander generation criteria, the input test signal, and wander generation test procedures can be found in Section 5.3 of GR-1244-CORE.


-
- R5-152 [257v2]** OC-N/OC-M and STS-N/STS-M electrical outputs shall meet the MTIE wander mask in [Figure 5-16](#) when timed with a wander-free reference having bandlimited white-noise phase modulation as described below.

- R5-153 [258v2]** OC-N/OC-M and STS-N/STS-M electrical outputs shall meet the TDEV wander mask in [Figure 5-17](#) when timed with a wander-free reference having bandlimited white-noise phase modulation as described below.
-

Conformance to these requirements is tested by providing an external or OC-N reference with bandlimited (with 3-dB cutoffs at 10 Hz and 150 Hz) white-noise phase modulation of 1 μ s peak-to-peak. (Also see GR-1244-CORE.)


Note that although the input test signal defined above is used in the jitter tolerance requirement that applies to externally timed SONET NEs (see **R5-171 [1015]**), the amount of jitter on that signal is greater than a line-timed NE (or through-timed ADM) is required to tolerate (see **R5-173 [1017v2]**). Thus, it is possible that an NE with “conforming” jitter tolerance characteristics will not accept the wander generation input test signal, and therefore will not be able to be tested using that signal. Alternate methods for verifying that such an NE *conforms* to the wander generation requirements are FFS. On the other hand, it may be possible to determine that an NE *does not conform* to those requirements by using an input test signal containing less than 1 μ s of white-noise phase modulation (e.g., a “clean” signal or a signal containing jitter at the maximum level tolerated by the NE). Specifically, if the amount of wander generated by the NE does not meet the applicable MTIE or TDEV masks when the input OC-N reference signal contains less than 1 μ s of white-noise phase modulation, then it would also not meet those masks if the NE could be tested using the specified input test signal.

Figure 5-16 MTIE for SONET Clocks



Observation Time, S (seconds)	MTIE (nanoseconds)
$S < 0.1$	N/A
$0.1 \leq S \leq 1$	20
$1 < S \leq 10$	$20 \times S^{0.48}$
$10 < S$	60

Figure 5-17 Time Deviation for SONET Clocks



Integration Time, τ (seconds)	TDEV (nanoseconds)
$\tau < 0.1$	N/A
$0.1 \leq \tau \leq 2.5$	$3.2 \times \tau^{-0.5}$
$2.5 < \tau \leq 40$	2
$40 < \tau \leq 1000$	$0.32 \times \tau^{0.5}$
$1000 < \tau$	10

5.4.4.3.3 Phase Transients

As discussed in Section 5.6 of GR-1244-CORE, phase transients are typically traceable to some causal event or activity, and need to be controlled in order to minimize the impact on downstream NEs. In some cases the phase transient criteria in GR-1244-CORE, which apply to clocks in general applications, are also sufficient for clocks in SONET NEs (see **R5-138 [247v2]**). However, in other cases the somewhat more stringent criteria provided in this section are needed to control jitter on asynchronously mapped payloads and the generation of pointer adjustment bursts. In addition, this section defines a SONET-specific synchronization rearrangement operation that may occur at a line-timed NE (or through-timed ADM). That operation is a switch between the incoming signals at the active SONET “interface”,²³ and combined with the operations listed in GR-1244-CORE results in the following possible synchronization rearrangement operations at a SONET NE:

- Manual switching between clock hardware units
- Automatic switching between clock hardware units
- Removal of a clock hardware unit, either active (in which case automatic switching between units would occur) or standby
- Manual timing reference switching
- Automatic timing reference switching
- Switching between working line 1 and the protection line at the active reference OC-N “interface” (for NEs that support line APS)
- Automatic clock diagnostics
- Entry into holdover
- Recovery from holdover
- A phase transient of up to 1000 ns on the active reference.

As indicated in **R5-158 [1182]**, the characteristics of the input phase transient in the last bullet item above depend on the type of reference signal (i.e., external DS1 or OC-N). In addition, those characteristics are consistent with the characteristics of the input phase transient that the clock is required to tolerate (see **R5-169 [272v3]**).

R5-154 [259v4] For SONET NEs that contain stratum 3 clocks or SMCs, the MTIE of OC-N/OC-M and STS-N/STS-M electrical outputs during any of the synchronization rearrangement operations listed above except a phase transient on the active reference shall be no greater than the “requirement” mask in [Figure 5-18](#).


²³Note that although the same phase transient criteria apply, when an NE that is timed from an OC-N “interface” stops taking timing from the incoming signal on one line (e.g., working line 1) and begins to take timing from the signal on another line at the same “interface” (e.g., the protection line), it is not considered to have performed a timing reference switch (see [Section 5.4.3.2](#)).

As noted above, it is desirable that phase transients from SONET NEs be small to minimize the number of STS pointer adjustments generated. The following objective (which is derived from the requirement in GR-1244-CORE that applies to SONET NEs containing stratum 2 or 3E clocks) would allow for phase hits that cause no more than one STS pointer adjustment.

-
- O5-155 [260v2]** For SONET NEs that contain stratum 3 clocks or SMCs, the MTIE of the SONET outputs during any of the synchronization rearrangement operations listed above except a phase transient on the active reference should be no greater than the “objective” mask in [Figure 5-18](#).
-

Consistent with statements in GR-1244-CORE, and ANSI T1.101 and T1.105.09, it is assumed that for a SONET NE, the output phase transient caused by any synchronization rearrangement operation other than a transient on the active reference will occur within 64 seconds if the NE is equipped with stratum 3 or 3E clocks or SMCs, or within 5000 seconds if the NE contains stratum 2 clocks. Based on this, the masks in [Figure 5-18](#) extend to a maximum observation time of 64 seconds, and those in GR-1244-CORE extend to 64 or 5000 seconds. While this is different than the 280-second maximum observation time that was specified in the corresponding figure in previous issues of this document and currently appears in several related standards documents, it is consistent with the “Phase Transient for Entry into Holdover” figure that appeared in Section 5.4.4.2.2 of GR-253-CORE Issue 3 (which now appears as [Figure 5-19](#)). In addition, the difference in the maximum observation times is not expected to have any impact on the conformance or nonconformance of a product to **R5-154 [259v4]** or **O5-155 [260v2]**.

Figure 5-18 MTIE for Phase Transients from SONET Clocks



Observation Time, S (seconds)	“Requirement” MTIE (nanoseconds)	“Objective” MTIE (nanoseconds)
$S < 0.0140$	N/A	N/A
$0.014 \leq S < 0.16$	$7.6 + 885 \times S$	$7.6 + 885 \times S$
$0.16 \leq S < 0.5$	$7.6 + 885 \times S$	150
$0.5 \leq S < 2.33$	$300 + 300 \times S$	150
$2.33 \leq S < 64$	1000	150
$64 \leq S$	N/A	N/A


As discussed in GR-1244-CORE, in determining an NE’s conformance to output phase transient criteria such as **R5-154 [259v4]** and **O5-155 [260v2]** it is necessary to distinguish between the phase transient caused by the rearrangement operation and any additional phase changes caused by differences in the operating frequencies of the NE’s clocks before and after the operation. For most types of rearrangement operations, this can be accomplished by using test configurations in which the NE’s clocks operate at essentially the same frequency before and after the event. However, that is not possible in the case of a clock’s entry into holdover. In that case, the clock may display an initial fractional frequency offset relative to the pre-holdover reference frequency, and that offset may make it difficult to interpret the measured phase data. One method of possibly simplifying this process is to restate the phase transient MTIE criteria as shown below and in [Figure 5-19](#). Note that unlike the corresponding masks in GR-1244-CORE, the masks in [Figure 5-19](#)

only include allowances for frequency offsets after they reach the maximum MTIE values shown in Figure 5-18 (i.e., 1000 or 150 ns). This was done for consistency with the issues of several related standards [i.e., T1.101, T1.105.09, option 2 in ITU-T Rec. G.813, *Timing characteristics of SDH equipment slave clocks (SEC)*] that were current when this document was being prepared, and also with previous issues of this document.

R5-156 [922v2] For SONET NEs that contain stratum 3 clocks or SMCs, the MTIE of the SONET outputs during entry into holdover shall be no greater than the “requirement” mask in Figure 5-19.

O5-157 [1181] For SONET NEs that contain stratum 3 clocks or SMCs, the MTIE of the SONET outputs during entry into holdover should be no greater than the “objective” mask in Figure 5-19.

Figure 5-19 Phase-Transient for Stratum 3 Clock and SMC Entry into Holdover




Observation Time, S (seconds)	“Requirement” MTIE (nanoseconds)	“Objective” MTIE (nanoseconds)
$S < 0.014$	N/A	N/A
$0.014 \leq S < 0.16$	$7.6 + 885 \times S$	$7.6 + 885 \times S$
$0.16 \leq S < 0.5$	$7.6 + 885 \times S$	$150 + 50 \times S$
$0.5 \leq S < 2.33$	$300 + 300 \times S$	$150 + 50 \times S$
$2.33 \leq S < 64$	$884 + 50 \times S$	$150 + 50 \times S$
$64 \leq S$	N/A	N/A

R5-158 [1182] For all SONET NEs containing LTE, the MTIE during any phase transient on an active DS1 or OC-N reference with an MTIE up to the corresponding DS1 or OC-N input phase transient mask in Figure 5-20 shall not exceed the output mask in the same figure.

O5-159 [1183] For all SONET NEs containing LTE, the MTIE during any phase transient on an active OC-N reference with an MTIE up to the “DS1 Input” phase transient mask in Figure 5-20 should not exceed the output mask in the same figure.

Note that in the situation covered in **R5-158 [1182]**, **O5-159 [1183]** and Figure 5-20, a relatively small amount of overshoot (i.e., up to 2%, consistent with the maximum phase gain allowed by the wander transfer requirement **R5-150 [256v3]** and the corresponding criteria in GR-1244-CORE) in a clock’s phase transient response is considered acceptable. Thus, the measured MTIE could be slightly greater than 1000 ns.

Figure 5-20 MTIE Masks for Input/Output Phase Transients



Observation Time, S (seconds)	DS1 Input MTIE (ns)	OC-N Input MTIE (ns)	Output MTIE (ns)
$S < 0.001326$	N/A	N/A	N/A
$0.001326 \leq S < 0.014$	$61000 \times S$	N/A	N/A
$0.014 \leq S < 0.0164$	$61000 \times S$	$7.6 + 885 \times S$	$7.6 + 885 \times S$
$0.0164 \leq S < 0.5$	1000	$7.6 + 885 \times S$	$7.6 + 885 \times S$
$0.5 \leq S < 2.33$	1000	$300 + 300 \times S$	$300 + 300 \times S$
$2.33 \leq S$	1000	1000	1020

5.4.4.3.4 Reference Validation and Transitions From Self-Timing to Normal Mode

If an NE is in a self-timing mode (i.e., holdover or free-run) and a “good” reference becomes available, the user would normally want the NE to use that reference as the synchronization source. However, in some trouble-shooting and maintenance situations it may be appropriate for an NE with a stratum clock or SMC to continue self-timing. Therefore, the following criteria apply.

R5-160 [265v2] Unless it has been inhibited by the user (see **CR5-161 [266]**), recovery from self-timing (holdover or free-run) shall be automatic.

CR5-161 [266] An NE with a stratum clock or SMC may be required to provide the ability to inhibit automatic recovery from self-timing.

R5-162 [267v3] An NE with a stratum clock or SMC shall conform to the reference validation criteria in Section 3.7 of GR-1244-CORE. An SMC shall conform to the criteria for a stratum 3 clock.

R5-163 [268] Automatic restoration from the free-run mode shall occur within two seconds of the presence of a validated reference signal.

In addition, a limit on the maximum rate of change of frequency during an SMC’s or stratum clock’s recovery from holdover is necessary to avoid the creation of excessive amounts of jitter on the SONET asynchronous payloads.

R5-164 [930v2] The maximum rate of frequency change during holdover recovery shall be less than 2.9 ppm/second.

Information regarding the method used to calculate an NE’s rate of frequency change during holdover recovery is provided in ANSI T1.105.09. That method is the same as the method given in Section 2.7.5 of GR-1244-CORE for calculating drift, with a regression time or measurement period ($N\tau_0$) equal to 1 second.

5.4.4.3.5 Jitter and Errors During Synchronization Rearrangement Operations

In general, the MTIE criteria in [Section 5.4.4.3.3](#) could be interpreted to allow nearly instantaneous phase jumps of up to 20 ns. This is clearly undesirable from a jitter point of view. The following objective addresses this concern.

O5-165 [928v3] The SONET outputs of an NE should meet the jitter generation requirement in [Section 5.6.2.3.6](#) during the synchronization rearrangement activities listed in [Section 5.4.4.3.3](#), and also during entry into and recovery from free-run.

R5-166 [261v2] Except for clock hardware protection switching, the synchronization rearrangement activities listed in [Section 5.4.4.3.3](#) shall cause no errors on payload traffic.

Note that **R5-166 [261v2]** does not apply for the case of traffic being carried on a failed reference or line. That traffic will either be lost (e.g., in a system with linear APS where both the working and protection lines fail) or may be temporarily interrupted while a protection switch is being completed (e.g., in a UPSR NE). Similarly, the requirement does not apply for traffic that is normally carried on the failed reference, was restored via a protection switch when the reference failed (and the NE entered holdover), and is subject to a revertive protection switch after the reference is restored (and the NE recovers from holdover).

O5-167 [262v2] Clock hardware protection switching should cause no errors on payload traffic.

5.4.4.3.6 Input Tolerance

In general, GR-1244-CORE allows suppliers flexibility in determining when an NE will consider a reference failed (see **R1244-30**). The NE may consider a reference failed as soon as it detects an LOS, AIS, Out Of Frame (OOF) or LOF defect, or it may wait several seconds up to the point when a failure is declared, to see if the defect persists. If the NE does wait to see if the defect persists, its output synchronization performance must not be degraded.

R5-168 [271v3] For interruptions (i.e., periods of no pulses, AIS or “bad” framing, not phase or frequency transients) of reference signals that do not cause reference switches or switches between lines (at an NE that supports line APS), the output criteria of [Figure 5-16](#) shall be met.

R5-169 [272v3] The NE shall tolerate phase transients on external DS1 and OC-N reference signals with MTIEs up to the corresponding DS1 and OC-N input phase transient masks in [Figure 5-20](#) (i.e., with magnitudes and slopes as defined in ANSI T1.101 for those types of signals).

Based on **R5-169 [272v3]**, a SONET NE that is timed using an OC-N signal is not required to be able to tolerate a phase transient on that signal that is equivalent to the worst-case transient specified for external DS1 signals. On the other hand, most SONET NEs are expected to be able to tolerate that type of input phase transient on any type of reference (i.e., they are not expected to be designed to be less tolerant to phase transients on OC-N signals than to phase transients on external DS1 signals), and therefore they can be expected to meet the following objective.

O5-170 [1184] An NE that is timed from an incoming OC-N signal should tolerate any phase transient on that signal with an MTIE up to the DS1 input phase transient mask in [Figure 5-20](#) (i.e., with magnitudes and slopes as defined in ANSI T1.101 for external DS1 signals).

R5-171 [1015] Clocks synchronized to an external DS1 timing signal shall tolerate, as a minimum, the jitter specified for the input test signal in the wander generation requirements in [Section 5.4.4.3.2](#).

CR5-172 [1016] Clocks synchronized to an external DS1 timing signal may be required to tolerate, as a minimum, input jitter applied according to the mask in Figure 7-1 of GR-499-CORE.

R5-173 [1017v2] Clocks that are timed from an incoming OC-N signal shall tolerate (for timing purposes) jitter applied according to the Category II jitter tolerance requirement mask in Figure 5-29.

The term “tolerate” used in the previous criteria is defined to mean no errors on payload signals, no indication of improper operation (i.e., alarms), no reference rejection, and remaining frequency-locked to the reference so that the output phase variations, relative to the input reference, are bounded. In addition, in the case of the input phase transient criteria (**R5-169 [272v3]** and **O5-170 [1184]**) it means that the amplitude of the output phase change is essentially equal to the amplitude of the input phase change.

5.4.4.3.7 Phase Build-Out

The phase build-out function is defined in Section 5.7 of GR-1244-CORE. As indicated in that section, it is required to be supported by stratum 2 and 3E clocks, and is required to not be supported by lower quality clocks.

R5-174 [1185] A SONET NE shall conform to the applicable (based on the quality of its internal clocks) criteria on phase build-out in Section 5.7 of GR-1244-CORE. An NE containing SMCs shall meet the criteria for a stratum 3 clock.

5.4.5 Timing Distribution

This section describes SONET NE criteria for SONET-based network synchronization distribution. SONET-based network synchronization distribution can occur two ways, either with DS1 signals derived from a terminating OC-N,²⁴ or with retimed traffic DS1 signals (i.e., DS1s carried in the SONET payload and then retimed). The derived DS1 is the preferred method for timing distribution, as retimed payload DS1s are subject to the possibility of slips at the retiming buffer and should only be used for special applications. (For example, for remote locations it may be preferable to provide a retimed signal rather than a separate timing-only signal that requires extra facilities.) In addition, DS1s that are carried on SONET and are not retimed are not recommended for network synchronization distribution to BITS clocks, because those signals may not meet the ANSI T1.101 synchronization interface specifications.

²⁴Although some NEs may support the capability to use an OC-M “interface” as a source for a derived DS1, that capability is not required in this document.

5.4.5.1 Timing Distribution on Derived DS1 Signals

In order for SONET-based network synchronization distribution to occur, BITS clocks must be able to receive synchronization signals from SONET NEs. Since BITS clocks accept DS1 signals as input references, SONET NEs that contain LTE need to be able to generate timing signals in the DS1 format.

R5-175 [273v4] A SONET NE that contains LTE shall have the capability to supply two DS1 timing reference signals. The NE shall be capable of deriving both of these DS1s from a single line-side OC-N “interface” (see Figure 5-21) and, if more than one OC-N “interface” is supported, of deriving each DS1 from a different OC-N “interface” (see Figure 5-22). As a minimum, the derived DS1 signals shall be in the Superframe format, and shall meet the electrical specifications (including those related to cable lengths) and pulse-density criteria in GR-499-CORE.

The reason for specifying that both of the above methods of deriving DS1 signals must be supported is that different applications have different uses for the derived DS1 signals. For example, if a SONET ring is being used for interoffice synchronization distribution and contains externally timed NEs that support DS1 synchronization status messages, then it would generally be beneficial to have each DS1 signal derived from a different OC-N “interface”. Conversely, if the NEs in a SONET ring are line-timed (in which case the derived DS1 signals are presumably being used for some purpose other than timing of a local BITS) or do not support DS1 synchronization status messages, then it would generally be beneficial to have both DS1 signals derived from the same OC-N “interface”.

Figure 5-21 DS1 Timing References Derived from a Single OC-N “Interface” Example

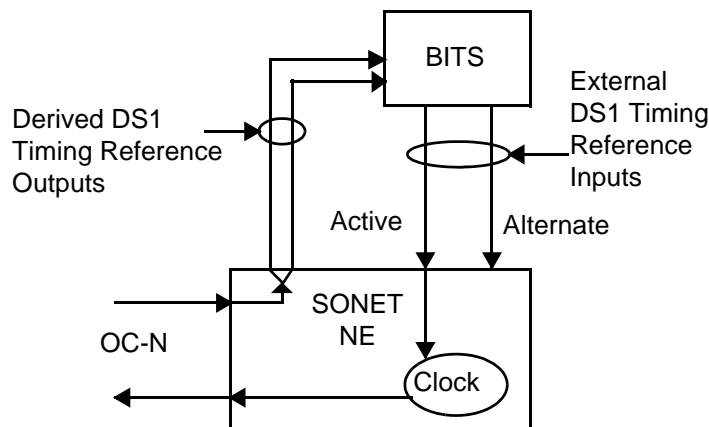
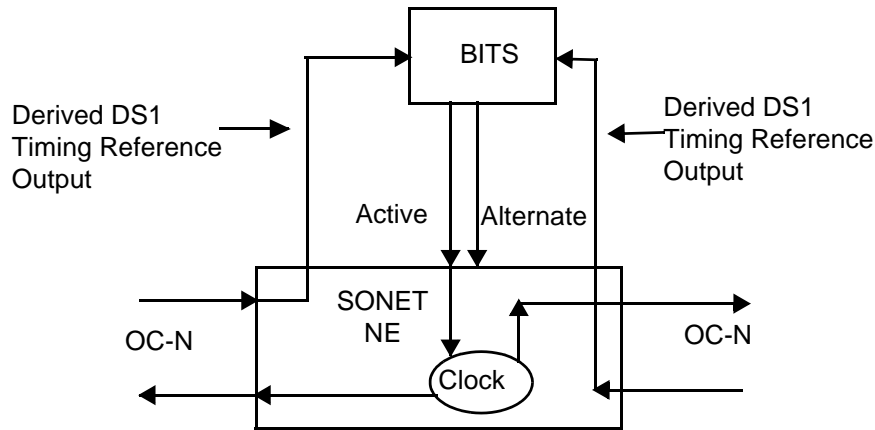


Figure 5-22 DS1 Timing References Derived from Different OC-N “Interfaces” Example



Note that in the case illustrated in [Figure 5-21](#), if the NE supports line APS it may either (as a default) derive both DS1s from the OC-N signal on a single line (i.e., working line 1 or the protection line), or derive one DS1 from an OC-N on working line 1 and the other DS1 from the OC-N on the protection line. In any case, just as the reliability of SONET is enhanced by the ability to smoothly switch traffic to a protection line if a problem occurs on a working line, a similar derived-DS1-source protection switching capability can be used to improve the robustness of the synchronization distribution network.

- CR5-176 [274v3]** An NE that supports line APS may be required to switch the source of timing for a derived DS1 to the OC-N on the protection line when the OC-N on working line 1 is considered failed or unavailable as a derived DS1 source according to [R5-212 \[1022\]](#) (i.e., due to an LOS, LOF, or AIS), and vice versa if one of the DS1s is normally derived from the OC-N on the protection line.
- O5-177 [1113]** If a SONET NE that meets [CR5-176 \[274v3\]](#) normally derives one DS1 from the signal on working line 1 and the other DS1 from the signal on the protection line, then it should use revertive switching for derived-DS1-source switches between those signals.
- O5-178 [1114]** If a SONET NE that meets [CR5-176 \[274v3\]](#) normally derives both DS1s from the signal on one line (i.e., working line 1 or the protection line), then unless it protects traffic and the derived DS1 source together and is using revertive switching for the traffic, it should use nonrevertive switching for derived-DS1-source switches between the working and protection line signals.

In addition to switching the source of a derived DS1 between the working and protection lines at a single OC-N “interface”, in some situations it may be desirable to support switching between separate OC-N “interfaces”. Note that if this capability is supported, it implies that the capability is provided for the user to provision two

(or possibly more) “interfaces” as potential sources for each derived DS1. Also note that if switching between “interfaces” is not desired in a particular application, it would normally be possible for the user to disable it (e.g., by provisioning only a single “interface” as the source for each derived DS1).

CR5-179 [1018v2] An NE that provides more than one OC-N “interface” may be required to support switching of the source of timing for a derived DS1 to a different (secondary) OC-N “interface” when the OC-N signal(s) at the original (primary) “interface” is considered failed or unavailable as a derived DS1 source according to **R5-212 [1022]** (i.e., due to an LOS, LOF, or AIS), or when switching is appropriate based on the received synchronization status messages (see **R5-195 [285v3]**).

R5-180 [1019] If an NE that provides line APS supports switching between OC-N “interfaces” as the source of timing for a derived DS1, then it shall conform to **CR5-176 [274v3]**. In addition, a switch between “interfaces” in response to an LOS, LOF or AIS shall occur only if the OC-N signals on both working line 1 and the protection line have failed.

As discussed in [Section 5.4.5.2.1](#), switching the source of a derived DS1 from one OC-N “interface” to another at an externally timed NE is not recommended. Thus, any NE that is provisioned to support derived DS1 source switching is expected to also be provisioned to be line (or possibly through) timed. In addition, it is generally assumed that the same “interfaces” will be provisioned as potential timing references and potential derived DS1 sources, and that it is desirable for an “interface” that is currently being used as a timing reference to also be active as a derived DS1 source. To maintain this parallel use of OC-N “interfaces”, the following requirements apply.

R5-181 [1020] A line-timed NE that is provisioned to support switching between OC-N “interfaces” as the source of timing for a derived DS1, and to (as a default) derive all of its active DS1s from the same OC-N “interface”, shall use the same type of switching (i.e., nonrevertive or revertive) as it uses for timing reference switching.

R5-182 [1021] A through-timed ADM that is provisioned to support switching between OC-N “interfaces” as the source of timing for a derived DS1, and to (as a default) derive each DS1 from a different OC-N “interface”, shall use revertive switching.

Note that if an NE’s timing reference and derived DS1 source lists are identical, then maintaining the same active/inactive status for those two uses of a particular OC-N “interface” could be achieved by simply linking the derived DS1 source switching feature to the NE’s timing reference and mode switching (i.e., through-timing to line-timing) features. Possible methods of automatically ensuring identical timing reference and derived DS1 source lists have been examined (see GR-253-ILR, Issue ID 253-51); however, they were not considered feasible. In addition, in some applications the user may want the lists to be different, at least temporarily. Therefore, it is left to the user to provision each list appropriately.

Also note that since switching between OC-N “interfaces” as the source for a derived DS1 is not recommended at externally timed NEs (see [Section 5.4.5.2.1](#)), no criteria have been included in this document to indicate if such switching (if supported) should be revertive or nonrevertive. Similarly, no revertive/nonrevertive switching criteria have been included to cover situations where parallel use of OC-N interfaces for timing and deriving DS1s is not desired (e.g., where a line-timed NE has been provisioned so that it normally derives its DS1s from the incoming signals at two different OC-N “interfaces”, or where a through-timed ADM has been provisioned to normally derive both of its DS1s from a single OC-N interface).

In order to maintain the parallel use of an NE’s OC-N “interfaces” for timing reference and derived DS1 source purposes in cases where the commands related to timing references defined in [Section 5.4.6](#) and GR-1244-CORE are used, it would be necessary for the NE to support the equivalent commands (e.g., a Manual Derived DS1 Source Switch command) for switching or locking out derived DS1 sources.

CR5-183 [1052v2] A SONET NE that supports the capability to switch between OC-N “interfaces” as the source for a derived DS1 may be required to support commands for derived DS1 source switching that are equivalent to the corresponding commands related to timing reference switching (see [Section 5.4.6](#)) that it supports.

CR5-184 [1115] A SONET NE that supports the capability to switch between the incoming signals on working line 1 and the protection line (at the same OC-N “interface”) as the source for a derived DS1 may be required to support commands for derived DS1 source switching that are equivalent to any corresponding commands related to timing reference switching (see [Section 5.4.6](#)) that it supports.

R5-185 [1053v2] If a SONET NE supports one or more commands related to derived DS1 source switching, the effects of those commands shall be functionally equivalent to the effects of the corresponding commands related to timing reference switching.

For example, a Manual Derived DS1 Source Switch command would need to be denied if it would cause a switch to a failed source or a source with a lower synchronization status message, while the effect of a Lockout a Derived DS1 Source command or a Forced Derived DS1 Source Switch command would be to effectively remove one or all but one source from the provisioned source list until, for example, the corresponding clear command is entered.

O5-186 [276] The derived DS1 should be a framed all-ones signal.

CR5-187 [277] The NE may be required to provide the capability for the user to provision the derived DS1 in the ESF format (in addition to the required Superframe format).

R5-188 [278] The SONET NE shall be capable of supporting all of its timing modes when providing derived DS1s from an OC-N.

R5-189 [279v3] Unless a switch to a different (available) OC-N signal is performed (see **CR5-176** [274v3], **CR5-179** [1018v2], **R5-180** [1019] and **R5-195** [285v3]), DS1 AIS (i.e., unframed all-ones) shall be inserted into the derived DS1 when the OC-N

signal is determined to be failed or unavailable as a derived DS1 source according to **R5-212 [1022]** (i.e., due to an LOS, LOF, or AIS) or **R5-213 [315v3]** (i.e., due to the lack of a validated synchronization status message). The DS1 AIS shall be generated no later than the declaration of the failure (see [Section 6.2.1](#)).

R5-190 [280v3] If an LOS, LOF or AIS causes AIS to be generated on the derived DS1, automatic restoration of the derived DS1 signal (i.e., the removal of DS1 AIS) shall occur within 12.5 seconds of the termination (and continued absence) of the defect on the derived DS1 source.

Note that the 12.5-second maximum restoration time specified in **R5-190 [280v3]** is consistent with the limit that appeared in Issue 3 of this document (i.e., 10.5 seconds to clear the LOS, LOF or AIS failure, plus 2 seconds), while also covering the case where the LOS, LOF or AIS defect does not persist long enough to cause the corresponding failure to be declared. Also note that in cases where the generation or removal of DS1 AIS is based on changes in the received synchronization status messages, the applicable criteria related to the timing of those actions appear in [Section 5.4.7.3](#).

In order to interface with BITS clocks and Customer Premise Equipment (CPE), the derived DS1 signals have several performance criteria. It is expected that the derived DS1 from the SONET NE will meet ANSI T1.101 specifications. Some of the following performance criteria are more stringent than those in T1.101. Conforming with the following requirements is expected to ensure that the derived DS1 signal, which may be generated at the end of a chain of clocks, meets the T1.101 specifications.

R5-191 [281v2] The MTIE of the derived DS1 shall be less than 50 ns for observation times longer than 0.1 second.

R5-192 [282] The TDEV of the derived DS1 shall be below the mask in [Figure 5-23](#).

The conformance of an NE to these requirements is measured using a 10-Hz first-order low-pass filter, relative to an ideal, jitter-free and wander-free source that is also providing timing to the OC-N. Also, it should be noted that the MTIE requirement does not specifically prohibit filtering incoming jitter from the OC-N, while the intent of the TDEV requirement is to force any wander that is generated on the derived DS1 to be at higher frequencies (lower integration times in terms of TDEV) so that it can be filtered by downstream BITS clocks.


R5-193 [283] The jitter on the derived DS1 shall be less than $1.0 U_{Ipp}$.

In this case, conformance testing is performed using a jitter-free OC-N source signal that meets the OC-N output wander TDEV mask in [Figure 5-14](#) in [Section 5.4.4.2.4](#). [For the purposes of this testing, an OC-N source signal can be considered to be “jitter-free” if the energy at observation times shorter than 0.1 s is such that the measured TDEV (in ns) is less than $100 \times \tau$ (with τ in seconds).]

R5-194 [284v2] The derived DS1 shall meet the MTIE requirement for DS1 phase transients in ANSI T1.101 during rearrangements (e.g., switching the derived DS1 source between working line 1 and the protection line in a system that supports line APS, or between “east” and “west” OC-Ns as per [Section 5.4.5.2.1](#)). For observation periods up to 280 seconds, a phase transient on a derived DS1 shall not exceed a magnitude of 1 μ s with a slope of 81 ns for a measurement period of 1.326 ms.

This measurement is made with a 100-Hz first-order low-pass filter, relative to an ideal, jitter-free and wander-free source.

Figure 5-23 Time Deviation for Derived DS1 Signals



Integration Time, τ (seconds)	TDEV (nanoseconds)
$\tau < 0.1$	N/A
$0.1 \leq \tau \leq 2.5$	$3.2 \times \tau^{-0.5}$
$2.5 < \tau \leq 40$	2
$40 < \tau \leq 1000$	$0.32 \times \tau^{0.5}$
$1000 < \tau$	10

5.4.5.2 Synchronization Status Messages for Derived DS1 Signals

This section describes and contains the criteria specific to the use of synchronization status messages on derived DS1 signals. This is in addition to the information and criteria in [Section 5.4.7](#), some of which is general and some of which is specific to the use of synchronization status messages on timing reference signals.

5.4.5.2.1 Switching

Ideally, the derived DS1 signals should always be referenced to the OC-N with the highest traceability as indicated by synchronization status messages. However, switching between OC-N “interfaces” used to derive a DS1 at an externally timed NE (i.e., for use in the interoffice timing distribution network) would cause the maintenance and administration of the network, especially the avoidance of timing loop creation, to be extremely difficult. Conversely, switching between OC-N “interfaces” used to derive a DS1 at line-timed NEs and through-timed ADMs will not cause timing loops. Therefore, automatic switching of the OC-N “interfaces” used to derive the DS1 signals (in order to achieve the highest traceability) is viable only at line-timed NEs and through-timed ADMs.

R5-195 [285v3] A line-timed NE or through-timed ADM shall automatically select (from the OC-N “interfaces” provisioned as possible sources for each derived DS1) the OC-N “interface” with the highest quality synchronization status message as the source for that derived DS1 signal.

Note that the preceding requirement covers source switching between OC-N “interfaces”, but does not include switching between the signals on the working and protection lines at a single OC-N “interface”. In general, working/protection line switching based on synchronization status messages is not required to be supported (although it may be provided by some NEs). The reason for this is that any NE that supports line APS is required to transmit the same synchronization status messages on working line 1 and the protection line (at a single “interface”) in all situations. Thus, the messages that are received by the far-end NE on those two lines can be expected to be identical, and it is considered sufficient for that NE to monitor the incoming messages on only one of those lines (e.g., the line whose signal is currently being used as the derived DS1 source).

5.4.5.2.2 Message Translation

Two modes of translating synchronization status messages from the terminating OC-N to the derived DS1 outputs are defined for SONET NEs. These are:

1. “Threshold AIS generation” mode, to be used when the receiving BITS clocks do not have synchronization status messaging capabilities
2. “Message pass-through” mode, to be used when BITS clocks have synchronization status messaging capabilities.

“Threshold AIS generation” mode is similar to the functionality provided by NEs that do not support synchronization status messaging. Such NEs generate an AIS on the derived DS1 when the OC-N becomes unacceptable as a reference (e.g., during an LOS or AIS on the OC-N). Since BITS clocks recognize AIS as an unacceptable reference signal, the generation of AIS ensures that the BITS does not take synchronization from an OC-N that has been disqualified. With the “threshold AIS generation” mode, OC-N signals that degrade in terms of synchronization traceability also cause AIS to be generated, ensuring that the reference is rejected by the BITS. This mode allows the derived DS1 signal to be used as a synchronization reference in the cases where synchronization status messaging in the DS1 ESF format is not supported by the BITS clock.

“Message pass-through” mode is appropriate in cases where the BITS clock supports synchronization status messaging. In this mode, the message on the OC-N is translated onto the derived DS1 and passed to the BITS. The BITS is then the element that decides which reference is appropriate to use, based upon the synchronization status message. In order to support this mode, the DS1 signal must be in the ESF format.

Some service providers are expected to upgrade their BITS clocks to support synchronization status messaging and therefore require the SONET NE to support the “message pass-through” mode. Other service providers may not upgrade their BITS clocks and consequently do not require this capability.

R5-196 [286] A SONET NE shall support the “threshold AIS generation” mode.

CR5-197 [287v3] A SONET NE may be required by some service providers to support the “message pass-through” mode (i.e., it may be required to support the generation of synchronization status messages on the derived DS1).

R5-198 [288] If an NE supports both message translation modes, the format of the derived DS1 signal shall indicate which mode is used. If the derived DS1 is in the ESF format, the “message pass-through” mode shall be used. If the derived DS1 is in the SF format, the “threshold AIS generation” mode shall be used.

The purpose of the previous requirement is to minimize the amount of provisioning that is necessary. However, based on the requirement, if an ESF-formatted, derived DS1 signal is input to a BITS clock that does not support synchronization status messages, the BITS may select an unsuitable reference (e.g., a reference with a DUS message) to be active. Therefore, it is strongly recommended that the network provider use a derived DS1 in the ESF format only when the BITS supports synchronization status messages.

In “threshold AIS generation” mode, the following additional requirements apply.

R5-199 [289] In “threshold AIS generation” mode, the NE shall insert AIS into the derived DS1 output when the synchronization status message in the OC-N signal that is being used as a reference for that derived DS1 is at or below a user-selectable quality level. The default for the threshold shall be quality level 7, “SMC Traceable.”

In “message pass-through” mode, the following additional requirements apply.

-
- R5-200 [291]** In “message pass-through” mode, the NE shall have the capability to generate the synchronization status messages listed in [Table 5-9](#) in the ESF data link of the derived DS1 signals.
- R5-201 [292v2]** The derived DS1 output shall carry the synchronization status message that corresponds to the message in the OC-N that is being used as the source for that derived DS1.
- R5-202 [294]** The synchronization status message shall be sent continuously in the ESF data link.
-

5.4.5.3 Timing Distribution on Traffic-Carrying DS1 Payload Signals

This section provides criteria for a traffic-carrying DS1 signal that may be used for synchronization distribution to locations remote from the SONET NE. For timing distribution at locations where the SONET NE resides, the derived DS1 described in [Section 5.4.5.1](#) is the preferred method. However, at remote locations it may be preferable to provide a signal retimed by a slip buffer, rather than a separate timing-only signal that would require extra facilities. DS1s that are byte-synchronously mapped may already be retimed by a slip buffer (see [Section 3.4.1.1.1](#)). If they are retimed, then they can be used for timing distribution purposes. However, if they are not retimed, or the NE only supports the asynchronous mapping, the following criteria apply.

-
- CR5-203 [295]** An NE may be required to provide a retiming slip buffer for timing distribution on a traffic-carrying, payload DS1 interface.
-

The slip buffer retimes the output DS1 from the SONET NE’s internal clock (which would be synchronized to an external reference or one of the NE’s terminating OC-Ns). Bits are written into the buffer from the VT SPE, and are read out under the control of the SONET NE’s clock. This buffer effectively smooths any phase transients due to VT pointer adjustments. It should be realized that the timing of the DS1 at the output of the slip buffer will be the same as the timing for the SONET NE. Therefore, it is important that the SONET NE have reliable timing. For example, in a line-timed access ring without synchronization status messages an NE with an SMC may end up in holdover during fiber cuts, which may cause unacceptable performance due to slips at the retiming buffer.

-
- R5-204 [296]** The slip buffer shall be at least 1 frame (125 μ s) plus a minimum of 18 μ s of hysteresis. (More hysteresis is desirable.) When a slip occurs, an entire DS1 frame (i.e., 193 bits, including the framing bit) shall be slipped unless the DS1 is byte-synchronously mapped and the VT PTE is generating a new DS1 framing pattern. If a new DS1 framing pattern is being generated, then only the 192 data bits shall be slipped and the framing pattern shall not be affected.
-

Note that whenever the buffer slips for an asynchronously mapped DS1, the framing pattern is interrupted. This interruption may cause downstream reframes, but is considered necessary to provide a clear-channel service. These slip buffers are not considered a universal solution to the timing distribution problem, but may be useful in a few specific applications. It is critical that the signal being retimed be traceable to a PRS so that slips do not occur (except during failures).

-
- R5-205 [297]** If a retiming slip buffer is provided, the NE shall accumulate slip counts as performance monitoring data according to the criteria in GR-820-CORE, *OTGR Section 5.1: Generic Digital Transmission Surveillance*.
-

5.4.6 SONET Timing Reference Switching and Entry Into Holdover

Sections 3.4 and 8.3 of GR-1244-CORE contain various criteria related to timing reference switching (both automatic and manually initiated) and entry into holdover or the free-run mode. In general, those criteria, along with a number of SONET-specific criteria contained in this section, are applicable to SONET NEs.

-
- R5-206 [298v4]** A SONET NE shall meet the criteria related to timing reference switching and entry into holdover or the free-run mode in Sections 3.4 and 8.3 of GR-1244-CORE, but with the triggers for automatic switching and mode changes expanded to include the additional conditions under which the NE would disqualify a reference (e.g., that it contains the DUS synchronization status message, see **R5-218 [311]**).
- R5-207 [305]** If the active timing reference is an OC-N “interface”, switching to an alternate timing reference (if available) shall only take place after it has been determined that any available protection switching of the OC-N line and its terminating circuitry has failed to end the outage. The clock shall maintain accuracy and meet the phase transient criteria in [Section 5.4.4.3.3](#) during the protection switch. Traffic and timing need not be protected together.
-

Note that several of the criteria referenced in **R5-206 [298v4]** refer to switching between references. Those particular criteria are applicable only to SONET NEs that support timing modes that utilize more than one possible reference (e.g., externally timed NEs, line-timed NEs with two OC-N “interfaces” provisioned as possible references). Also, as indicated in GR-1244-CORE, in some cases an NE may need to enter the free-run mode instead of holdover when all of its references fail (e.g., if the holdover value has been corrupted, see **O1244-130**).

Also note that in order for an NE that does not support provisioning of an OC-N/M “interface” as a single reference (see **R5-129 [1051]**) to meet **R5-207 [305]**, the working and protection lines would need to be listed consecutively in the NE’s timing reference list. In addition, such an NE may use the same type of switching (i.e., revertive or nonrevertive) between references at a single “interface” as between references at separate “interfaces”. (See **O5-216 [1118]** regarding the type of switching to be used for switches between lines at a single “interface” when that “interface” is considered a single reference.)

Finally, note that a number of the criteria related to manually initiated timing reference switching that now appear in GR-1244-CORE were based on criteria that appeared in Issue 3 of this document. In order to reduce the redundancies (and the potential for inconsistencies) between the two documents, this section has been modified (as of Issue 4) to address only those topics not covered by the criteria that were added to GR-1244-CORE (i.e., primarily SMCs and synchronization status messages). In addition, various details related to the functionality of the Forced Reference Switch command that were under study at the time that Issue 3 of GR-1244-CORE was released are now specified below.

R5-208 [300v5] In addition to the cases discussed in GR-1244-CORE, a Manual Reference Switch command supported by a SONET NE shall be denied if it would cause a switch to a reference with a lower synchronization status message (including a DUS message), and shall be preempted if the incoming synchronization status messages change such that the NE needs to perform a timing reference switch according to **R5-217 [310]**.

R5-209 [1186] If the Forced Reference Switch command is supported, its impact shall be to effectively cause all of the provisioned timing references except for the specified reference to be suspended from the NE's timing reference list. In addition, the command shall be denied only if the specified reference has already been locked out via the Lockout a Reference command (if supported), and the other references shall remain suspended until the Forced Reference Switch command is either cleared, or preempted by a Forced Reference Switch command specifying a different reference or a Lockout a Reference command specifying the same reference.

Consistent with the Lockout a Reference case discussed in GR-1244-CORE, if the reference specified in a Forced Reference Switch command is considered failed or unavailable, the NE would need to enter holdover. Also, if an NE's timing reference list supports a maximum of two entries, then the Forced Reference Switch command functionality defined above and the Lockout a Reference command functionality defined in GR-1244-CORE would be redundant (i.e., locking out one of the two references would be equivalent to performing a forced switch to the other reference), and therefore such an NE would not be expected to support both of those commands.

In addition to the reference-switching-related commands described in GR-1244-CORE, in some cases it may be useful for a line-timed NE (or possibly a through-timed ADM) to provide one or more commands that would allow the user to control whether the incoming signal on working line 1 or the protection line is used for timing.

O5-210 [1116] An NE that supports the use of the incoming signals on working line 1 and the protection line at an OC-N/M "interface" as a single reference should also support one or more commands that allow the user to control which of those signals is used for timing purposes.

Note that unless one of them is an AIS-L signal (and thus is considered to be failed), the signals received on the working and protection lines at a particular OC-N/M "interface" are originated by the same far-end SONET NE and can be expected to

have the same synchronization characteristics (e.g., they are synchronized to the same clock and contain the same synchronization status message). Therefore, the capability to manually switch between those signals is not considered to be as important as the capability to manually switch between “interfaces”, and the particular commands that should be supported are not specified.

Also note that an NE may, but is not required to, protect traffic and timing together (see **R5-207 [305]**). If traffic and timing are protected together, then the NE can meet **O5-210 [1116]** simply by supporting the required traffic-related protection switch commands (i.e., a manually initiated traffic protection switch or lockout will result in the corresponding timing-reference-related switch or lockout).

-
- R5-211 [1117]** If a SONET NE supports one or more commands to allow the user to control whether the incoming signal on working line 1 or the protection line at an OC-N/M “interface” is used for timing, the functionalities of those commands shall be clearly documented.
-

5.4.6.1 Timing Reference Failure Conditions

In general, the conditions for a SONET NE to consider a timing reference failed are a superset of those discussed in Section 3.4.1 of GR-1244-CORE. Therefore the criteria in that section are applicable to SONET NEs (see **R5-206 [298v4]**), along with the additional criteria discussed below.

-
- R5-212 [1022]** An incoming SONET signal shall be considered failed or unavailable for timing purposes under the following conditions:
- Loss of signal energy (e.g., LOS defect detection/failure declaration)
 - Loss of framing (e.g., LOF defect detection/failure declaration)
 - Line AIS (e.g., AIS-L defect detection/failure declaration at an NE containing LTE).
- R5-213 [315v3]** For SONET signals, if no validated synchronization status message is detected (e.g., due to transmission errors or the receipt of an undefined message, see [Section 5.4.7.1](#)) for a period of greater than 10 seconds, the NE shall consider the signal failed for timing purposes.
- R5-214 [317v3]** For DS1 references in the ESF format, if no validated synchronization status message is detected (see [Section 5.4.7.1](#)) for a period of greater than 10 seconds, the SONET NE shall consider the reference to be failed (unless the reference has been provisioned as not supporting synchronization status messages, see **R5-224 [318]**).
- R5-215 [302v2]** A timing reference shall be considered failed or unavailable upon receipt of synchronization status message indicating that the reference is traceable to a source in holdover/free-run that is of worse quality than the local internal clock.²⁵
-

Note that (consistent with the discussion in GR-1244-CORE) **R5-212 [1022]** provides some flexibility with respect to the time that an NE takes to consider a reference failed. Specifically, an NE is not required to consider a signal to be failed for timing purposes immediately upon detecting a particular defect, and is also not required to delay considering the signal as failed until the corresponding failure is declared. What is important is that the NE conform to the applicable criteria independent of whether or not it considers the signal failed (e.g., to **R5-168 [271v3]** in [Section 5.4.4.3.6](#) in situations where a timing reference signal is not considered failed).

Also note that although the receipt of a DUS message disqualifies a reference or derived DS1 source from being active (see **R5-218 [311]**), this condition does not cause the reference to be considered “failed or unavailable for timing purposes”. While this may not seem logical, declaring a Loss of Synchronization failure and generating an alarm upon the receipt of a DUS message could lead to standing alarms at SONET NEs, and it is widely agreed that standing alarms should be avoided. Furthermore, the receipt of a single DUS message is a normal condition in many applications. The purpose of generating an alarm is to alert people that something is abnormal and needs to be corrected.

Finally, while the criteria in Section 3.4.1 of GR-1244-CORE and most of this section are specific to signals provisioned for use as timing references, **R5-212 [1022]**, **R5-213 [315v3]** and the preceding discussion also applies in cases where a SONET signal is provisioned to be used as a derived DS1 source (and thus those requirements use the general phrase “for timing purposes” instead of a specific phrase such as “for use as a timing reference”).

As identified in Section 3.4.1 of GR-1244-CORE, a SONET NE may be required to reject a reference based on frequency offset by some service providers in some applications. In particular, some service providers are concerned that payload performance is not guaranteed for frequency offsets greater than 4.6 ppm, while clocks in SONET NEs may be pulled off-frequency to offsets greater than this. However, even stratum 3 clocks cannot provide frequency rejection at 4.6 ppm.

5.4.6.2 Performance During Timing Reference Switching

Refer to [Section 5.4.4.3.5](#) for the rearrangement phase transient limits that are applicable to reference switches. The MTIE requirement does not apply for reference switches caused by references that fail due to off-frequency conditions.

25. For example, a SONET NE with an internal stratum 3 clock would consider a reference failed if it is carrying a synchronization status message indicating that it is traceable to an SMC or stratum 4 clock.

5.4.6.3 Revertive and Nonrevertive Timing Reference Switching

The applicable criteria for revertive and nonrevertive switching between references are contained in Section 3.4.3 of GR-1244-CORE (see **R5-206 [298v4]** in this document). Note that those criteria are specifically intended to apply to switching between references, not switching between the working and protection lines at a single OC-N “interface”. The applicable objective for that case appears below.

-
- O5-216 [1118]** Unless it protects traffic and timing together (see **R5-207 [305]**) and is using revertive switching for the traffic, a SONET NE that supports the use of the incoming signals on working line 1 and the protection line at an OC-N/M “interface” as a single reference should use nonrevertive switching for timing-reference-related switches between those signals.
-

5.4.6.4 Synchronization Status Messages and Timing Reference Switching

The following reference switching requirements apply for SONET NEs that support synchronization status messages.

-
- R5-217 [310]** The available, provisioned reference with the highest quality synchronization status message shall be selected as the active reference.
- R5-218 [311]** The NE shall not select a reference with a DUS message as the active synchronization reference.
-

Note that the preceding requirements cover switching between references (e.g., OC-N “interfaces”), but do not include switching between the signals on the working and protection lines at a single OC-N “interface”. In general, working/protection-line switching based on synchronization status messages is not required to be supported (although it may be provided by some NEs). The reason for this is that based on the existing requirements, any NE that supports line APS must transmit the same synchronization status messages on working line 1 and the protection line (at a single “interface”) in all situations. Thus, the messages that are received by the far-end NE on those two lines can be expected to be identical, and it is considered sufficient for that NE to monitor the incoming messages on only one of those lines (e.g., the line whose signal is currently being used as the timing source).

Also note that requirement **R5-217 [310]** implies that when the active reference has a degradation in synchronization quality or an alternate reference has an improvement in synchronization quality, a reference switch may be triggered. In many cases, it would be appropriate to perform this switch “immediately”. However in some cases, the absence of a “holdoff” time before the switch is performed will result in extraneous reference switches. In particular, this is expected to be an issue for externally timed NEs that are timed from a BITS that does not change the synchronization status messages on all of its outputs simultaneously.

-
- R5-219 [1023v2]** An externally timed NE shall refrain for at least 10 seconds from performing a switch between its primary and secondary external DS1 reference
-

signals based on a change in the synchronization status messages contained in those signals, unless that change causes the NE to consider its currently active reference to be failed or unavailable (see **R5-215 [302v2]**), or the message in the active reference changes to the DUS message (see **R5-218 [311]**). If a reference switch is still appropriate at the end of this 10-second holdoff period, the NE shall perform the switch within an additional 1-second period.

Table 5-10 contains an example in which the synchronization status messages on two references go through a series of changes and indicates which reference would be selected by the NE to be active. Two scenarios, revertive and nonrevertive reference switching, are illustrated. Note that the entries in the table are not static, individual examples. Instead, the table represents a series of events starting with the conditions of the first row, and proceeding through various changes to the conditions of the last row.

Table 5-10 Example Reference Selection Using Synchronization Status Messages

Sync Message on Reference A	Sync Message on Reference B	Active Reference Nonrevertive	Active Reference Revertive (default A)
PRS	STU	A	A
PRS	PRS	A	A
ST2	PRS	B	B
ST2	STU	B	B
ST2	PRS	B	B
PRS	PRS	B	A

5.4.7 Synchronization Status Message Validation and Generation

This section discusses the synchronization status messaging criteria for SONET NEs. The specific criteria describe how SONET NEs validate, react to, and generate synchronization status messages.

5.4.7.1 Message Validation

In general, it is important that the scheme used to validate the synchronization status messages be robust to occasional errors in the messaging channel.

R5-220 [312] The supplier shall document the scheme used to validate messages on the ESF data link for DS1 synchronization references and on the S1 byte for SONET signals.

R5-221 [313v2] A SONET NE shall validate the message carried in a SONET signal if the same defined message (see [Table 5-9](#)) is detected in 8 consecutive samples (these may or may not be consecutive SONET frames) of bits 5 through 8 of the S1 byte. The sampling rate shall be such that the maximum time to detect a change of the incoming message is 1 second (assuming no transmission errors).

R5-222 [314v2] A SONET NE shall validate the message carried in an ESF DS1 synchronization signal if the same defined message is detected in at least 7 out of 10 consecutive samples of the 16-bit codewords in the ESF data link.

In cases where an NE is not able to validate a synchronization status message on a timing reference or derived DS1 source signal, **R5-213 [315v3]** and **R5-214 [317v3]** apply. In addition, if a DS1 reference is in the SF format, messaging cannot be expected and the following requirement is applicable.

R5-223 [316] For DS1 references in the SF format, the SONET NE shall consider the reference to have a “Synchronized – Traceability Unknown” message.

If a DS1 reference is in the ESF format, the NE can initially assume that the reference supports synchronization status messaging. Therefore, the NE can expect to see a valid message on the reference. However, an external reference in the ESF format may sometimes be used without synchronization status messages. In this case, the user would want the NE to behave differently. A provisioning step is necessary to make the distinction between an ESF signal that supports synchronization status messages and one that does not.

R5-224 [318] The user shall be able to provision the NE to accept an external DS1 reference in the ESF format that does not support synchronization status messages. For such a provisioned reference, the SONET NE shall consider the reference to have a “Synchronized – Traceability Unknown” message.

5.4.7.2 Message Reaction

For the most part, the criteria related to the actions that a SONET NE needs to take in response to changes that it detects in the synchronization status messages on its incoming signals are provided in other sections (e.g., [Sections 5.4.5.2.1](#) and [5.4.6.4](#) for switching between timing reference signals or derived DS1 source signals). This section provides an additional “message reaction” requirement related to the reporting of synchronization status message changes to an Operations System (OS) for network management purposes. In addition, it also contains an objective to support the ability of the user to determine the messages received at different interfaces to the NE.

R5-225 [319v2] As a default, the NE shall automatically generate a status report to an OS when the synchronization status message on any provisioned reference or derived DS1 source changes. The report shall indicate which reference or derived DS1

source changed, the time of the change, the old synchronization status message, and the new synchronization status message.²⁶

- O5-226 [320v4]** The NE should report to the user, on demand, the synchronization status message at any of its SONET “interfaces” (output and input, including those where processing of the incoming S1 byte has been disabled), and on its external DS1 reference signals and derived DS1 signals (when supported).

5.4.7.3 Message Generation

The criteria in this section apply to all NEs that support synchronization status messaging, while [Sections 5.4.7.3.1](#) through [5.4.7.3.3](#) contain related criteria that apply to NEs provisioned to operate in specific timing modes.

- R5-227 [321v2]** When the synchronization status message in a signal that is provisioned to be used as a timing reference and/or derived DS1 source changes, the NE shall validate the change or determine that there is no valid message as per [Section 5.4.7.1](#), react appropriately (e.g., as per the criteria in [Sections 5.4.5.2](#), [5.4.6.4](#) and [5.4.7.2](#)) and, except in cases covered by **O5-228 [1187]** (if met), **R5-231 [323v2]**, **R5-232 [324]** or **R5-236 [1024v2]**, insert the appropriate synchronization status messages in all transmitted SONET and/or derived DS1 signals within 10 seconds.²⁷

- O5-228 [1187]** When a SONET NE performs a timing reference switch and the quality level of the synchronization status message on the newly active reference is higher than that of the message that was being transmitted at the time of the switch (based on the message received in the previously active reference), the NE should not change its outgoing synchronization status messages to indicate the quality of the new reference signal until it has completely synchronized to that signal.

²⁶. Note that in an office with a BITS that supports synchronization status messages, a change in the message being transmitted by that BITS could cause multiple externally timed NEs to report changes on their references. In such cases, the user may want to disable the automatic generation of these status reports. GR-474-CORE, *OTGR Section 4: Network Maintenance: Alarm and Control for Network Elements*, contains the applicable criteria related to such provisioning. (Also see [Section 6.2.1.8.5](#) of this document.)

²⁷. “All transmitted signals” includes any SONET signal transmitted from an “interface” where synchronization status messages are supported. See [Section 5.4.2](#) for the criteria on supporting synchronization status messages at line-side and drop-side “interfaces”. Also, for an NE that is operating in the threshold AIS generation mode (see [Section 5.4.5.2.2](#)) or the case where there is no validated message either before or after the incoming message changes, the phrase “insert the appropriate synchronization status message” is to be interpreted as “generate or stop the generation of AIS, as appropriate” (at the affected derived DS1 output). Finally, note that in this and several other requirements (e.g., **R5-229 [1188]**, **R5-231 [323v2]**), “10 seconds” is an upper limit and therefore a faster response would generally be considered to be “better”. These cases should not be confused with other criteria in which a required or minimum delay of 10 seconds is specified (e.g., see **R5-219 [1023v2]**).

- R5-229 [1188]** If **O5-228 [1187]** is met, the time that the NE takes to change the message shall be no longer than the locking time (as specified in GR-1244-CORE) plus 10 seconds.
- R5-230 [322v2]** When the NE enters holdover or free-run, the synchronization status message on all of its transmitted SONET signals (other than those provisioned to carry the DUS message, see **R5-232 [324]**) shall be changed within 10 seconds to indicate the holdover level of the SONET NE's internal clock (e.g., "Stratum 3 Traceable" or "SMC Traceable").
- R5-231 [323v2]** When the NE recovers from holdover or free-run, the synchronization status message shall not change to indicate the quality of the reference signal until the NE has completely resynchronized. The time to change the message in that case shall be no longer than the recovery time [i.e., the sum of the requalification time (as specified in [Section 5.4.4.3.4](#)) and the locking time (as specified in GR-1244-CORE)] plus 10 seconds.

Note that **O5-228 [1187]** and **R5-231 [323v2]** are specifically applicable for cases where an outgoing synchronization status message is being changed to indicate the quality of the newly active reference. For SONET "interfaces" where the new message is going to be the DUS message, **R5-234 [326v3]** and **R5-242 [330]** are the applicable requirements. Based on those requirements, the new message must be inserted with no significant delay (i.e., within 10 seconds after the reference has been requalified). Also note that **O5-228 [1187]** does not apply in cases where the message on the active reference changes but no reference switch is performed, and that if **O5-228 [1187]** is not met (so **R5-229 [1188]** is not applicable) then **R5-227 [321v2]** applies.

-
- R5-232 [324]** The NE shall allow the user to individually provision each SONET "interface" so that the transmitted signals from that "interface" carry the DUS message instead of the message that reflects the actual traceability of the signals.
-

The primary purpose of **R5-232 [324]** is to allow service providers to block the synchronization status messages on the outgoing signals at certain interfaces so that those signals cannot be used for timing purposes at the far end. Typical applications could be at carrier-to-carrier or carrier-to-customer interfaces.

5.4.7.3.1 Externally Timed NEs

In [Section 5.4.7.1](#) it was described how the NE reacts differently depending upon the format of the external DS1 reference. If the external reference is in the SF format, the NE knows that messaging cannot be expected on the signal. Conversely, the NE assumes that an ESF DS1 reference supports synchronization status messages. A similar distinction based on the format of the external DS1 reference is made for the message generation criteria.

If the external DS1 reference is in the ESF format, the following requirements apply (unless the NE has been provisioned to not expect messages on the external ESF DS1 reference).

R5-233 [325] If none of the terminating signals at a particular SONET “interface” are being used to derive a DS1, then the NE shall insert the synchronization status message from the active external ESF DS1 reference on the SONET signals transmitted from that “interface”.

R5-234 [326v3] If a terminating signal at a particular SONET “interface” is being used to derive a DS1 and the synchronization status message on the active external ESF DS1 reference matches the synchronization status message being transmitted on the derived DS1 (translated from the terminating SONET signal), then the NE shall insert the DUS message on the SONET signals transmitted from that “interface” (unless the automatic generation of the DUS message has been disabled for all of the DS1s being derived from that interface, see **CR5-237 [1061]** and **R5-238 [1062]**).

R5-235 [327v4] If a terminating signal at a particular SONET “interface” is being used to derive a DS1 and (in the steady-state, see **R5-236 [1024v2]**) the synchronization status message on the active external ESF DS1 reference does not match the synchronization status message being transmitted on that DS1 (translated from the terminating SONET signal) or the automatic generation of the DUS message has been disabled for all of the DS1s being derived from that interface (see **CR5-237 [1061]** and **R5-238 [1062]**), then the NE shall insert the synchronization status message from the active external reference on the SONET signals transmitted from that “interface”.

The effect of these requirements is illustrated in [Figure 5-24](#), in which OC-N #1 carries the DUS message in accordance with **R5-234 [326v3]**, and OC-N #2 carries the PRS message in accordance with **R5-235 [327v4]**.

R5-236 [1024v2] When an NE that is transmitting the DUS message at one of its SONET “interfaces” (because it meets the condition described in **R5-234 [326v3]**) detects a change in the incoming synchronization status message at that “interface”, it shall continue to generate a DUS message for 10 seconds after the synchronization status message has been translated to the derived DS1 (i.e., after the message on the outgoing derived DS1 has been changed to reflect the new message received at that SONET “interface”).

In general, the purpose of the delay described in **R5-236 [1024v2]** is to allow a BITS time to detect changes on its references and react appropriately (e.g., change the messages on its outputs). After the period of 10 seconds, **R5-234 [326v3]** and **R5-235 [327v4]** are again the applicable criteria for determining the message to be generated on the transmitted SONET signals.

In certain applications (e.g., in applications where a particular derived DS1 is not being used as a reference signal for the BITS), it may be desirable for an externally timed NE to provide the user with the capability to disable the automatic generation of the DUS message according to the requirements shown above. To support such applications, the following criteria apply. Note that similar criteria related to disabling the automatic generation of the DUS message by a line-timed SONET NE do not exist, and that it is recommended that such a feature not be provided.

CR5-237 [1061] An externally timed SONET NE may be required to be capable of being provisioned such that it does not automatically generate the DUS message according to **R5-234 [326v3]**.

R5-238 [1062] A SONET NE that can be provisioned to derive each of its DS1s from a different SONET “interface” (see **R5-175 [273v4]**) and that also provides the capability to disable the automatic generation of the DUS message according to **R5-234 [326v3]**, shall provide that capability on a per-derived-DS1 basis.

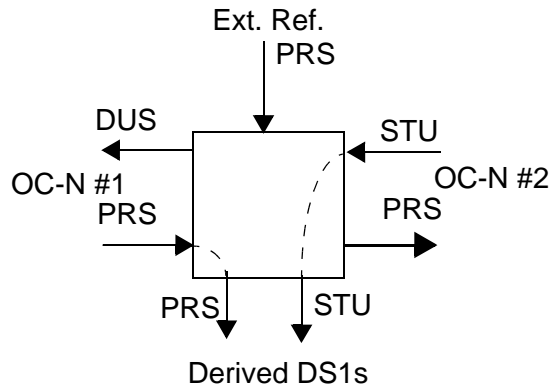
R5-239 [1063] If a SONET NE provides the capability to disable the automatic generation of the DUS message according to **R5-234 [326v3]**, its default shall be that the automatic generation is enabled.

If the external DS1 reference is in the SF format or the NE has been provisioned to not expect messages on the external ESF DS1 reference, the following requirement applies.

R5-240 [328v2] If the active external reference does not carry synchronization status messages (e.g., an external DS1 reference is in the SF format or the NE has been provisioned not to expect synchronization status messages on the ESF DS1), the NE shall insert the “Synchronized – Traceability Unknown” message on all transmitted SONET signals that have not been provisioned to continuously carry the DUS message (see **R5-232 [324]**).

Based on **R5-240 [328v2]** and **R5-230 [322v2]** (which applies if the NE enters holdover or free-run), a SONET NE that is provisioned to be externally timed using DS1 signals that do not carry synchronization status messages will not automatically generate the DUS message at any of its interfaces, even if the incoming signals at one or more of those interfaces are being used to derive DS1 signals. The reason for this is that the primary purpose of automatic DUS generation is to support automatic reconfiguration of the synchronization distribution network (in response to failures), and such reconfiguration activities are not feasible in a network utilizing BITS clocks, externally timed NEs and derived DS1 signals unless synchronization status messages are supported at both the SONET and DS1 interfaces. In cases where a service provider wants the DUS message transmitted in the outgoing signals at a SONET interface where an incoming signal is being used to derive a DS1 (e.g., to preclude the formation of a timing loop due to provisioning errors at other NEs), the NE can be provisioned to continuously transmit the DUS message at that interface according to **R5-232 [324]**.

Figure 5-24 Example Implementation of R5-234 and R5-235




5.4.7.3.2 Line-Timed NEs

When the NE is line timed, the following requirements apply.

-
- R5-241 [329]** At all SONET “interfaces” that are not the active synchronization reference, the line-timed NE shall insert the synchronization status message from the active synchronization source on the transmitted SONET signals (see Figure 5-25, OC-N #2).
- R5-242 [330]** The line-timed NE shall generate a DUS message in the transmitted signals from the OC-N “interface” that is the active synchronization reference (see Figure 5-25, OC-N #1).
-

Figure 5-25 Example Implementation of R5-241 and R5-242




5.4.7.3.3 Through-Timed NEs

When a SONET ADM is through timed, the following requirements apply.

-
- R5-243 [331]** A through-timed ADM shall insert the synchronization status message from the terminating OC-N in the transmitted OC-N in the same direction (see Figure 5-26 for an example).

- R5-244 [332]** A through-timed ADM shall insert the synchronization status message of the appropriate timing source in all dropped SONET signals.

Figure 5-26 Example Implementation of R5-243



5.5 Framing For SONET Signals

This section contains criteria related to the monitoring of incoming SONET signals for SEF defects. These criteria are applicable to all SONET NEs with STE functionality, and also to any non-STE SONET NEs that need to frame on the SONET signal for monitoring and AIS-L insertion purposes (e.g., physical layer regenerators). Also see [Section 6.2.1.1.2](#) for related criteria on LOF defects and failures (which occur when SEF defects persist).

- R5-245 [333v2]** The framing pattern (or patterns) observed by a SONET NE shall include a subset of the A1 and A2 bytes contained in the incoming STS-N electrical or OC-N signal.

Note that it is not necessary for a SONET NE to use the same framing pattern for detecting SEF defects as it uses for terminating SEF defects (i.e., it may monitor different subsets of the A1 and A2 bytes for detecting and terminating the SEF defect). In addition, a framing pattern does not necessarily have to consist of full bytes. For example, an NE could monitor certain bits of a particular A1 byte, along with all or parts of one or more other A1 and A2 bytes for detecting SEF defects.

In general, it is important for a SONET NE's SEF defect detection and termination algorithms to be robust. In particular, the detection algorithm needs to allow the NE to quickly detect an SEF defect when there has been a change of frame alignment on an incoming signal, while avoiding the detection of SEF defects that are caused by simple errors in the framing pattern (e.g., by a high BER). Similarly, the termination algorithm needs to allow the NE to quickly locate the actual A1 and A2 bytes and terminate the SEF defect, while avoiding framing on a "false" framing pattern elsewhere in the signal. The intent of the following criteria is to promote the development of SEF defect detection and termination algorithms that meet these four goals without unnecessarily constraining the possible designs.

- R5-246 [334]** An SEF defect shall be detected when the incoming signal has a minimum of four consecutive errored framing patterns. The maximum SEF detection time shall be 625 μ s for a random signal.

R5-247 [335] The framing algorithm used to check the alignment shall be such that an SEF defect is not detected more than an average of once every 6 minutes while the BER of the STS-N electrical or OC-N signal is 10^{-3} .²⁸

R5-248 [336] Once an SEF defect has been detected, the SONET NE shall terminate the SEF defect upon detecting two successive error-free framing patterns.

Any implementation of the frame recovery circuitry that, following a detected SEF defect, achieves realignment within the 250- μ s interval implied by the preceding requirement is acceptable.

O5-249 [1119] A SONET NE should provide some method of “reducing” the probability that it will frame on any false framing patterns that may be contained in an incoming STS-N electrical or OC-N signal.

R5-250 [1120] Any methods that the SONET NE uses to reduce the probability that it will frame on false framing patterns shall be documented.

Note that in order for an NE to frame on a false framing pattern, it would be necessary for the bits containing that pattern to remain constant from one SONET frame to the next. For most SONET signals, the bits that are most likely to mimic the contents of the A1 and A2 bytes and remain constant for relatively long periods of time are those contained in the H2 bytes. The reason for this is that each H2 byte can have any value from ‘00’ to ‘FF’ (hex), and its value remains constant unless the STS pointer value needs to be changed to accommodate a change in the alignment of the transported STS SPE. In an STS-N electrical or OC-N signal there are N consecutive (and generally independent) H2 bytes, and at the higher bit rates this results in a number of combinations of STS pointer values that can (depending on the particular algorithm) result in the presence of false framing patterns.

In general, there are a number of methods that an NE could use to “reduce” the probability that it will frame on a false framing pattern. These include (but are not limited to) the following.

- The framing pattern used by the NE in its STS-N electrical or OC-N SEF defect termination process could contain more than N bytes (i.e., more bytes than the number of H2 bytes contained in the signal and available to carry a false framing pattern). Note that this method would be applicable primarily in cases where N is relatively small, as the use of a large number of bytes in the SEF defect termination process could significantly delay that process in cases where the BER of the incoming signal is high.
- The framing pattern used by the NE in its SEF defect termination process could contain enough of the A1 and A2 bytes that the probability of those bytes being mimicked elsewhere in the signal is extremely small (i.e., significantly smaller than for a three-byte or four-byte pattern).
- The framing pattern used by the NE in its SEF defect termination process could include bytes that are separated by more than N bytes (e.g., the first A1 byte and

²⁸A Poisson distribution of bit errors is assumed.

the last A2 byte) so that the framing pattern cannot be mimicked solely in the H2 bytes. Note that this method cannot be used for terminating SEF defects on OC-768 signals, as there are only 126 bytes in between the first A1 byte and the last A2 byte.

- When an incoming signal contains two or more framing patterns (i.e., the actual pattern and one or more false patterns), the NE could use the contents of some other byte (e.g., the BIP value contained in a byte position that should be occupied by a B2 byte) to determine which pattern it should frame on.

While a BIP value could be used as described in the last bullet item listed above, the use of such values in an NE's "normal" framing process (i.e., when the incoming signal contains only a single apparent framing pattern) is generally not recommended. The reason for this is that it could significantly delay the SEF defect termination process in cases where the BER of the incoming signal is relatively high.

5.6 Jitter

This section provides jitter criteria for SONET transport systems and equipment. In general, these criteria are consistent with the specifications that appear in ANSI T1.105.03, *Synchronous Optical Network (SONET) – Jitter and Wander at Network and Equipment Interfaces*, and ITU-T Rec. G.825, *The control of jitter and wander within digital networks which are based on the synchronous digital hierarchy (SDH)*. However, since neither of those standards addressed OC-768/STM-256 signals at the time that this document was being prepared, the OC-768 criteria that appear here are based on a combination of the specifications for OTU3 signals in ITU-T Rec. G.8251, *The control of jitter and wander within the optical transport network (OTN)*, and extrapolations based on the criteria that appear in this document for lower bit-rate SONET signals. If in the future, different specifications are provided for the OC-768 case in T1.105.03 and/or Rec. G.825, those specifications will supersede the criteria provided here.

As indicated in GR-499-CORE, jitter is defined as the short-term variations of a digital signal's significant instants (e.g., optimum sampling points) from their ideal positions in time. In addition, short-term variations are defined as phase oscillations of frequency greater than 10 Hz (i.e., phase modulations that, after demodulation, pass through a high-pass filter with a cutoff frequency of 10 Hz and a 20 dB/decade roll-off).²⁹ Long-term variations (i.e., phase modulations that, after demodulation, pass through a low-pass filter with a cutoff frequency of 10 Hz and a roll-off of 20 dB/decade) are defined as wander and are addressed as part of synchronization criteria in [Section 5.4](#), and the phase variation criteria in [Section 5.7](#).

5.6.1 Network Interface Jitter Criteria

The following requirement limits the levels of jitter on OC-N and STS-N electrical signals. These “network interface jitter limits” restrict the wideband and highband jitter appearing on a SONET signal anywhere in the network. Therefore, the requirement applies at interfaces between users and carriers or between two carriers, as well as at the output ports of individual NEs.

R5-251 [337v2] Timing jitter at network interfaces shall not exceed $A_1 U_{I,pp}$ when measured over a 60-second interval with a bandpass filter having a high-pass cutoff frequency of B_1 and a roll-off of 20 dB/decade, and a low-pass cutoff frequency of B_3 and a roll-off of 60 dB/decade.³⁰

Timing jitter at network interfaces shall not exceed $A_2 U_{I,pp}$ when measured over a 60-second interval with a bandpass filter having a high-pass cutoff frequency B_2 and a roll-off of 20 dB/decade, and a low-pass cutoff frequency of B_3 and a roll-off of 60 dB/decade.

The values for the A_1 , A_2 , B_1 , B_2 and B_3 parameters referred to in this requirement appear in [Table 5-11](#). Parameter values are given for all of the currently defined OC-N and STS-N electrical interfaces except OC-24. (The applicable values for that bit rate appear in ANSI T1.105.03.) Note that two sets of values are given for OC-48 interfaces. The reason for this is that for that bit rate, allowances were made in ITU-T Rec. G.958, *Digital line systems based on the synchronous digital hierarchy*

29. Note that although the demarcation frequency between jitter and wander is defined here as 10 Hz, the frequency range of interest in a particular jitter or wander test or measurement will not necessarily include the entire defined range (or only the defined range). Therefore, measurement filter cutoff frequencies and minimum or maximum test frequencies are generally included in the criteria (e.g., see [Table 5-11](#) for a number of different measurement filter cutoff frequencies to be used in wideband and highband SONET network interface jitter measurements). Also note that as long as the appropriate filter and/or test frequencies are specified, the particular label that is given to phase variations within a particular frequency range is not important. For example, in a case where the cutoff frequency of the low-pass filter specified for use in a wander generation test is lower than the cutoff frequency of the high-pass filter specified for use in the jitter generation test for a particular type of equipment, whether the phase variations at frequencies between those two cutoff frequencies are considered to be wander or jitter will have no impact on a tested equipment's conformance to the criteria.

30. In most cases, the amount of jitter that will occur at frequencies above the low-pass cutoff frequencies of the filters specified for use in this and other jitter measurements is expected to be very small. Therefore, some documents (including previous issues of this GR) use/used the term “at least” when specifying those cutoff frequencies, meaning that the use of filters with higher low-pass cutoff frequencies is/was considered acceptable. On the other hand, most of the current ANSI and ITU-T standards related to jitter provide very specific values for both the cutoff frequencies and roll-off characteristics of the measurement filters. To reduce the chance of disagreements that can be caused by the application of different specifications or the use of different test equipment, “at least” has been removed from the various filter specifications in this document (as of Issue 4). In addition, the specified characteristics of the low-pass filters used in these measurements are those of a maximally flat Butterworth filter.

for use on optical fiber cables, for two different types of regenerators. These typically exhibit very different timing recovery circuit Q factors, and are referred to as:

1. Type A regenerators, in which timing recovery circuits employ electronic devices [e.g., Surface Acoustic Wave (SAW) or dielectric resonator filters] and the output is timed directly from the recovered clock with no additional filtering
2. Type B regenerators, in which timing circuits typically employ Phase Lock Loops (PLLs).

Type A regenerators are generally characterized by greater jitter tolerance, especially at high jitter frequencies, and by more high frequency jitter accumulation due to their relatively broad jitter transfer characteristics (see [Sections 5.6.2.1](#) and [5.6.2.2.2](#) for definitions of jitter transfer and jitter tolerance). Conversely, Type B regenerators are typically characterized by less tolerance to high frequency jitter due to the high Q factors associated with PLLs, and by less high frequency jitter accumulation due to their narrow phase transfer characteristics. Thus, the second set of OC-48 network interface values [OC-48(B)] corresponds to a tighter jitter limit that must be met to ensure proper performance in the presence of NEs exhibiting “reduced” jitter tolerance (i.e., Type B regenerators, or OC-48 Line terminating NEs that are not designed to operate in line systems containing Type A regenerators).

Also note that unlike the OC-48 case, the various jitter criteria that apply to OC-192 and OC-768 interfaces and equipment were developed assuming a single type of regenerator. In addition, it was assumed that the OC-192 or OC-768 regenerator would be a hybrid of the OC-48 Type A and B regenerators, employing a highly tolerant timing recovery circuit at the input, followed by a buffer and PLL at the output. This results in OC-192 and OC-768 network interface jitter and jitter tolerance requirements that are comparable to those for OC-48 systems utilizing Type A regenerators, and a jitter transfer requirement that is somewhat comparable to the specifications for Type B OC-48 regenerators. This in turn is intended to optimize the OC-192 and OC-768 line system robustness.

Finally, it should be noted that subsequent to the inclusion of the dual sets of OC-48 jitter criteria in this document, ITU-T Rec. G.958 was withdrawn and the relevant information was added to ITU-T Rec. G.783. In Rec. G.783 the discussion related to the different types of regenerators and normal/reduced jitter tolerance was expanded to include STM-1 and STM-4 interfaces. The resulting additional specifications for those two interfaces have not been reflected in ANSI T1.105.03, and in the view of Telcordia are not necessary for SONET systems. Thus, no additional sets of OC-3 or OC-12 jitter criteria are expected to be added to this GR.

Table 5-11 Parameters for Network Interface Jitter Requirements

OC-N/STS-N Level	B ₁	B ₂	B ₃	A ₁ (UI _{pp})	A ₂ (UI _{pp})
1	100 Hz	20 kHz	400 kHz	1.5	0.15
3	500 Hz	65 kHz	1.3 MHz	1.5	0.15
12	1 kHz	250 kHz	5 MHz	1.5	0.15
48	5 kHz	1 MHz	20 MHz	1.5	0.15
48(B)	5 kHz	12 kHz	20 MHz	1.5	0.15
192	20 kHz	4 MHz	80 MHz	1.5	0.15
768 ^a	20 kHz	16 MHz	320 MHz	6.0	0.15

Note:

- a. The values in this row should be considered preliminary. See the discussion in [Section 5.6](#).

To ensure that the accumulated jitter at an interface does not exceed the network limits, and that SONET line systems do not experience performance degradations due to jitter, the criteria in [Section 5.6.2](#) are applicable to SONET NEs.

5.6.2 SONET NE Jitter Criteria

SONET equipment jitter criteria are specified in the following areas:

- Jitter transfer ([Section 5.6.2.1](#))
- Jitter tolerance ([Section 5.6.2.2](#))
- Jitter generation ([Section 5.6.2.3](#)).

The criteria in these sections specify limits for SONET equipment interfaces that fall into the following two categories (see GR-499-CORE for the definitions of these categories):

- Category I – Asynchronous DS_n interfaces to SONET NEs are considered Category I interfaces
- Category II – OC-N, STS-N electrical and synchronous³¹ DS₁ interfaces to SONET NEs are considered Category II interfaces.

The Category I jitter generation and transfer criteria are critical to control DS_n jitter accumulation through SONET “islands”. A SONET island is a collection of SONET NEs that create a continuous path for a digital signal that is asynchronously multiplexed at its entry point and asynchronously demultiplexed at its exit point. As a DS_n signal enters and exits these SONET islands, jitter can be mapped from one

31. “Synchronous DS₁ interfaces” refers to DS₁ interfaces where an incoming DS₁ is byte-synchronously mapped into a VT SPE that is synchronized to the SONET NE’s clock, or where a transmitted DS₁ signal is retimed to the NE’s clock (see [Sections 3.4.1.1.1](#) and [5.4.5.3](#)). Other cases are considered to be Category I.

island to the next by the stuffing mechanism in the asynchronous mapping. This can lead to jitter enhancement (i.e., the accumulation of jitter), which may degrade error performance. Islands are an issue for networks as they transition from asynchronous transport to SONET. In an all-SONET network there would be only a single island, and the accumulation of jitter on asynchronous payloads would be less of a concern. In addition, the performance of the synchronization network, as well as the performance of the SONET NEs' internal clocks, may impact the rate at which pointer adjustments are generated and affect jitter accumulation through SONET islands. [Section 5.6.2.4](#) discusses timing jitter for tandem connections of digital equipment in terms of jitter enhancement.

5.6.2.1 Jitter Transfer

Jitter transfer is defined in Section 7 of GR-499-CORE, which also contains Category II to Category I, Category I to Category I, and Category II to Category II jitter transfer criteria for non-SONET NEs. In SONET, Category II to Category I jitter transfer (e.g., the OC-N line jitter that appears on an asynchronous DS_n output) is not expected to be significant. Therefore, this section contains only "Category I" criteria (i.e., Category I to Category I jitter transfer, such as that through a multiplexer/demultiplexer pair), and "Category II" criteria [i.e., Category II to Category II jitter transfer, such as that through a SONET regenerator or Electronic Digital Signal Cross-connect (EDSX)³²].

Note that as discussed in GR-499-CORE, measurement methodology is critical for accurate and useful jitter transfer measurements. It is extremely important that the jitter generation characteristics of the NEs and the noise floor of the jitter measurement system be ascertained and taken into account when jitter transfer measurements are made. In all cases, the input jitter amplitudes that are used need to be large enough to result in output jitter amplitudes that are greater than the noise floor of the measurement system and are not significantly affected by generated jitter at the same jitter frequency. If this is not possible at certain jitter frequencies (e.g., if the necessary level of input jitter is greater than the system's jitter tolerance), then any results obtained at those frequencies are erroneous and should be discarded. Note that the noise floor is particularly important at high jitter frequencies, where input jitter levels are limited to lower values by the jitter tolerance mask and the jitter attenuation is required to be large. Jitter transfer tests would normally be expected to concentrate on frequencies within two decades or less of the break point in the jitter transfer mask.

For a system with a linear jitter transfer function, jitter transfer measurements can be made (and identical results can be obtained) using sinusoidal jitter applied to the input signal at any level up to the jitter tolerance level for that interface and that

32. The Category II jitter transfer criteria could also be applied to the OC-N, STS-N electrical, and synchronous DS1 outputs from other line-, loop-, or through-timed SONET NEs. However, if the synchronous outputs are timed from an NE's clocks (which in turn are timed from one or more incoming signals), then the synchronization criteria in [Section 5.4](#) are much more stringent than the jitter transfer criteria in this section, and jitter transfer does not need to be measured. Conversely, if the transmitted signals are directly controlled by the timing of the received signals (as in a SONET regenerator or EDSX) rather than by the NE's clocks, then many of the synchronization criteria are not applicable, and jitter transfer must be measured.

specific jitter frequency. However, SONET systems typically do not have linear jitter transfer functions (both by design and due to inherent factors such as the limited number of stuff opportunity bits available in the asynchronous DS_n to VT or STS SPE mappings), and therefore the results obtained in any jitter transfer tests are likely to depend on the particular input amplitudes used. In general, the primary purpose of the jitter transfer requirements is to prevent performance degradations by limiting the accumulation of jitter through a series of systems such that it does not exceed the network interface jitter requirements (or the jitter tolerance of any of the NEs involved). Thus, it is more important that a system meet the jitter transfer criteria for relatively high input jitter amplitudes (e.g., amplitudes close to the network interface jitter or jitter tolerance limits) than for very low input amplitudes. Therefore, for testing the conformance of a system to the jitter transfer requirements in this document (e.g., to **R5-253 [338]** or **R5-252 [339]**), the input jitter amplitude range is limited to 0.1 to 1.0 times the amplitude given by the appropriate jitter tolerance mask. (That is, the jitter transferred through the system must be under the jitter transfer mask for any input jitter amplitude within this range, but is not required to be under the jitter transfer mask for input amplitudes outside of the range.)

5.6.2.1.1 Category I Jitter Transfer

The jitter transfer criteria given in this section limit the amount of jitter on an input DS_n signal at one NE that can be transferred (via the bit stuffing mechanisms used in the asynchronous DS_n mappings) to the output DS_n at another NE. In the development of these criteria, it was assumed that any jitter appearing on an input DS_n signal would be mapped into the STS or VT SPE, and that all of the phase smoothing would occur in the desynchronizer when the DS_n is extracted from the SPE.³³ Therefore, the intent of the criteria is to establish the minimum level of phase smoothing that can be performed in desynchronizers. Of course it is also possible for a supplier to design an NE's synchronizer to provide significant attenuation of the input jitter. In a single-supplier network, this could result in acceptable performance even if the NE's desynchronizers do not meet the intent of the criteria. However, for interoperability in a multi-supplier network, a desynchronizer must meet the jitter transfer criteria when it receives a DS_n (within an STS or VT SPE) that was mapped at any synchronizer that conforms to the bit stuffing jitter criteria in [Section 3.4](#).


Note that the criteria in this section are not sufficient to ensure that the desynchronizers in SONET NEs will provide adequate overall jitter attenuation. Specifically, attenuation of the jitter arising from decoded pointer adjustments places more stringent criteria on the desynchronizer characteristics (see [Sections 5.6.2.3.2 through 5.6.2.3.5](#)).

R5-252 [339] For Category I DS1 and DS3 interfaces, the jitter transfer function shall be under the mask in [Figure 5-27](#).

³³ Since the bit stuffing mechanism is a form of sampling, input jitter above a certain frequency will be aliased to a lower frequency. For example, in the DS1 mapping, the stuff opportunity bits occur every 0.5 ms, and therefore input jitter at frequencies greater than approximately 1 kHz will be aliased.

R5-253 [338] For Category I DS_n interfaces other than DS1 and DS3 (e.g., asynchronous DS1C and DS2 interfaces), the jitter transfer function shall meet the jitter transfer requirement in GR-499-CORE, Section 7.3.2.

Figure 5-27 Category I DS1 and DS3 Jitter Transfer Mask



Phase smoothing circuits are also needed in cases where a DS_n is demultiplexed from an STS or VT SPE and then multiplexed into a higher bit-rate DS_n or another STS or VT SPE within a single SONET NE (e.g., within a DCS that performs DS1 level cross-connections between two STS-1 electrical interfaces), even though the lower bit rate or asynchronous DS_n payload does not appear at an external (Category I) interface. These circuits are necessary to filter bit stuffing jitter, intrinsic payload mapping jitter and pointer adjustment jitter. All of these types of jitter appear at the output of the desynchronizer and, if not filtered, can appear as bit stuffing jitter when the DS_n is multiplexed into a new STS or VT SPE.

R5-254 [340v2] Phase smoothing circuits shall be employed when an asynchronous DS_n payload is demultiplexed from an STS or VT SPE and then multiplexed into a higher bit-rate DS_n or another STS or VT SPE within a single SONET NE.

Given the design of many desynchronizers, the most critical type of jitter for an NE to filter is expected to be pointer adjustment jitter. Therefore, more specific criteria on the phase smoothing characteristics of the circuits required by **R5-254 [340v2]** appear after the pointer adjustment jitter generation criteria, in [Section 5.6.2.3.7](#).


5.6.2.1.2 Category II Jitter Transfer

The jitter transfer requirement in this section limits the amount of jitter on an input OC-N or STS-N electrical signal that can be transferred to the output OC-N or STS-N electrical signal. As noted in [Section 5.6.2.1](#), jitter transfer is primarily of concern

within two decades or less of the breakpoint in the jitter transfer mask. This is reflected in the values of f_L and f_H that appear in Figure 5-28 and limit the jitter frequency range over which the following requirement applies. In particular, each value of f_L is a factor of 100 less than the breakpoint frequency (f_c), and each value of f_H is either 10 or 80 times larger than f_c (and in all cases is equal to B_3 from the network interface jitter limits given in Section 5.6.1).

R5-255 [341] For Category II interfaces, the jitter transfer function shall be under the mask in Figure 5-28.

Figure 5-28 Category II Jitter Transfer Mask



OC-N/STS-N Level	f_L (kHz)	f_c (kHz)	f_H (MHz)	P (dB)
1	0.4	40	0.4	0.1
3	1.3	130	1.3	0.1
12	5	500	5	0.1
48	20	2000	20	0.1
192	10	1000	80	0.1
768 ^a	40	4000	320	0.1

Note:

- a. The values in this row should be considered preliminary. See the discussion in Section 5.6.

5.6.2.2 Jitter Tolerance

5.6.2.2.1 Category I Jitter Tolerance

For Category I interfaces to SONET NEs, the definition of input jitter tolerance given in Section 7.3.1 of GR-499-CORE is the applicable definition.

R5-256 [342] Category I interfaces to SONET NEs shall meet the Category I input jitter tolerance requirement in Section 7.3.1 of GR-499-CORE.

5.6.2.2.2 *Category II Jitter Tolerance*


For OC-N interfaces, jitter tolerance is defined as the maximum peak-to-peak amplitude of sinusoidal jitter applied on the input OC-N signal that causes no more than a 1-dB power penalty and does not cause any failures to be declared. This is a stress test intended to ensure that no additional penalty is incurred under operating conditions. For STS-N electrical interfaces, application of the above definition is not feasible, and the jitter tolerance definition provided in Section 7 of GR-499-CORE is used instead.

R5-257 [343v4] A SONET NE's STS-N electrical and OC-N interfaces, with the exception of OC-48 interfaces that are specified as having "reduced" jitter tolerance (see the discussion below), shall tolerate, as a minimum, input jitter applied according to the mask in [Figure 5-29](#) (with the parameter values specified in the figure for the particular rate signal).

Note that for jitter frequencies below f_4 in [Figure 5-29](#), ANSI T1.105.03 indicates that the phrase "causes no more than a 1-dB power penalty" in the OC-N jitter tolerance definition provided above is supposed to be replaced by the phrase "causes no errors". However, that substitution can cause the results of the low-frequency portion of a jitter tolerance test to be dependent upon two unspecified variables - the optical power level at which the test is conducted, and the dwell time at each jitter frequency/amplitude combination. [In particular, if the optical power level of the test signal is relatively low and the dwell time is relatively long, the jitter tolerance of the receiver may appear to be low as a result of the normal (non-zero) BER associated with that power level.] By utilizing the "no more than 1-dB power penalty" definition over the entire range of jitter frequencies, the optical power level and dwell time are specified [i.e., the power level is set to 1 dB more than the level that results in a BER of 1×10^{-10} or 1×10^{-12} (depending on the application, see [Section 4.2.4.1](#)) with no jitter applied, and the dwell time is controlled by the time required to determine (with some appropriate level of confidence) whether the BER is better than, worse than, or essentially equal to the baseline BER of 1×10^{-10} or 1×10^{-12}]. In addition, it is the view of Telcordia that the error bursts caused by buffer overflow and underflow events (which are more likely to occur in the presence of large-amplitude, low-frequency jitter and are apparently the events that the revised definition in T1.105.03 was intended to address) will be adequately reflected in the measured BER, and therefore in the apparent power penalty.³⁴

Also note that in several cases the parameter values that define the mask in [Figure 5-29](#) are different than those that appeared in earlier issues of this document (i.e., those prior to Issue 4). These changes were made to better align this document with the new or revised jitter tolerance specifications that appeared in the 2003 version of ANSI T1.105.03 (OC-1 through OC-192) and ITU-T Rec. G.8251 (OTU3, which is expected to be used as the basis for future STM-256/OC-768 specifications).

Figure 5-29 SONET Category II Jitter Tolerance Mask



OC-N/ STS-N Level	f ₀ (Hz)	f ₁ (Hz)	f ₂ (Hz)	f ₃ ^a (Hz)	f ₄ (kHz)	f ₅ ^a (kHz)	f ₆ ^a (MHz)	A ₄ (UI _{pp})	A ₃ (UI _{pp})	A ₂ ^a (UI _{pp})	A ₁ ^a (UI _{pp})
1	10	NA	41.3	100	2	20	0.4	NA	3.63	1.5	0.15
3	10	NA	68.7	500	6.5	65	1.3	NA	10.9	1.5	0.15
12	10	18.5	100	1000	25	250	5.0	27.8	15	1.5	0.15
48	10	70.9	500	5000	100	1000	20	106.4	15	1.5	0.15
192	10	296	2000	20000	400	4000	80	444.6	15	1.5	0.15
768 ^b	10	1184	8000	20000	400	16000	320	1776	15	6.0	0.15

Notes:

- a. With the exception of B_2 for the OC-48(B) interface listed in Table 5-11, the values of f_3, f_5, f_6, A_1 and A_2 provided here are identical to the values of B_1, B_2, B_3, A_2 and A_1 provided in that table.
- b. The values in this row should be considered preliminary. See the discussion above and in Section 5.6.

34. If the applied jitter at a particular frequency primarily causes buffer overflow or underflow events (rather than individual bit errors), it may not be possible to set the amplitude of that jitter such that the measured BER is equal to 1×10^{-10} or 1×10^{-12} . However, it should be possible to find the maximum amplitude that does not result in overflow/underflow events (and thus results in a BER that is clearly better than 1×10^{-10} or 1×10^{-12} due to the 1 dB of additional optical power that is provided to the receiver under test, compared to the no-jitter case). Any increase in the jitter amplitude will then cause overflow/underflow events and (typically) a BER that is clearly worse than 1×10^{-10} or 1×10^{-12} . In that case, the maximum amplitude that does not result in overflow/underflow events is also the maximum amplitude that “causes no more than a 1-dB power penalty”, and thus is the amplitude of jitter tolerated by the receiving NE (at that jitter frequency).

SONET NEs with OC-48 interfaces that meet **R5-257 [343v4]** are expected to be capable of interworking with all other NEs (in particular, with all SONET regenerators) that meet the OC-48 jitter transfer requirement in [Section 5.6.2.1.2](#). However, in some applications that interworking capability may not be necessary, and the user may only require the NE to meet a reduced jitter tolerance specification (i.e., the NE may not be required to meet **R5-257 [343v4]**). In particular, if Type B regenerators (which have tighter jitter transfer characteristics than those required in [Section 5.6.2.1.2](#)) are used in a SONET line system, then less jitter is transferred through each regenerator. Therefore, the regenerators and the associated terminal interfaces are subjected to less input jitter, and they can operate properly even if they have reduced jitter tolerance.

R5-258 [931] If a SONET NE has reduced OC-48 jitter tolerance, that shall be clearly documented.

Annex A of ANSI T1.105.03 contains the applicable jitter tolerance specification for an NE with reduced OC-48 jitter tolerance capabilities. In addition, it contains the jitter transfer specification for Type B regenerators, and discusses how NEs with reduced jitter tolerance may interwork with other equipment subject to certain application constraints.

For Category II DS1 interfaces to SONET NEs, the definition of input jitter tolerance given in Section 7.3.1 of GR-499-CORE is the applicable definition.

R5-259 [345] Category II DS1 interfaces to SONET NEs shall meet the Category II input jitter tolerance requirement in Section 7.3.1 of GR-499-CORE.

5.6.2.3 Jitter Generation

Jitter generation is defined in Section 7.3.3 of GR-499-CORE. In SONET, jitter generation criteria exist for both Category I and Category II interfaces. For Category I interfaces the jitter generation criteria are divided into two areas. [Section 5.6.2.3.1](#) contains the criteria in the area of “mapping” jitter generation, while [Sections 5.6.2.3.2](#) through [5.6.2.3.5](#) contain the criteria in the area of “pointer adjustment” jitter generation. In addition, [Section 5.6.2.3.6](#) contains the Category II jitter generation requirement, and [Section 5.6.2.3.7](#) contains “pointer adjustment jitter to bit stuffing jitter conversion” criteria.

All of the jitter generation tests are performed with no jitter or wander applied at any applicable input (i.e., at any input where the timing characteristics of the incoming signal might affect the jitter generated on the output signal of interest).³⁵ In addition,

³⁵Note that in cases where the timing of the output signal is synchronized to the NE’s internal clock, the (Category II) jitter generation requirement applies when that clock is locked to a jitter-free and wander-free reference, and also when it is operating in its free-run or holdover mode. Also note that in cases where there is an “applicable input”, the jitter generation criteria apply when the bit rate of the input signal is offset from nominal by any value within the specified or appropriate limits (e.g., within the required pull-in range of the NE’s clock for an OC-N signal being generated by an externally timed NE).

a bandpass filter is used to limit the jitter generation measurements to the jitter frequency range of interest. For DS_n interfaces (both Category I and Category II), the bandpass filter has a 10-Hz high-pass cutoff frequency with a roll-off of 20 dB/decade, and a low-pass cutoff frequency of F_4 and a roll-off of 20 or 60 dB/decade (as listed in Table 7-1 of GR-499-CORE).³⁶ For Category II STS-N electrical and OC-N interfaces, the bandpass filter has the cutoff frequencies listed in Table 5-16, and roll-offs of 20 dB/decade (high-pass) and 60 dB/decade (low-pass). In addition, the specified characteristics of the low-pass portion of the filter are those of a maximally flat Butterworth filter.

5.6.2.3.1 Category I Mapping Jitter

Mapping jitter is the sum of the intrinsic payload mapping jitter and the jitter that is generated as a result of the bit stuffing mechanisms used in all of the asynchronous DS_n mappings defined in SONET. Section 3.4 contains bit stuffing jitter criteria (for some of the asynchronous DS_n mappings) that are specifically applicable to the NE that maps the DS_n into a VT or STS SPE, while this section contains criteria that are applicable to a SONET system (e.g., for a signal that is mapped into an SPE at one NE and extracted from that SPE at another NE). To measure the mapping jitter it is necessary for the system to be configured so that no STS or VT pointer adjustments occur during the tests. In a single-product test, this can be accomplished by looping back the OC-N or STS-N electrical signal. In addition, it is necessary to measure the mapping jitter using a variety of DS_n bit rates that meet the DS_n interface criteria in Section 9 of GR-499-CORE.

For some DS_n interfaces (e.g., DS1C and DS2 interfaces), the Category I jitter generation requirement in Section 7.3.3 of GR-499-CORE is currently the applicable requirement. However, tighter requirements have been developed for asynchronous DS1 and DS3 interfaces to SONET NEs, due to concerns about jitter accumulation through SONET islands. In addition, for interoperability in a multi-supplier network, a desynchronizer must meet the mapping jitter generation requirements when it receives a DS_n (within an STS or VT SPE) that was mapped at any synchronizer that conforms to the bit stuffing jitter criteria in Section 3.4.

R5-260 [346] For Category I DS_n interfaces other than DS1 and DS3 interfaces, the mapping jitter shall meet the jitter generation requirement in Section 7.3.3 of GR-499-CORE.

³⁶ For DS1 interfaces, F_4 is equal to 40 kHz, and for DS3 interfaces it is 400 kHz. Thus, the measurement filters for the DS_n jitter generation criteria provided here are the same as those specified for use in wideband network interface jitter measurements in GR-499-CORE (for which the maximum allowable amplitude for both DS1 and DS3 interfaces is 5.0 UI_{pp}). Note that although a second set of jitter generation requirements and measurements utilizing the highband network interface jitter measurement filters from GR-499-CORE is not provided, the network interface jitter requirement in that document applies anywhere that a DS_n signal appears in a network. This includes an output port on a SONET NE, and therefore that requirement needs to be met under the various conditions specified in the following sections.

- R5-261 [347]** For a Category I DS1 or DS3 interface, the mapping jitter generation shall be less than the value given in [Table 5-12](#).

Table 5-12 Category I Mapping Jitter Limits

Interface	Jitter (UI _{pp})
DS1	0.70
DS3	0.40

5.6.2.3.2 *Category I Jitter Generation Due to Pointer Adjustments, General*

Pointer adjustment jitter criteria are necessary to ensure acceptable DS_n jitter performance during synchronization rearrangements or failure conditions, and for DS_n signals that traverse SONET islands. These criteria are tested using the “pointer test sequences” described by Figures 5-30 through 5-35 in the following sections.

To thoroughly test an NE it is necessary to perform the test sequences a number of times (i.e., using pointer increments, using pointer decrements, using a variety of values for the time variable “T”, and using a variety of DS_n bit rates). [Table 5-13](#) gives the limits for T, and along with the notes that follow it, is considered an integral part of the pointer test sequence specifications.

During the pointer adjustment jitter tests, it is important that no buffer spills occur in either the pointer processor circuits or the desynchronizer circuits. The following requirement applies during the single pointer adjustment tests in [Section 5.6.2.3.3](#), the burst pointer adjustment tests in [Section 5.6.2.3.4](#), and the periodic pointer adjustment tests in [Section 5.6.2.3.5](#) that simulate frequency offsets of up to 4.6 ppm in a SONET network. The objective applies during the periodic pointer adjustment tests ([Section 5.6.2.3.5](#)) that simulate frequency offsets from 4.6 ppm to 20 ppm. In addition, both criteria are applicable during all of the pointer adjustment tests in [Section 5.6.2.3.7](#).

- R5-262 [348]** Complete data integrity shall be maintained (i.e., no bit errors shall occur) through the SONET system during all of the pointer adjustment jitter generation tests where T is either in the “required range” given in [Table 5-13](#), or is not applicable (i.e., during the single and burst pointer adjustment tests).
- O5-263 [349]** Complete data integrity should be maintained through the SONET system during all of the periodic pointer test sequences where T is in the “objective range” given in [Table 5-13](#).

Table 5-13 Pointer Test Sequence Parameters

SPE/Payload	t (for burst and periodic tests)	T (Required range for periodic tests)	T (Objective range for periodic tests)
VT1.5/DS1	2 ms	1 s < T < 10 s	0.2 s < T < 1 s
STS-1/DS3	0.5 ms	34 ms < T < 10 s	7.5 ms < T < 34 ms

Pointer Test Sequence Notes:

- Annex G of ANSI T1.105.03 contains measurement methodology information that is useful for these tests.
- Each test sequence consists of an “initialization” period, a “cool down” period, and a “measurement” period. During the initialization period, a series of pointer adjustments of the same polarity as will be used in the measurement period is applied to fill or deplete any buffers provided by the NE under test.³⁷ In the case of tests that utilize a periodic sequence of pointer adjustments, these initialization pointer adjustments may be applied in the same pattern as will be used in the rest of the test sequence (as shown in Figures 5-33 through 5-35), or a different pattern. In any case, a cool down period of at least 30 seconds is then provided to allow the NE to reach its stable operating condition. Finally, the generated jitter is measured during the measurement period.
- For STS payloads (e.g., DS3), the pointer adjustments are applied to the STS pointers. For VT payloads (e.g., DS1), the pointer adjustments are applied to the VT pointers.
- For all of the sequences, the tests must be performed using a variety of DS_n bit rates that meet the DS_n interface criteria in Section 9 of GR-499-CORE.
- All tests must be run with all positive pointer adjustments, and also with all negative pointer adjustments.
- For periodic sequences, T is constant for each measurement and is determined by the frequency offset between the SPE and its carrier (e.g., the OC-N or STS-N electrical signal in the case of DS3 within an STS-1 SPE). The value of T must be varied over the ranges given in Table 5-13.
- For periodic sequences with added and canceled pointer adjustments, the tests must be run with only added pointer adjustments, and also with only canceled pointer adjustments.

³⁷ Most SONET NEs are expected to include a buffer to absorb many of the pointer adjustments resulting from wander that could occur in the network. For those NEs, pointer adjustment jitter is expected to occur only when the buffer is full (or depleted) and then a pointer decrement (increment) occurs. If the buffer is filled by a series of decrements and then an increment occurs, the buffer would absorb that increment (and also the next decrement since it would no longer be full), and no pointer adjustment jitter would occur.

Although the pointer adjustment patterns described in the following sections are patterns that have been observed in the network and are considered to be relatively likely to occur, a variety of other patterns have also been observed. While some of those patterns (e.g., “dithered” pointer adjustments) have resulted in improved system jitter performance (compared to the patterns described in this document) in both single-product and multi-product configurations, others have resulted in significant performance degradations. Specifically, degradations have occurred when some of the pointer adjustments generated by a pointer processor have been much more closely spaced (in time) than the average spacing.³⁸ The following requirements are intended to limit such degradations in two specific steady-state situations. They are not intended to limit the operations that a pointer processor may need to perform to accommodate such things as clock phase transients or wander, network configurations with multiple frequency offsets, or incoming pointer adjustment bursts.

R5-264 [1025] In the presence of a constant frequency offset (between the incoming and outgoing signals at a pointer processor) of up to ± 20 ppm and no pointer adjustments in the incoming signal, the minimum spacing (time) between consecutive pointer adjustments generated by a pointer processor shall be greater than or equal to 0.5 times the long-term average spacing.

R5-265 [1026] In the presence of incoming pointer adjustments in the patterns shown in [Figure 5-33b](#) (VT only), [Figure 5-34b](#) (STS only) or [Figure 5-35b](#) (STS and VT) and no frequency offset between the incoming and outgoing signals, the minimum spacing between consecutive pointer adjustments generated by a pointer processor shall be greater than or equal to 0.5 times the long-term average spacing.

For the preceding requirements, the time needed to determine the “long-term average” is not specified. However, it generally should be sufficient to monitor the generated pointer adjustments through a complete “cycle” (e.g., 783 STS pointer adjustments, or 104 VT1.5 pointer adjustments), and for frequency offsets and incoming pointer adjustment rates that are large enough that the resulting generated pointer adjustments cannot be considered to be single pointer adjustments (see [Section 5.6.2.3.3](#)).

5.6.2.3.3 *Category I Jitter Generation Due to Single Pointer Adjustments*

Single pointer adjustments are those that may occur due to wander in the SONET network or small timing reference frequency offsets. For testing purposes it is assumed that the jitter caused by one pointer adjustment will not be affected by the preceding or following pointer adjustments if all of the adjustments are separated by 30 seconds or more.

³⁸For example, in one case a pointer processor that was receiving evenly spaced pointer adjustments on a through STS-1 path was observed to buffer every other incoming adjustment until it received the next adjustment. At that time it would generate two very closely spaced pointer adjustments (e.g., separated by 4 frames), resulting in a pattern of double pointer adjustment bursts at approximately twice the spacing of the original incoming single adjustments.

R5-266 [350] The jitter appearing at a Category I DS_n interface shall be less than the corresponding value in [Table 5-14](#) when the pointer test sequence in [Figure 5-30](#) is applied.


Table 5-14 Jitter Due to Single Pointer Adjustments

SPE/Payload	Jitter (U_{Ipp})
VT1.5/DS1	$A_0 + 0.60^1$
VT2/DS1A	Under Study
VT3/DS1C	Under Study
VT6/DS2	Under Study
STS-1/DS3	$A_0 + 0.30^1$
STS-3c/DS4NA	Under Study

Note:

1 A_0 is the mapping jitter generated by the NE under test.

Figure 5-30 Single Pointer Adjustment Test Sequence



5.6.2.3.4 Category I Jitter Generation Due to Bursts of Pointer Adjustments

Jitter generation criteria for bursts of pointer adjustments are currently specified only for DS3 interfaces. Closely spaced pointer adjustments for lower bit-rate signals (e.g., DS1 signals) are not expected to occur due to the larger amount of phase movement necessary to generate VT pointer adjustments. In addition, jitter generation criteria for bursts of pointer adjustments for DS4NA payloads is not considered to be necessary.

The pointer adjustment burst sequences specified in the following criteria are intended to simulate the expected worst-case pointer activity due to phase transients caused by synchronization rearrangements (see [Section 5.4.4.3.5](#)).

R5-267 [351] The jitter appearing at a Category I DS3 interface shall be less than $1.3 U_{Ipp}$ when the pointer test sequence described in [Figure 5-31](#) is applied.

R5-268 [352] The jitter appearing at a Category I DS3 interface shall be less than $1.2 U_{Ipp}$ when the pointer test sequence described in Figure 5-32 is applied.

Figure 5-31 Maximum Rate Pointer Burst Test Sequence





Figure 5-32 Phase Transient Pointer Burst Test Sequence



5.6.2.3.5 Category I Jitter Generation Due to Periodic Pointer Adjustments

The pointer test sequences in this section are intended to simulate the pointer adjustments that could occur in a SONET network in the presence of timing reference offsets. The test sequences specified in the criteria simulate the adjustments that could occur for frequency offsets up to 4.6 ppm (i.e., the “required range” for T shown in [Table 5-13](#)), while those in the objectives cover offsets from 4.6 ppm to 20 ppm (i.e., the “objective range”). Currently, criteria exist only for DS1 and DS3 payloads.

R5-269 [353] The jitter appearing at a Category I DS1 or DS3 interface shall be less than the corresponding value given in [Table 5-15](#) when the pointer test sequences described in [Figures 5-33](#), [5-34](#) and [5-35](#) are applied with T in the required range.

O5-270 [354] The jitter appearing at a Category I DS1 or DS3 interface should be less than the corresponding value given in [Table 5-15](#) when the pointer test sequences described in [Figures 5-33](#), [5-34](#) and [5-35](#) are applied with T in the objective range.

In general, the 26-1 pointer adjustment sequence shown in [Figure 5-33b](#) (or a very similar sequence) would be expected to be generated by a VT1.5 pointer processor that determines whether or not a pointer adjustment needs to be performed based on a fixed threshold that is equal to the nominal time difference between consecutive bytes of the VT1.5 superframe within the SONET signal (i.e., $500 \mu\text{s}/108 \text{ bytes} = 4.63 \mu\text{s}/\text{byte}$). Similarly, the 87-3 sequence in [Figure 5-34b](#) would be expected to be generated by an STS pointer processor that utilizes a fixed threshold equal to the time difference between consecutive bytes of the STS-1 frame within the SONET signal [i.e., $125 \mu\text{s}/(9 \text{ rows} \times 90 \text{ bytes}/\text{row}) = 0.154 \mu\text{s}/\text{byte}$]. Neither of these patterns would be expected to be generated by the other type of pointer processor (i.e., the 26-1 pattern will not be generated by an STS pointer processor, and the 87-3 pattern will not be generated by a VT1.5 pointer processor), or by a VT2, VT3 or VT6 pointer processor (which might instead generate pointer adjustments in a 35-1, 53-1 or 107-1 pattern). The reason for this is that the “1” portion of the 26-1 pattern occurs due to the inability of the VT pointer to “point” to the byte positions within the VT superframe occupied by the V1, V2, V3 and V4 bytes, while the “3” portion of the 87-3 pattern occurs due to the inability of the STS pointer to point to the transport overhead bytes associated with that STS in the SONET frame. In the VT1.5 case, the V1 through V4 bytes are separated by 26 bytes at which the pointer can point, while in the STS case the transport overhead bytes occur in groups of 3 and are separated by 87 bytes at which the pointer can point. As a result of this, an NE that generates the 26-1 pattern will always generate exactly 26 pointer adjustments in between each gap of 1, and an NE that generates the 87-3 pattern will always generate exactly 87 pointer adjustments in between each gap of 3, independent of whether any of those adjustments occur much sooner or later than the average. This is reflected in the preferred values shown in [Figures 5-33c](#) and [5-34c](#) for the “added” pointer adjustment cases, and in [Figures 5-33d](#) and [5-34d](#) for the “cancelled” pointer adjustment cases. On the other hand, several early versions of the SONET standards and criteria specified a total of 27 (VT1.5) or 88 (STS) pointer adjustments between gaps in the “added” case, and 25 (VT1.5) or 86 (STS) pointer adjustments between a normal gap and the longer gap in the “cancelled” case. These differences are not expected to have a significant impact on the amount

of jitter that is generated by the receiving NE, and therefore the values shown in parentheses in the figures are also considered acceptable in testing an NE against the criteria.

In contrast to the cases described above, the “continuous” pattern shown in Figure 5-35 can be generated by both VT and STS pointer processors. In particular, that pattern would be expected to be generated by a VT or STS pointer processor that utilizes a fixed threshold equal to the average time difference between consecutive bytes of VT or STS SPE within the SONET signal (as opposed to consecutive bytes of the VT superframe or STS frame within the SONET signal). In effect, such pointer processors can be viewed to be compensating for the fact that the pointer cannot point to certain byte positions (as described above). This results in a slightly larger threshold of 4.81 μ s for a VT1.5 (i.e., 500 μ s/104 bytes), or 0.160 μ s for an STS [i.e., 125 μ s/(9 rows \times 87 bytes/row)].

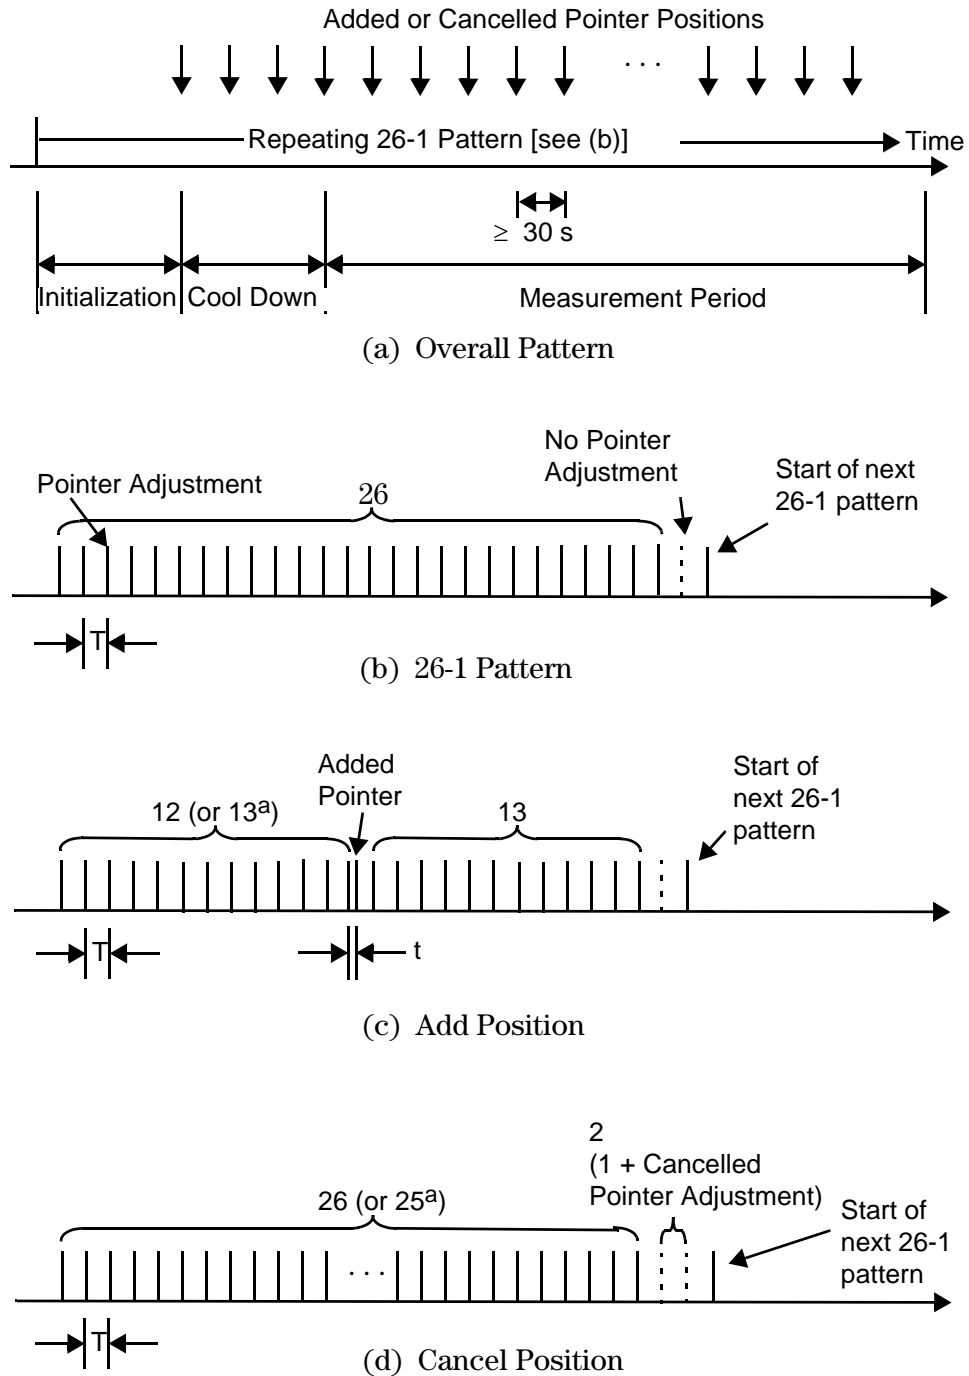
Table 5-15 Jitter Generation Limits for Periodic Pointer Adjustment Sequences

SPE/Payload	Figure Number of Test Sequence					
	Figure 5-33b	Figures 5-33c and 5-33d	Figure 5-34b	Figures 5-34c and 5-34d	Figure 5-35b	Figures 5-35c and 5-35d
VT1.5/DS1	1.3 UI _{pp}	1.9 UI _{pp}	-	-	1.3 UI _{pp}	1.9 UI _{pp}
STS-1/DS3	-	-	1.0 UI _{pp}	1.3 UI _{pp}	1.0 UI _{pp}	1.3 UI _{pp}

It has been observed that some SONET NEs generate pointer adjustments in patterns that do not match any of the pointer test sequences. The following criteria are included (in addition to R5-264 [1025] and R5-265 [1026]) to ensure acceptable jitter performance for those NEs in single-supplier configurations.

- R5-271 [355]** In a single-supplier configuration with a single timing reference offset of 0 to ± 4.6 ppm, the jitter appearing at a Category I DS1 or DS3 interface shall be less than 1.5 UI_{pp}.
- O5-272 [356]** In a single-supplier configuration with timing reference offsets equal to twice the specified free-run accuracy of the NEs’ internal clocks (e.g., 9.2 ppm for NEs with stratum 3 clocks, 40 ppm for NEs with SMCs), the jitter appearing at a Category I DS1 or DS3 interface should be less than 1.5 UI_{pp}.

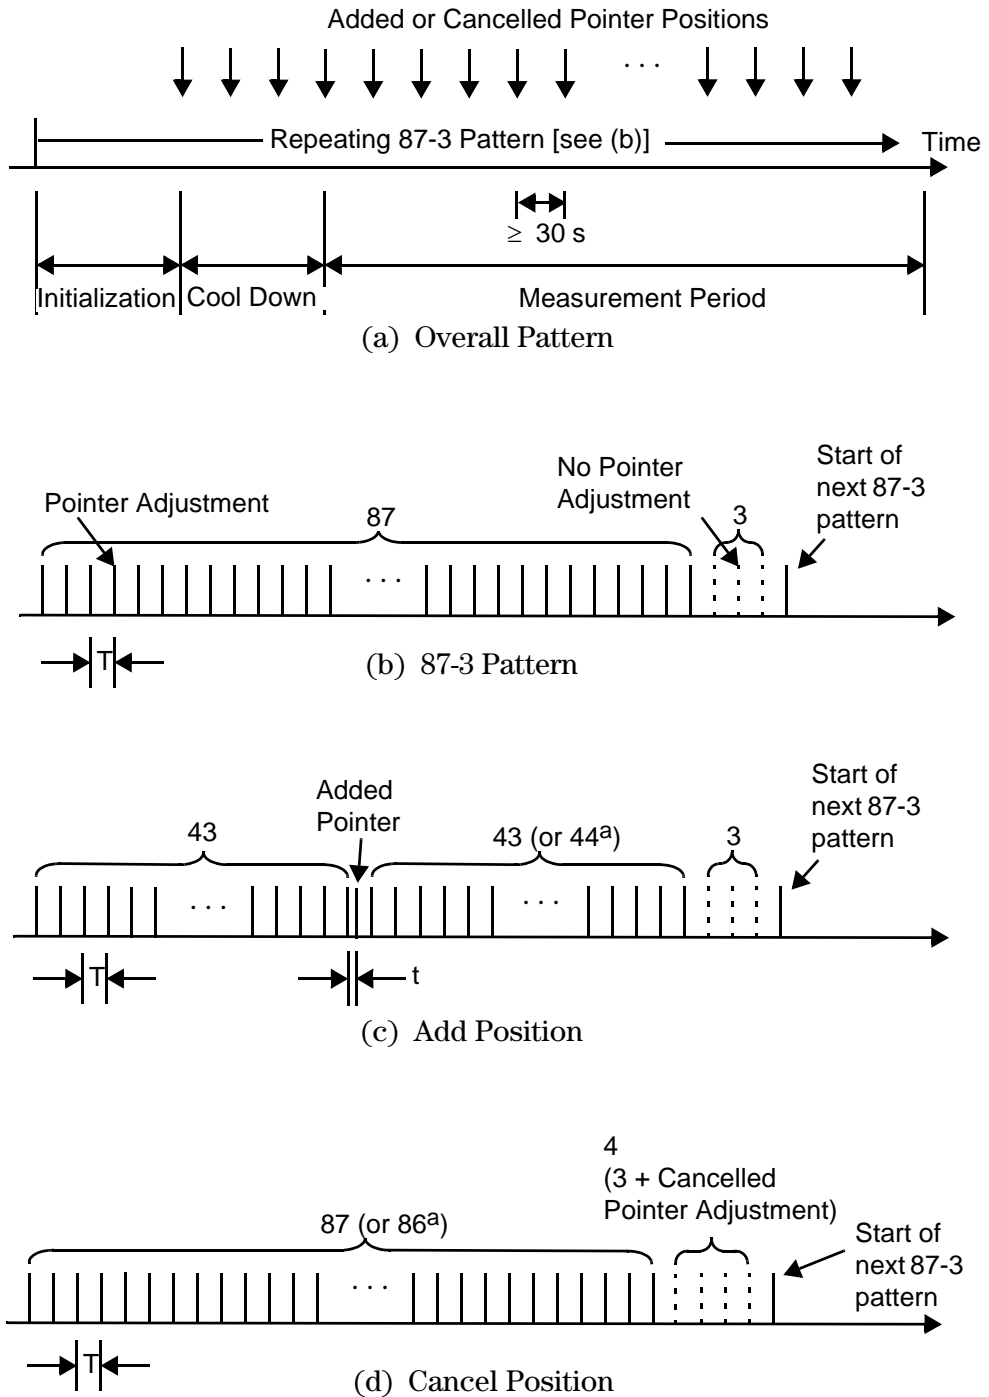
Figure 5-33 Periodic VT1.5 Pointer Adjustment Test Sequence (26-1 Pattern)



Note:

a The first value listed is preferred. See the discussion following **O5-270 [354]**.

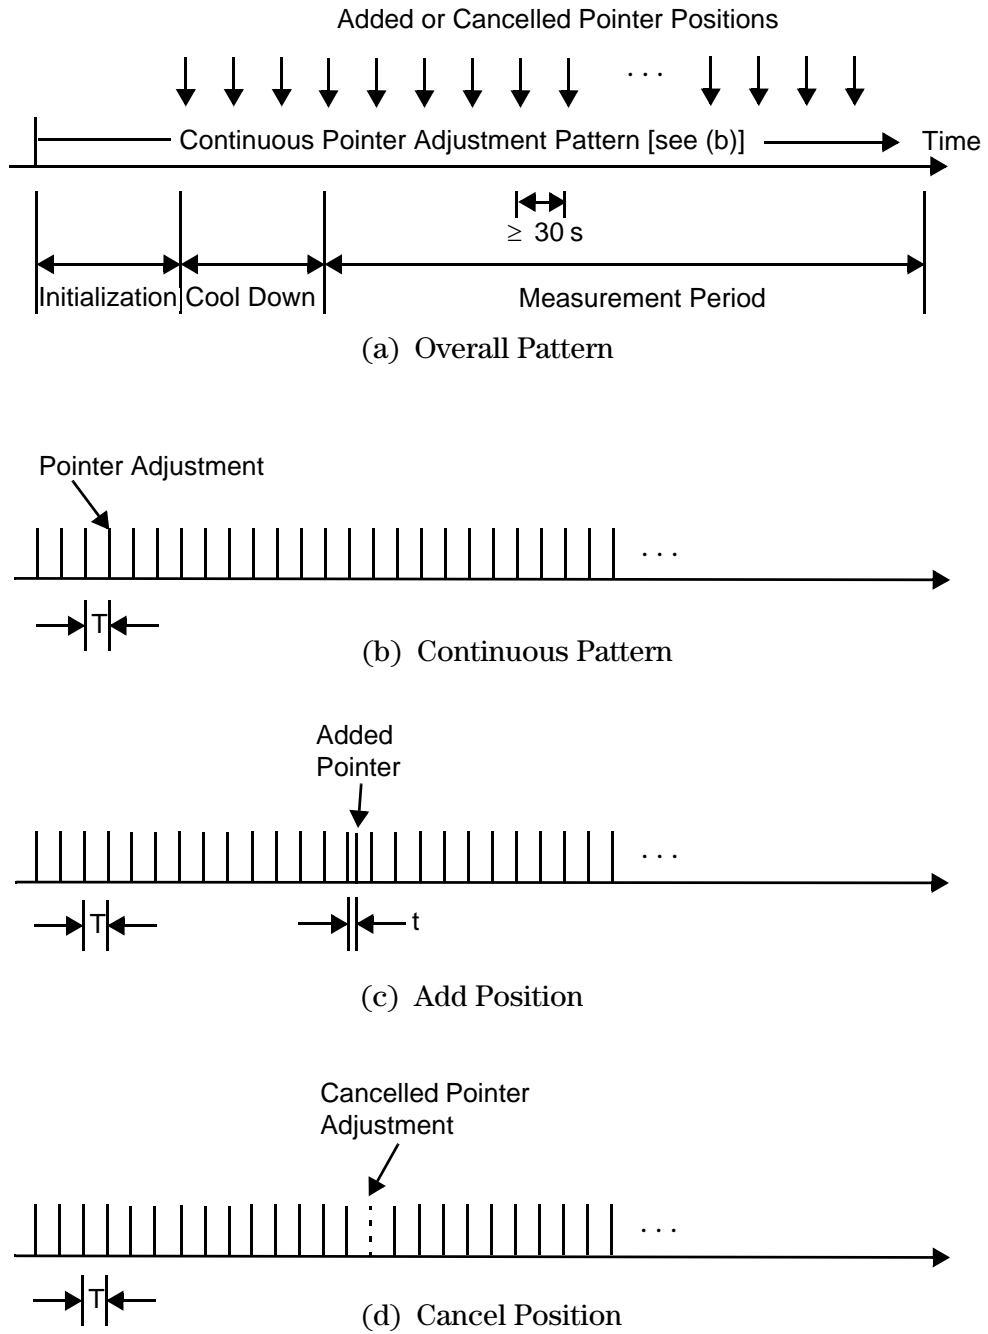
Figure 5-34 Periodic STS Pointer Adjustment Test Sequence (87-3 Pattern)



Note:

a The first value listed is preferred. See the discussion following **O5-270 [354]**.

Figure 5-35 Periodic Pointer Adjustment Test Sequence (Continuous Pattern)



5.6.2.3.6 Category II Jitter Generation

R5-273 [357v2] The jitter generated at non-SONET Category II interfaces shall be less than $0.01 U_{I_{rms}}$, and shall also be less than $0.10 U_{I_{pp}}$.³⁹ At SONET Category II interfaces, the generated jitter shall be less than the values listed in [Table 5-16](#).

See the discussion in [Section 5.6.2.3](#) for additional information on jitter generation measurements and the filters to be used in those measurements.

Table 5-16 SONET Jitter Generation

OC-N/STS-N Level	High-Pass Filter Cutoff	Low-Pass Filter Cutoff (B ₃)	Jitter Generation Limit
1	12 kHz	400 kHz	$0.1 U_{I_{pp}}$ and $0.01 U_{I_{rms}}$
3	12 kHz	1.3 MHz	$0.1 U_{I_{pp}}$ and $0.01 U_{I_{rms}}$
12	12 kHz	5 MHz	$0.1 U_{I_{pp}}$ and $0.01 U_{I_{rms}}$
48	12 kHz	20 MHz	$0.1 U_{I_{pp}}$ and $0.01 U_{I_{rms}}$
192	20 kHz	80 MHz	$0.3 U_{I_{pp}}$
	4 MHz		$0.1 U_{I_{pp}}$
768 ^a	20 kHz	320 MHz	$1.2 U_{I_{pp}}$
	16 MHz		$0.1 U_{I_{pp}}$

Note:

- a. The values in this row should be considered preliminary. See the discussion in [Section 5.6](#).

5.6.2.3.7 Pointer Adjustment Jitter to Bit Stuffing Jitter Conversion

As discussed in [Section 5.6.2.1.1](#), phase smoothing circuits are needed in cases where a single SONET NE demultiplexes a DS_n from an STS or VT SPE and then multiplexes that DS_n into a higher bit-rate DS_n or another STS or VT SPE. These

³⁹Note that for Category II DS1 interfaces, **R5-273 [357v2]** is less stringent than the corresponding specification for 1544 kb/s outputs that was added to ITU-T Rec. G.813 via Corrigendum 1 in June, 2005. That specification was copied directly from ITU-T Rec. G.812, *Timing requirements of slave clocks suitable for use as node clocks in synchronization networks*, which limits the jitter that can be generated on 1544 kb/s timing signals (as opposed to traffic signals) and is consistent with the jitter generation requirement applicable to external DS1 timing signals in GR-1244-CORE. No technical justification was provided to support a $0.05 U_{I_{pp}}$ limit for traffic signals, and a product that meets the new specification would also meet **R5-273 [357v2]** (which first appeared in a predecessor to GR-253-CORE in 1993). Therefore, there is currently no plan for this requirement to be changed to match the Rec. G.813 specification.

circuits are primarily needed to filter pointer adjustment jitter that is generated when the DS_n is demultiplexed, so that it does not appear as bit stuffing jitter when the DS_n is multiplexed into the higher bit-rate DS_n or new STS or VT SPE.

The following criteria are meant to limit the amount of jitter that an NE can convert (internally) from pointer adjustment jitter to bit stuffing jitter, to the same levels as would be applicable if the demultiplexing and multiplexing were performed by separate NEs (i.e., as would be applicable if the DS_n appeared as an external DS_n signal out of one NE, and then was mapped into a higher bit-rate DS_n or new STS or VT SPE by another NE). The criteria are given in terms of the jitter that would appear at the output of a “nominal” desynchronizer when the various pointer test sequences are input to the NE. For the cases where the DS_n is multiplexed into a new STS or VT SPE, the nominal desynchronizer is defined as one with filtering characteristics equal to the DS1 and DS3 jitter transfer mask shown in [Figure 5-27](#). Similarly for cases where the DS_n is multiplexed into a higher bit-rate DS_n, it is defined as one with filtering characteristics equal to the appropriate Category I-to-Category I (multiplexer/demultiplexer pair) jitter transfer mask from GR-499-CORE.

-
- R5-274** [932] The jitter appearing at the output of a “nominal” DS1 or DS3 desynchronizer shall be less than the corresponding value in [Table 5-14](#) when that desynchronizer receives a signal from the NE, and the input signal to the NE contains pointer adjustments in the pointer test sequence shown in [Figure 5-30](#).
- R5-275** [933] The jitter appearing at the output of a “nominal” DS3 desynchronizer shall be less than $1.3 U_{Ipp}$ when that desynchronizer receives a signal from the NE, and the input signal to the NE contains pointer adjustments in the pointer test sequence shown in [Figure 5-31](#).
- R5-276** [934] The jitter appearing at the output of a “nominal” DS3 desynchronizer shall be less than $1.2 U_{Ipp}$ when that desynchronizer receives a signal from the NE, and the input signal to the NE contains pointer adjustments in the pointer test sequence shown in [Figure 5-32](#).
- R5-277** [935] The jitter appearing at the output of a “nominal” DS1 or DS3 desynchronizer shall be less than the corresponding value given in [Table 5-15](#) when that desynchronizer receives a signal from the NE, and the input signal to the NE contains pointer adjustments in the pointer test sequences shown in [Figures 5-33](#), [5-34](#) and [5-35](#), with T in the required range.
- O5-278** [936] The jitter appearing at the output of a “nominal” DS1 or DS3 desynchronizer should be less than the corresponding value given in [Table 5-15](#) when that desynchronizer receives a signal from the NE, and the input signal to the NE contains pointer adjustments in the pointer test sequences shown in [Figures 5-33](#), [5-34](#) and [5-35](#), with T in the objective range.
-

5.6.2.4 Jitter Enhancement

Jitter enhancement is defined in Section 7.3.4 of GR-499-CORE, which also contains jitter enhancement criteria for non-SONET NEs. Those criteria help ensure acceptable network performance as jitter is accumulated on a signal that traverses a system consisting of multiple NEs; however, they are not readily translated into criteria for SONET NEs. The possibility of developing SONET jitter enhancement criteria has been considered (e.g., in GR-253-ILR, Issue ID 253-66), and at this time no such criteria are expected to be added to this document.

5.7 Phase Variations on Payload Signals

This section provides further criteria (in addition to the jitter criteria in [Section 5.6](#)) regarding phase variations on DS1 and DS3 payload signals transported on SONET networks, and is based on ANSI T1.105.03. In addition, Annex G in that document contains information related to measurement methodologies, and those methods are recommended for use in testing SONET NEs against these criteria.

DS_n signals carried on SONET networks will incur phase variations due to normal clock noise, any clock offsets (e.g., due to synchronization failures), and the bit-stuffing mechanisms used in the asynchronous DS_n to VT or STS SPE mappings. In general, these phase variations must be limited for interworking with existing equipment.

Previously, this section might have been labeled “Wander on Payload Signals.” However, the term “wander” is specifically defined to indicate phase variations with frequencies below 10 Hz, and some of the criteria in this section are not bounded by that frequency. Instead, the DS1 payload measurements are made using a 100-Hz first-order low-pass filter, and therefore the more generic term “phase variations” is used. (Note that the DS3 measurements are made using a 10-Hz first-order low-pass filter, and thus could be accurately considered “wander” measurements.)

In all of the phase variation tests, no jitter or wander is applied to the input DS1 or DS3 signal or to the timing references to the NEs. In addition, all of the results are given in terms of the MTIE (see [Section 5.4.1.2](#)).

5.7.1 Mapping Phase Variations

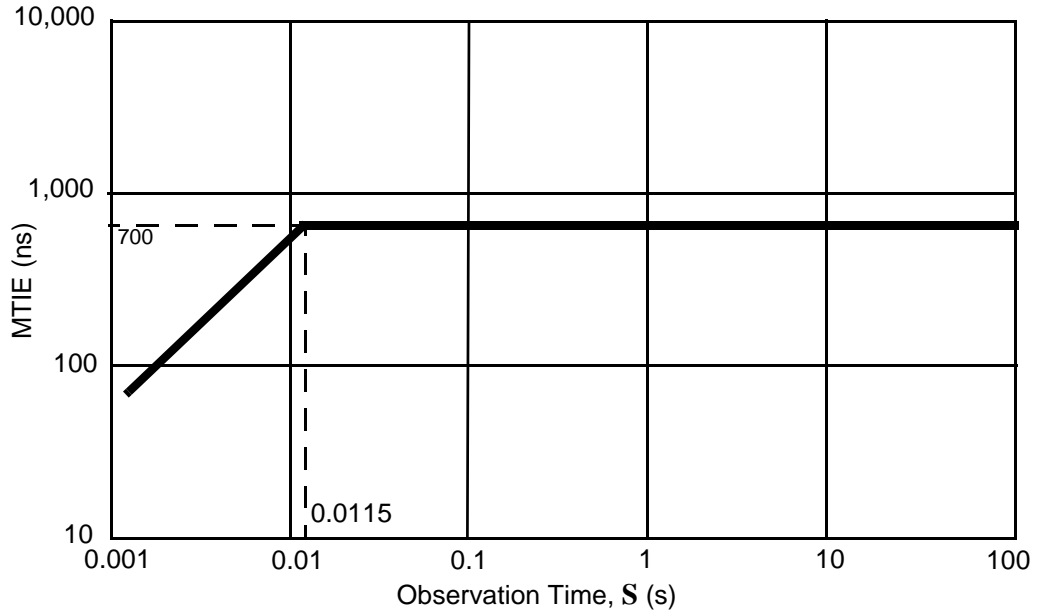
In the asynchronous DS1 to VT1.5 and DS3 to STS-1 mappings, a bit-stuffing mechanism is used to accommodate frequency differences between the DS1 or DS3 and the VT1.5 or STS-1 SPE. This mechanism causes phase variations on DS1 and DS3 signals carried on SONET networks, and these phase variations must be limited for interworking with other network equipment.

To measure the mapping phase variations, it is necessary for the system to be configured so that no STS or VT pointer adjustments occur during the tests. In a single-product test, this can be accomplished by looping back the OC-N or STS-N electrical signal. In addition, it is necessary to measure the mapping phase variations using a variety of DS1 and DS3 bit rates that meet the interface criteria in Section 9 of GR-499-CORE.

R5-279 [358] The mapping phase variations on a DS1 output from a SONET NE shall be below the mask in [Figure 5-36](#).⁴⁰

R5-280 [1027] The mapping phase variations on a DS3 output from a SONET NE shall be below the mask in [Figure 5-37](#).


Figure 5-36 DS1 Mapping Phase Variation Limits



Observation Time, S (seconds)	MTIE (nanoseconds)
$S < 0.001326$	N/A
$0.001326 \leq S \leq 0.0115$	$61000 \times S$
$0.0115 < S$	700

⁴⁰The equations for each of the MTIE masks in the figures in this section appear in tables in those figures.

Figure 5-37 DS3 Mapping Phase Variation Limits



Observation Time, S (seconds)	MTIE (nanoseconds)
$S < 0.1$	N/A
$0.1 \leq S \leq 0.25$	20
$0.25 < S \leq 1.0$	$7 + 53 \times S$
$1.0 < S \leq 10$	$23 + 37 \times S^{0.5}$
$10 < S \leq 100$	140
$100 < S$	N/A

5.7.2 Pointer Adjustment Phase Variations

The pointer adjustment activity in a SONET network is a function of the synchronization characteristics of that network. Clock noise exceeding the buffering capacity of the NEs results in pointer adjustments, which can cause significant phase variations on the DS_n payloads. Since the magnitude of a VT pointer adjustment is much larger than the magnitude of an STS pointer adjustment (e.g., one byte at the VT1.5 rate versus one byte at the STS-1 rate), VT pointer adjustments occur much less frequently but cause much larger phase variations.


Pointer adjustment statistics can vary widely; however, a set of standard test sequences was developed in T1X1.3 for use in testing the effects of pointer adjustments on the phase variations of DS1 and DS3 payloads. These sequences are a subset of the sequences used in the pointer adjustment jitter generation tests in Sections 5.6.2.3.3 through 5.6.2.3.5. As with the jitter tests, the pointer adjustment phase variation tests must be performed with both positive and negative pointer adjustments, and (for the periodic tests) “T” must be varied over the range given in

Table 5-13. Unlike the jitter tests, the mapping and pointer adjustment components of the phase variations are expected to be largely separable, and therefore the pointer adjustment tests need to be performed using only a single DS1 or DS3 bit rate (e.g., the bit rate that minimizes the mapping phase variations). The results of the mapping phase variation tests (at that bit rate) can then be subtracted from the cumulative results obtained in the pointer adjustment phase variation tests to determine the portion of the MTIE caused specifically by the pointer adjustments.

5.7.2.1 Single Pointer Adjustments


- R5-281 [359]** The MTIE of a DS1 output from a SONET NE shall be below the mask in [Figure 5-38](#) when the pointer adjustment test sequence in [Figure 5-30](#) is applied.
- R5-282 [1028]** The MTIE of a DS3 output from a SONET NE shall be below the mask in [Figure 5-39](#) when the pointer adjustment test sequence in [Figure 5-30](#) is applied.

Figure 5-38 Single VT Pointer Adjustment Phase Variation Limits



Observation Time, S (seconds)	MTIE (nanoseconds)
$S < 0.001326$	N/A
$0.001326 \leq S \leq 0.0164$	$61000 \times S$
$0.0164 < S \leq 0.93$	$925 + 4600 \times S$
$0.93 < S$	5200

Figure 5-39 Single STS-1 Pointer Adjustment Phase Variation Limits



Observation Time, S (seconds)	MTIE (nanoseconds)
$S < 0.1$	N/A
$0.1 \leq S \leq 0.18$	$945 \times S$
$0.18 < S \leq 100$	170
$100 < S$	N/A


5.7.2.2 Pointer Adjustment Bursts

As was the case for jitter generation, phase variation criteria for bursts of pointer adjustments are currently specified only for DS3 interfaces. Closely spaced pointer adjustments for lower bit-rate signals are not expected to occur due to the larger amount of phase movement necessary to generate VT pointer adjustments. The pointer adjustment burst sequences specified in the following criteria are intended to simulate the expected worst-case pointer activity due to phase transients caused by synchronization rearrangements (see Section 5.4.4.3.5).

R5-283 [1029] The MTIE of a DS3 output from a SONET NE shall be below the mask in Figure 5-40 when the pointer test sequence described in Figure 5-31 is applied.


R5-284 [1030] The MTIE of a DS3 output from a SONET NE shall be below the mask in Figure 5-41 when the pointer test sequence described in Figure 5-32 is applied.

Figure 5-40 Maximum Rate Pointer Burst Phase Variation Limits



Observation Time, S (seconds)	MTIE (nanoseconds)
$S < 0.1$	N/A
$0.1 \leq S \leq 0.28$	$1820 \times S$
$0.28 < S \leq 100$	510
$100 < S$	N/A

Figure 5-41 Phase Transient Pointer Burst Phase Variation Limits



Observation Time, S (seconds)	MTIE (nanoseconds)
$S < 0.1$	N/A
$0.1 \leq S \leq 0.70$	$1650 \times S$
$0.70 < S \leq 100$	1155
$100 < S$	N/A


5.7.2.3 Periodic Pointer Adjustments

Periodic pointer adjustments can be caused by frequency offsets between SONET islands, or by offsets resulting from synchronization reference failures. The magnitude of the offset determines the frequency of the pointer adjustments.

-
- R5-285 [360v3]** The MTIE of a DS1 output from a SONET NE shall be below the mask of [Figure 5-42](#) when the pointer adjustment test sequences in [Figures 5-33b](#) and [5-35b](#) are applied with T in the required range (see [Table 5-13](#)).
 - O5-286 [937v2]** The MTIE of a DS1 output from a SONET NE should be below the mask of [Figure 5-42](#) when the pointer adjustment test sequences in [Figures 5-33b](#) and [5-35b](#) are applied with T in the objective range.
 - R5-287 [1031v2]** The MTIE of a DS3 output from a SONET NE shall be below the mask of [Figure 5-43](#) when the pointer adjustment test sequences in [Figures 5-34b](#) and [5-35b](#) are applied with T in the required range.


O5-288 [1032v2] The MTIE of a DS3 output from a SONET NE should be below the mask of Figure 5-43 when the pointer adjustment test sequences in Figures 5-34b and 5-35b are applied with T in the objective range.

Figure 5-42 Periodic VT Pointer Adjustment Phase Variation Limits



Observation Time, S (seconds)	MTIE (nanoseconds)
$S < 0.001326$	N/A
$0.001326 \leq S \leq 0.0164$	$61000 \times S$
$0.0164 < S \leq 1.97$	$925 + 4600 \times S$
$1.97 < S$	10,000

Figure 5-43 Periodic STS-1 Pointer Adjustment Phase Variation Limits



Observation Time, S (seconds)	MTIE (nanoseconds)
$S < 0.1$	N/A
$0.1 \leq S \leq 0.44$	$1830 \times S$
$0.44 < S \leq 100$	800
$100 < S$	N/A

6 SONET Network Element Operations Criteria

This section contains Operations, Administration, Maintenance and Provisioning (OAM&P) criteria that are common to all types of SONET NEs. [Section 6.1](#) describes memory administration functions, while [Section 6.2](#) describes maintenance functions including alarm surveillance ([Section 6.2.1](#)), PM ([Section 6.2.2](#)) and testing ([Section 6.2.3](#)). GRs and TRs that cover specific types of SONET NEs refer to details within this section, and may expand or qualify the descriptions as appropriate.

6.1 Memory Administration

Memory administration deals with the functions necessary to control and administer the databases resident in NEs. The memory administration functions discussed in this section include data manipulation, memory backup and restoration, and system administration (including security).

6.1.1 Memory Administration Data

SONET NEs contain provisionable data structures that are common to all SONET NEs, as well as data structures unique to each type of SONET NE. For memory administration and other system and network management functions (e.g., fault management), it is necessary to know the termination, cross-connection and/or multiplex configuration within an NE. Moreover, data specific to a particular client signal, such as the mapping of a digital signal payload (e.g., asynchronous), are also necessary.

In the Open Systems Interconnection (OSI) communications environment described in [Section 8](#), an OS and NE can exchange messages by using standard services and protocols. These services depend on an information model that is described in terms of managed object classes and their attributes. An information model for SONET memory administration is an abstraction of the physical components in a SONET NE that perform the memory administration functions. GR-836-CORE, *OTGR Section 15.2: Generic Operations Interfaces Using OSI Tools – Information Model Overview: Transport Configuration and Surveillance for Network Elements* and GR-836-IMD, *OTGR Section 15.2: Generic Operations Interfaces Using OSI Tools – Information Model Details: Transport Configuration and Surveillance for Network Elements*, define managed object classes, attributes and associated Common Management Information Service Element (CMISE) service mappings that may be needed to perform memory administration functions. GR-1042-CORE, *Generic Requirements for Operations Interfaces Using OSI Tools – Information Model Overview: Synchronous Optical Network (SONET) Transport Information Model* and GR-1042-IMD, *Generic Requirements for Operations Interfaces Using OSI Tools – Information Model Details: Synchronous Optical Network (SONET) Transport Information Model*, describe managed object classes and attributes specific to all types of SONET NEs.

For Transaction Language 1 (TL1), the NE resident database can be viewed by the memory administration OS as a collection of logical matrices referred to as administrative views. Each row of the matrix represents an entity (i.e., a logical

service or resource associated with an NE) of the view, and the columns represent the set of attributes that each entity may have. Data dictionaries for TL1 entities are based on administrative views. GR-472-CORE, *OTGR Section 2.1: Network Element Configuration Management* gives a thorough discussion of administrative views for TL1. GR-199-CORE, *OTGR Section 12.2: Operations Application Messages – Memory Administration Messages*, provides administrative views of the data dictionaries for SONET entities.

6.1.2 Data Manipulation

Data manipulation deals with entering, editing, deleting and retrieving data in the database resident in an NE. These functions are required to provision new services and equipment in the network.

-
- R6-1 [366]** Any provisionable feature or parameter shall be provisionable locally by a craftsperson and remotely from an OS.
- R6-2 [361v2]** A SONET NE shall use the two-level “Group #, VT #” convention shown in [Section 3](#) of this document (Figures [3-13](#), [3-15](#), [3-17](#) and [3-19](#)) for numbering VTs within an STS-1. This numbering convention shall be used on OS/NE, Workstation (WS)/NE [including any Graphical User Interface (GUI) display] and NE/NE interfaces.
- R6-3 [938]** A SONET NE shall use either the two-level “STS-3 #, STS-1 #” convention or the single-level “1 to N in order of appearance at the input to the byte-interleaver” convention for numbering STS-1s within an OC-N. These numbering conventions shall be used on OS/NE, WS/NE (including any GUI display) and NE/NE interfaces.
-

The STS-1 numbering conventions are illustrated in Figures [5-1](#) and [5-2](#), while [Table 6-1](#) shows and discusses the conversions between the two conventions for OC-N signals. Note that an NE must not use a single-level STS-1 numbering convention where the STS-1s are numbered from 1 to N in the order that they appear in the byte-interleaved OC-N signal [i.e., corresponding to the values in the former C1 (now the J0 and Z0) bytes]. Also note that ITU-T Rec. G.707 contains specifications related to multi-level, Administrative Unit Group-N (AUG-N)-based numbering schemes that can be used to identify each Administrative Unit-n (AU-n) within an STM-N signal.¹ Whether or not equivalent schemes should be included in a future issue of this document is currently under study.

1. For example, the scheme used to identify an AU-3 within an STM-256 is a five-level scheme consisting of an AUG-64 number (1 through 4), an AUG-16 number (1 through 4), an AUG-4 number (1 through 4), an AUG-1 number (1 through 4), and an AU-3 number (1 through 3).

Table 6-1 STS-1 Numbers in OC-N Signals

STS-1 Numbers in an OC-3		STS-1 Numbers in an OC-12		STS-1 Numbers in an OC-48 ^a			
STS-3 #/ STS-1 # Scheme	1 to N Scheme	STS-3 #/ STS-1 # Scheme	1 to N Scheme	STS-3 #/ STS-1 # Scheme	1 to N Scheme	STS-3 #/ STS-1 # (cont.)	1 to N (cont.)
1,1	1	1,1	1	1,1	1	9,1	25
1,2	2	1,2	2	1,2	2	9,2	26
1,3	3	1,3	3	1,3	3	9,3	27
Order of transmission in a byte-interleaved OC-3 is (1,1), (1,2), (1,3) or 1, 2, 3.				2,1	4	2,1	4
				2,2	5	2,2	5
				2,3	6	2,3	6
				3,1	7	3,1	7
				3,2	8	3,2	8
				3,3	9	3,3	9
				4,1	10	4,1	10
				4,2	11	4,2	11
				4,3	12	4,3	12
Order of transmission in a byte-interleaved OC-12 is (1,1), (2,1), (3,1), (4,1), (1,2), ..., (4,3) or 1, 4, 7, 10, 2, ..., 12.				5,1	13	5,1	13
				5,2	14	5,2	14
				5,3	15	5,3	15
				6,1	16	6,1	16
				6,2	17	6,2	17
				6,3	18	6,3	18
				7,1	19	7,1	19
				7,2	20	7,2	20
				7,3	21	7,3	21
				8,1	22	8,1	22
				8,2	23	8,2	23
				8,3	24	8,3	24
						13,1	37
						13,2	38
						13,3	39
						14,1	40
						14,2	41
						14,3	42
						15,1	43
						15,2	44
						15,3	45
						16,1	46
						16,2	47
						16,3	48
				Order of transmission in a byte-interleaved OC-48 is (1,1), (2,1), ..., (16,1), (1,2), ..., (16,3) or 1, 4, ..., 46, 2, ..., 48.			

Note:

- a. For OC-192, the pattern shown for OC-48 is continued with (17,1) and 49, (17,2) and 50, ..., up to (64,3) and 192. Similarly, for OC-768 the pattern for OC-192 is continued with (65,1) and 193, (65,2) and 194, ..., up to (256,3) and 768. The order of transmission in a byte-interleaved OC-192 is then (1,1), (2,1), ..., (64,1), (1,2), ..., (64,3) or 1, 4, ..., 190, 2, ..., 192, while in a byte-interleaved OC-768 it is (1,1), (2,1), ..., (256,1), (1,2), ..., (256,3) or 1, 4, ..., 766, 2, ..., 768.

-
- R6-4 [939]** For numbering an STS-Mc within an OC-N, a SONET NE shall use the number of the STS-1 in which the STS-Mc starts. This numbering convention shall be used on OS/NE, WS/NE (including any GUI display) and NE/NE interfaces.
-

The STS-Mc numbering convention is illustrated in [Figure 5-2](#), in which the STS-3c that starts in STS-1 number 4,1 is referred to as “STS-3c Number 4,1” and the STS-12c that starts in STS-1 number 5,1 is referred to as “STS-12c Number 5,1”.

6.1.3 Administration of Operations Communications Information

To enable a SONET NE to communicate with management systems and other NEs, the NE needs to be initialized at installation time with communications-related information supplied by the network provider. This information includes options for tailoring each layer of DCC, Local Communications Network (LCN) or OS/NE communications protocol stacks, network addresses, and Target Identifiers (TIDs) for NEs employing TL1 messages.

-
- R6-5 [367]** A SONET NE shall provide, via the local craftsperson interface, the ability to initialize the NE with communications-related information upon installation of the NE by the network provider. Such information shall include protocol options for the various operations communications interfaces supported by the NE, the NE’s TID (when TL1 messages are supported), and the NE’s network address (i.e., its Network Service Access Point, or NSAP) for communications purposes. The supplier shall clearly specify in user documentation those data items that need to be provided upon installation of the NE to ensure proper operations communications involving the NE.
-

6.1.4 Regenerators

A line with regenerators may contain either physical-layer regenerators or STE regenerators. Regenerators on a line where the primary or secondary Section Embedded Operations Channel (EOC) is, or will be, active must be STE regenerators.

-
- R6-6 [368]** The LTE shall be provisionable to indicate either physical-layer regenerators or STE regenerators are being used for a given line.
-

Such an indication allows the LTE to identify lines that are suitable for supporting Section EOCs. This is especially useful for LTE with features that allow software reconfiguration of protection systems.

6.1.5 Memory Backup and Restoration

-
- R6-7 [369v2]** A SONET NE shall provide a local nonvolatile memory backup.
-

Such local backup is typically provided within the equipment frame. GR-472-CORE contains details on memory backup. A SONET NE may optionally provide an additional (or secondary) local, nonvolatile memory backup.

R6-8 [370] Data shall be backed up in at least one nonvolatile backup memory automatically after each primary memory update.

R6-9 [371] Restoration of data from the local backup memory, once initiated, shall be completed within 5 minutes.

O6-10 [372] Restoration of data from the local backup memory, once initiated, should take no more than one minute.

These time requirements represent upper bounds. GRs and TRs covering specific type of NEs may provide more stringent restoration time requirements for those NEs. GR-472-CORE contains additional criteria related to memory backup and restoration.

There may be catastrophic failure conditions under which even the local nonvolatile memory backup is lost. To guard against such instances, network providers will designate a management application (e.g., a memory administration or configuration management OS) to be responsible for remotely restoring the memory of NEs so affected. Such management applications must constantly maintain an accurate view of the database changes made in the NE.

R6-11 [373] A SONET NE shall be able to have its configurable memory restored by a remote memory restoration application identified by the network provider.

R6-12 [374v2] A SONET NE shall be able to determine whether or not the source of an update to its configurable memory is the same management application (e.g., OS) as the one responsible for restoring lost local nonvolatile memory backups.

R6-13 [375] When the source of a memory update is different from the memory restoration application (e.g., a local craftsman or other management application), the SONET NE shall send an autonomous indication to the memory restoration application detailing this “hidden update.”

R6-14 [376] If communications to the memory restoration application are not available, “hidden updates” shall be logged by the SONET NE and reported when asked for by the memory restoration application after communications are restored.

Historically, the remote restoration of an NE’s memory has involved the transmission of individual memory administration messages to re-create the last view of the NE’s total configuration. While this process can be manageable for small NEs, the process can be quite burdensome and time-consuming for large NEs. Memory restoration via bulk files is becoming a critical operations feature for memory restoration applications and larger NEs.

CR6-15 [377] A SONET NE may be required to support the ability to have its nonvolatile memory backup restored via bulk file transfer methods by a remote memory restoration management application (e.g., an OS or controller).

As indicated in the GRs and TRs covering specific types of SONET NEs, this feature may be considered an objective or a requirement for certain types of NEs. If the SONET NE supports the optional bulk memory restoration feature, the following requirements apply.

R6-16 [378v2] If a SONET NE supports the optional bulk memory restoration feature, then the NE shall support the feature via a full 7-layer OSI-based operations interface to a bulk memory restoration application using the memory backup functions and File Transfer Access and Management (FTAM) protocol or File Transfer Protocol (FTP) requirements described in GR-1250-CORE, *Generic Requirements for Synchronous Optical Network (SONET) File Transfer*.

R6-17 [379] If a SONET NE supports the optional bulk memory restoration feature, then the NE shall be able to report a bulk “snapshot” of its nonvolatile memory backup to the memory restoration application upon request.

GR-836-CORE and GR-836-IMD contain the information model and CMISE service mappings, and GR-199-CORE and GR-833-CORE, *OTGR Section 12.3: Network Maintenance: Network Element and Transport Surveillance Messages*, contain TL1 messages to support transaction-oriented memory backup and restoration. Additional requirements for the management of bulk memory restoration processes may be found in GR-2915-CORE, *Application of Software Management OSI Information Model to Software Download and Memory Restoration*.

6.1.6 System Administration and Security

System administration deals with housekeeping functions needed for proper operation of the NE in service-provider networks. Examples of system administration functions include setting the date and time, and NE identification. NE requirements to support system administration functions are provided in GR-472-CORE.

A major area that impacts system administration is security. Security requirements involve routing functions (within the operations communication network), logins, passwords and security levels (screening options). GR-815-CORE, *Generic Requirements for Network Element/Network System (NE/NS) Security* specifies generic requirements for NE security functions, and TR-NWT-000835, *OTGR Section 12.5: Network Element and Network System Security Administration Messages* specifies TL1 messages that can be used to administer NE security. Preliminary requirements for the information model and service mappings for NE security administration are provided in GR-1253-CORE, *Generic Requirements for Operations Interfaces Using OSI Tools: Telecommunications Management Network Security Administration*.

- R6-18 [380]** A security mechanism shall be provided within a SONET NE to prevent unauthorized communication to the NE via any ports and communications channels accepting operations-related command inputs,² and to allow secure access to the database of the NE. Such a mechanism shall adhere to the security requirements of GR-815-CORE.
- R6-19 [381]** The data necessary to support the security mechanism within a SONET NE shall be provided and administered only by authorized security administrators via either WS/NE or OS/NE communications.
- R6-20 [382]** A SONET NE shall support security administration functions in conformance with GR-815-CORE.

It is a goal to centralize security administration using Telecommunications Management Network (TMN) Public Key Infrastructure (PKI) Recommendations, although the requirements for TMN PKI functionality are not currently fully specified. On an interim basis, a centralized security server may be used in conjunction with TL1-manageable NEs using the requirements found in NSIF-038-2000, *NSIF Requirements for a Centralized Security Server*.

In addition, applications in every NE need to employ specific security measures such as logins and passwords to further protect against unauthorized access. Access control features such as authorization levels are needed to limit authorized access to the appropriate functions. The following sections describe additional criteria regarding security mechanisms to be provided in SONET NEs.

6.1.6.1 NE Security Mechanism

GR-815-CORE documents the security requirements for NEs and Mediation Devices (MDs). As discussed in GR-815-CORE, a security mechanism implemented within a SONET NE provides answers to the following questions.

- Is the session requester a valid user?
- Does the calling address of the request origination conform with that of the said valid user?
- Is the user authorized to issue the command being issued?
- Is the user authorized to access the particular portion of the NE database that the command is directed to impact?

Unless all the answers are in the affirmative, a transaction cannot be successfully completed. The security mechanism also provides an audit capability so that unauthorized activities and attempts can be investigated. This mechanism contains the following components:

- Identification and authentication

2. Operations-related commands include functions that allow direct access to any NE resources, such as database, memory, software, processes, etc. In this context, access refers to the following operations: create, retrieve, update and delete.

- System access control
- Resource access control
- Audit.

Identification is the process whereby a session requester's unique and auditable identity (such as the user ID) is recognized. Authentication is the process of verifying the claimed identity of the session requester (such as password check). System access control uses features related to a user's session, such as session "timeout" and real-time detection of password cracking attempts by intruders, thus decreasing the risk of an intruder gaining access by posing as a valid user. Resource access control is enforced after an authorized session has been established, ensuring that no access to the NE database is allowed without proper permission. It serves a dual purpose of accomplishing confidentiality and integrity of the data. Audit features provide the data necessary to support the detection and investigation of unauthorized access.

Telcordia has also issued security requirements for OSI-based TMN interfaces. That document (GR-1469-CORE, *Security for OSI-Based Telecommunications Management Network Interfaces*) specifies security requirements for OS/NE, OS/OS and NE/NE interfaces.

R6-21 [383] A SONET NE supporting interfaces conforming to 7-layer OSI protocol stacks shall conform to the security requirements described in GR-1469-CORE for layers 1 through 6. In addition:

- When TL1-based interfaces are used in the NE for WS/NE, NE/NE, or OS/NE communications, the NE shall support the data and messages provided in TR-NWT-000835 for the administration of the NE security mechanism
- When CMISE interfaces are used on the application layer, the NE shall support data, messages and mechanisms to conform to the requirements provided by GR-1469-CORE.

Data communications is an integral part of the operations and control infrastructure for the SONET architecture. The security of Local Area Networks (LANs) and WANs that interconnect the SONET NEs, OSs and MDs has a bearing on the security of the SONET transport service. Telcordia has issued security requirements (GR-1332-CORE, *Generic Requirements for Data Communications Network Security*) for the data communications network component of the TMN.

R6-22 [384] The data communications network component of the SONET TMN shall conform to the security requirements in GR-1332-CORE.

6.1.6.1.1 Identification and Authentication

R6-23 [385] A SONET NE shall support identification and authentication for all users, for all ports accepting operations-related command inputs, in conformance with GR-815-CORE.

R6-24 [386] A SONET NE that supports remote access for operations-related command inputs shall provide a feature for additional strong authentication, beyond reusable passwords, such as:

- Third-party authentication
- Public/private key encryption technology.

O6-25 [1122] It is an objective that SONET NEs provide a feature for additional strong authentication for all users, beyond reusable passwords, for all ports accepting operations-related command inputs.

The intent of the above objective is to provide the same level of security for all NE access ports (e.g., local craftsperson access) as is required by **R6-24 [386]** for remote access ports.

It is a goal that **R6-24 [386]** and **O6-25 [1122]** be met using TMN PKI technology. On an interim basis, **R6-24 [386]** may be met using a centralized security server as specified by NSIF-038-2000.

R6-26 [387] When TL1 interfaces are used, a SONET NE shall enforce that a session requester accessing the NE via a TL1/X.25-based OS/NE interface must pass identification information based on an X.25 calling address or Permanent Virtual Circuit (PVC) identifier.

6.1.6.1.2 *System Access Control*

R6-27 [388] A SONET NE shall support system access control functions in conformance with GR-815-CORE.

R6-28 [389] A SONET NE shall not grant a user remote access unless the user is authenticated via strong authentication.

R6-29 [390] A SONET NE shall employ features corresponding to the timeout interval function that GR-815-CORE describes.

R6-30 [391] A SONET NE shall break down an OSI application association if an attempted session request is unsuccessful after a provisionable number of tries. The default number of tries shall be three.

6.1.6.1.3 *Resource Access Control*

R6-31 [392] A SONET NE shall support resource access control and authorization functions in conformance with GR-815-CORE.

R6-32 [393v2] A SONET NE supporting one or more restricted DCCs shall support user identification and access control privileges (see GR-815-CORE) that limit the functionality available to valid outside users accessing the NE via a restricted DCC. The functions allowed via such DCCs shall be definable by local service providers.

6.1.6.1.4 Audit

R6-33 [394] A SONET NE shall support audit features in conformance with GR-815-CORE.

6.1.6.2 DCC Security

The location of an NE in the network can necessitate the enforcement of certain security-related restrictions on one or more of the DCCs the NE supports. For example, an NE supporting line-side interfaces that cross an administrative boundary needs to apply restrictive message routing functions on the DCCs that cross that boundary. Restrictions are needed to:

- Minimize the possibility of unauthorized access to one network's operations functions in NEs and OSs by people, systems or equipment in another network
- Limit the broadcast of one network's OS/NE, NE/NE and WS/NE communications into the other network, thereby reducing the need to control network congestion and also reducing the opportunity to monitor another network's operations communications.

If a SONET network within an administrative domain supports broadcast routing, the network provider may want to disable the DCC crossing the boundary to eliminate excessive broadcasting of messages from outside the routing domain and potential flooding of the network. However, if a network provider wants an active DCC crossing the boundary to allow, for example, its OS to monitor far-end performance data from the NE in the other domain, the DCC crossing the boundary is labeled restricted and the NE terminating that DCC has to support selective routing. This also eliminates the problem of broadcasting all messages from one domain into another domain. To use the restricted DCC, the boundary NE has to restrict access to the network based on source address and destination address. The network provider should take the particular application and needs into consideration in choosing one of these two methods (i.e., disabled DCC or restricted DCC).

R6-34 [395v2] An NE shall be capable of disabling a SONET DCC. The default shall be to enable an equipped DCC.

R6-35 [397v2] Only properly authorized system administrators shall be allowed to enable or disable an equipped SONET DCC.

R6-36 [398] On reinitialization of the NE after failure, DCCs shall maintain the enabled or disabled state they held before the failure of the NE.

- CR6-37 [399]** To provide NE/NE and indirect OS/NE communications paths that may be required by a network provider across administrative boundaries, an NE may be required to terminate one or more enabled DCCs that cross an administrative boundary.
- R6-38 [400]** If an NE has the capability to terminate a DCC that crosses an administrative boundary, it shall be capable of classifying it as either restricted or unrestricted for operations security purposes. The default setting shall be unrestricted.
- R6-39 [402]** Only properly authorized system administrators shall be allowed to change the classification of a DCC from restricted to unrestricted or vice versa.
- R6-40 [403]** A change in classification of a restricted DCC shall not be allowed by communications over that or any other restricted DCC.
- R6-41 [404]** On reinitialization of the NE after failure, DCCs shall maintain the restricted or unrestricted states they held before the failure of the NE.
-

Single-ended maintenance and other functions may still be needed between NEs that span an administrative boundary. An Inter-Carrier Interface (ICI) is an example of a SONET interface that spans an administrative boundary. OSs on one side of the administrative boundary may need, in a tightly controlled and agreed-upon manner, to communicate with an NE on the other side of the boundary via the restricted DCC. Selective routing is one function that must be employed at the network layer in an NE supporting such restricted DCCs to satisfy security concerns and still allow operation of the links that cross administrative boundaries.

The following requirements apply to a SONET NE having one or more active DCCs that have been classified as restricted.

- R6-42 [405]** A SONET NE shall not forward received Network Protocol Data Units (NPDUs) across an administrative boundary to the underlying Data Link Service for a restricted DCC. The NE shall “terminate” the DCC.
- R6-43 [406]** A SONET NE shall support, for each restricted DCC, a list that itemizes NSAP address pairs that are to be allowed in source and destination address fields of NPDUs allowed into the NE’s network.
- R6-44 [407]** The list described in **R6-43 [406]** shall be established via either WS/NE or OS/NE communications, and only by authorized system administrators via an unrestricted DCC or direct interfaces.
- R6-45 [408]** A SONET NE shall screen the NPDUs received from a restricted DCC and only accept the PDU if the source address and destination address matches an allowable source/destination address pair from the list described in **R6-43 [406]**.
- R6-46 [409]** A SONET NE shall enforce access control on the restricted DCC to restrict the requested operations functions that will be permitted on a per-user basis.
-

- R6-47 [410]** A SONET NE shall provide the ability for an authorized administrator to specify the access control privileges assigned to a user for restricted DCC use.
-

6.1.7 Software Generics

- R6-48 [411]** The initial software generic shall be entered in the SONET NE at or before installation.
- O6-49 [412v2]** A SONET NE should be able to receive its initial software load and later software generics via either an OS/NE or NE/NE (or MD/NE) interface using FTAM or FTP.
-

Requirements for software download using the SONET Operations Communications network are provided in GR-1250-CORE. Generic requirements for the support of Remote Software Management are provided in GR-472-CORE.

- R6-50 [413]** The NE shall provide the ability to retrieve (locally via a WS and remotely via an OS) the current version ID of software.
- O6-51 [414]** Software updates, or patches, should be identified and included in the current version report.
-

6.1.8 Self Inventory

Self inventory is a feature of an NE by which it inventories its own equipment. The NE self-inventory feature is used by Element Management Systems (EMSs) in order to perform “automatic discovery” of the subnetwork of NEs being managed by the EMS. In order for an EMS to have a complete view of the subnetwork, it is desirable for NEs to report information not only about their own equipage, but also information about which NEs they are physically and logically connected to (i.e., adjacency information).

- R6-52 [415v2]** A SONET NE shall be able to report to a management application or a craftsperson its equipage (including plug-ins, common equipment and software), option settings and cross-connect configuration.
- O6-53 [1123]** It is an objective that a SONET NE be able to report to a management application or craftsperson information about which NEs it is physically and logically connected to (i.e., adjacency information).
-

6.2 Maintenance

This section provides SONET NE maintenance criteria necessary to maintain the NE and the network. Maintenance requirements include alarm surveillance, PM, testing and control features that are essential in the normal operation of the NE. The requirements in this section address functions that are used to perform the following maintenance tasks.

- **Trouble detection** deals with detecting defects and declaring failures. [Section 6.2.1.1](#) defines defects and failures related to SONET.
- **Trouble or repair verification** is the process of verifying the continued existence or nonexistence of a problem before beginning or closing out work on that problem.
- **Trouble sectionalization** deals with sectionalizing the failure to one of the terminating NEs or the facility that connects them. This process is achieved through alarms, maintenance signals (e.g., AIS), PM data, test access and loopbacks.
- **Trouble isolation** deals with the isolation of failures down to a replaceable circuit pack, module or fiber. Test access, loopbacks, performance data and diagnostics available within the SONET NEs are used to achieve this isolation.
- **Restoration** permits service to be restored even though the failure may not have been repaired. Protection switching, described in [Section 5.3](#), and rerouting of traffic are examples of how restoration can be achieved.

In formulating generic maintenance criteria common to all SONET NEs, it is useful to take a functional view of equipment. Thus, maintenance criteria common to all SONET NEs may be specified in terms of STE, LTE, STS PTE and VT PTE.

Suppliers are referred to SR-NWT-002723, *Applicable TL1 Messages for SONET Network Elements*, for information on how specific TL1 messages apply to the maintenance functions discussed here.

6.2.1 Alarm Surveillance

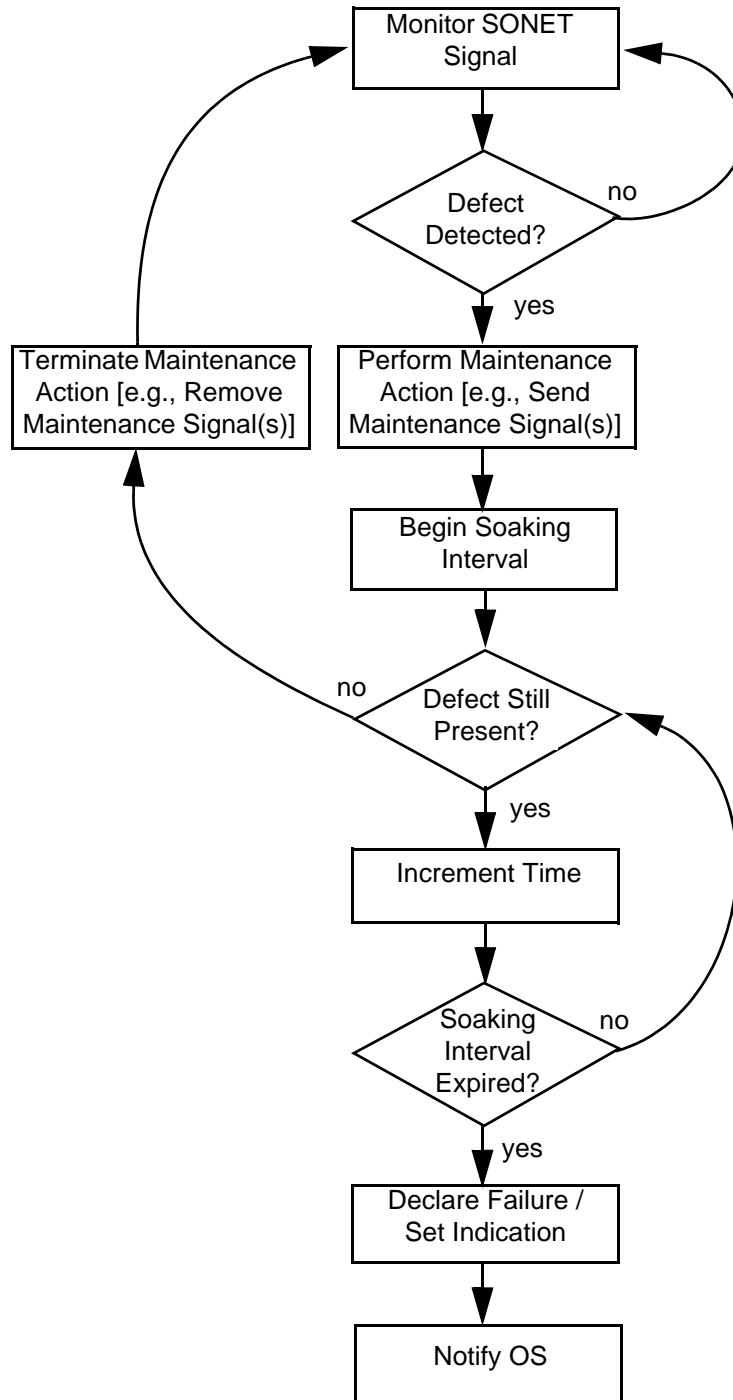
Alarm surveillance deals with the detection and reporting of certain degraded conditions in the network. This section enumerates those conditions that need to be detected within the SONET signal and the NE. In addition to defining the various conditions, this section also describes the NEs' actions in response to detecting those conditions.

This document refers to occurrences in the network that are detected as "defects". A defect is defined to be a limited interruption in the ability of an item to perform a required function, and SONET NEs are required to "detect" and "terminate" certain defects on the incoming signals relevant to the layers of functionality they provide. Detection of a defect may cause a particular action to be performed (e.g., the transmission of a maintenance signal), while termination of the defect generally causes that action to be halted (e.g., removal of the maintenance signal). When a defect persists for a period of time (i.e., a soaking interval), a corresponding failure is generally "declared" and the NE sets a failure indication. Once a failure indication has been set, if the defect is terminated and remains absent for a period of time, then

the failure is “cleared”. Failure indications may or may not be automatically reported to the OS, and the reported indications may be alarmed or nonalarmed. Failure indications, whether or not they are automatically reported to the OS, are available and retrievable by the OS or some other user system interface. Some failure indications may also result in audible and/or visible alarm indications locally at the NE. For more details on the failure monitoring process and alarm strategy, including definitions of critical, major and minor alarms, see GR-474-CORE and GR-820-CORE.

The defects and failures discussed in [Section 6.2.1.1](#) are considered to be directly detected defects and failures, and represent root-cause problems on the incoming SONET bitstream. [Sections 6.2.1.2](#), [6.2.1.3](#) and [6.2.1.4](#) describe symptomatic defects, which are detected via maintenance signals on the incoming SONET bitstream that are generated as a result of an upstream or downstream SONET NE detecting one of the directly detected defects described in [Section 6.2.1.1](#). A general model of defect detection and failure declaration is illustrated in [Figure 6-1](#).

Figure 6-1 General Defect Detection and Failure Declaration Model



6.2.1.1 Directly Detected Defects and Failures

6.2.1.1.1 Loss of Signal

To detect a failure that occurs at the source (e.g., laser failure) or the transmission facility (e.g., fiber cut), all incoming SONET signals are monitored for loss of physical-layer signal (optical or electrical). The detection of an LOS defect must take place within a reasonably short period of time for timely restoration of the transported payloads.

R6-54 [416v2] A SONET NE shall monitor all incoming SONET signals (before descrambling) for an “all-zeros patterns,” where an all-zeros pattern corresponds to no light pulses for OC-N optical interfaces and no voltage transitions for STS-1 and STS-3 electrical interfaces. An LOS defect shall be detected when an all-zeros pattern on the incoming SONET signal lasts 100 μ s or longer. If an all-zeros pattern lasts 2.3 μ s or less, an LOS defect shall not be detected.

The treatment of all-zeros patterns lasting between 2.3 μ s and 100 μ s for the purpose of LOS defect detection is not specified and is therefore left to the choice of the equipment designer. For testing conformance to the LOS detection requirement, it is sufficient to apply an all-zeros pattern lasting at most 2.3 μ s, and to apply an all-zeros pattern lasting at least 100 μ s.

Note that although an all-zeros pattern that lasts for less than 2.3 μ s must not cause the detection of an LOS defect, an NE that receives a relatively long (in terms of the number of bit periods) all-zeros pattern of less than 2.3 μ s is not necessarily expected to continue to operate error-free through that pattern.³ For example, in such cases it is possible that the NE’s clock recovery circuitry may drift off frequency due to the lack of incoming pulses, and therefore the NE may be “looking in the wrong bit positions” for the SONET framing pattern after the all-zeros pattern ends. If this occurs, it will continue for approximately 500 μ s, at which point the NE will detect an SEF defect. The NE would then perform the actions associated with SEF defect detection (e.g., initiate a search for the “new” framing pattern position), rather than the actions associated with LOS defect detection (e.g., AIS and RDI insertion, possible protection switch initiation).

In addition to monitoring for all-zeros patterns a SONET NE may also detect an LOS defect if the received signal level (e.g., the incoming optical power) drops below an implementation-determined threshold.

3. It should also be noted that there are no explicit criteria concerning the number of consecutive zeros (or ones) that an NE must be able to receive without errors. This topic, as well as the possibility of reducing the 2.3 μ s minimum LOS defect detection limit for high bit-rate SONET signals, has been considered on several occasions in the standards bodies and in the GR-253-ILR. However, no consensus has been reached that would indicate that changes in these areas are needed (or desirable). If such agreements are reached in the future, it is expected that they will be reflected in this document. In any case, in order to meet other criteria such as those related to long-term error performance, an NE would have to be able to tolerate all-zeros and all-ones patterns that have any significant probability of occurring in a “normal” SONET signal.

O6-55 [940] If a SONET NE monitors the received signal level for the purpose of detecting LOS defects, then its signal level threshold should be selected such that an LOS defect will not be detected if the BER is still acceptable (e.g., no LOS defect if the BER is better than the SF BER threshold used for protection switching in linear APS, see [Section 5.3.3.1](#)).

R6-56 [1124] A SONET NE that monitors the received signal level for the purpose of detecting LOS defects shall not terminate an existing LOS defect unless the level of the incoming signal is greater than some LOS defect termination threshold that is in turn greater than the signal level at which the NE detected the LOS defect.

The intent of **R6-56 [1124]** is to prevent an NE from repetitively detecting and terminating an LOS defect as a result of minor fluctuations in the level of the incoming signal around the NE's LOS defect detection threshold. While no specific value is required for the difference between an NE's LOS defect detection and termination thresholds, that difference needs to be large enough to prevent LOS defect oscillations, without being so large that it hinders the reestablishment of traffic after maintenance procedures.

R6-57 [417v2] The SONET NE shall terminate an LOS defect when one of the following occurs:

- The incoming signal level is above the NE's LOS defect termination threshold (if applicable), the signal has two consecutive valid framing alignment patterns and, during the intervening time (one frame), no all-zeros pattern qualifying as an LOS defect exists
 - The incoming signal level is above the NE's LOS defect termination threshold (if applicable), and no pulse-free intervals of length T' occur during a time period equal to the greater of 125 μ s and $2.5 \times T'$, where $2.3 \leq T' \leq 100 \mu$ s.
-

Note that the major difference between the two LOS defect termination methods described in **R6-57 [417v2]** is that in the first method (which is based on the requirement that appeared in early issues of this document) the incoming signal must contain valid framing, while in the second method (which is based on the specifications that first appeared in the 1997 version of T1.231, *Layer 1 In-Service Digital Transmission Performance Monitoring*,) valid framing is not required. In general, the second method is now considered to be the preferred method, since an NE that uses it can in some cases provide more accurate indications of incoming signal problems. For example, suppose that an NE that has previously declared an LOS failure begins to receive a signal with too much optical power, causing its receiver to become overloaded. If the NE supports the first method then it will probably continue to indicate the LOS failure, possibly hindering maintenance efforts. On the other hand, an NE that supports the second method would generally be expected to terminate the LOS defect, clear the LOS failure and declare an LOF failure (see [Section 6.2.1.1.2](#)).

R6-58 [418v3] A SONET NE shall declare an LOS failure when the LOS defect persists for 2.5 (± 0.5) seconds or an LOS defect is present when the conditions for declaring an

LOF failure (see **R6-66 [426v3]**) are met. Upon declaring the failure, it shall set an LOS failure indication and send an alarm message to an OS.

- R6-59 [419]** For the purposes of trunk conditioning, a SONET NE that contains DS0 PTE or VT PTE that supports the rearrangement of the DS0 channels in byte-synchronously mapped DS1s shall declare an LOS failure if it is subject to a period of short, intermittent LOS defects. For failures resulting from the NE intermittently detecting and terminating the defect, the NE shall integrate the time during which the defect persists, using a 4:1 to 15:1 count-up/count-down ratio. During a sustained LOS defect, the integrator shall count up to reach the alarm threshold in 2.5 (± 0.5) seconds. Upon reaching the alarm threshold, the NE shall declare the LOS failure, set an LOS failure indication and send an alarm message to an OS. If the defect is terminated before the threshold is reached, the integrator shall count down at a slope 1/4 to 1/15 of the count-up slope.

SONET NEs that contain DS0 PTE or VT PTE that supports the rearrangement of the DS0 channels in byte-synchronously mapped DS1s are also required to use this integration technique for declaring certain other failures (e.g., see **R6-67 [427]** for LOF failures). In addition, SONET NEs may also apply this (or some similar) integration technique to declare failures due to intermittent defects in applications where trunk conditioning is not required.

-
- R6-60 [420]** A SONET NE shall clear an LOS failure when the LOS defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the SONET NE shall clear the LOS failure indication and send a clear message to an OS.
- R6-61 [421v2]** SONET NEs interfacing with DS1, DS1C, DS2 or DS3 signals shall detect LOS on those signals according to the requirements in GR-499-CORE.

Criteria for LOS on other transport signal interfaces is FFS.

6.2.1.1.2 Loss of Frame

As discussed in [Section 5.5](#), all SONET NEs are required to monitor their incoming SONET signals for SEF defects. In general, if an SEF defect persists, then an LOF defect must be detected according to the criteria in this section. Among the types of problems that could result in the detection of LOF defects are certain source (transmitter) failures and extremely high BER conditions. Such high BER conditions could in turn be caused by problems such as excessively high received optical power levels (i.e., receiver overload) or very low received optical power levels (if low received power levels are not detected as LOS defects, see [Section 6.2.1.1.1](#)).

-
- R6-62 [422v2]** A SONET NE shall detect an LOF defect when an SEF defect (see [Section 5.5](#)) on the incoming SONET signal persists for 3 ms.

A SONET NE may optionally implement a 3-ms integration timer to deal with intermittent SEF defects when monitoring for LOF. Such a 3-ms integration timer consists of an SEF timer and an in-frame timer that operate as follows.

1. The in-frame timer is activated (accumulates) when an SEF defect is absent. It stops accumulating and is reset to zero when an SEF defect is detected.
2. The SEF timer is activated (accumulates) when an SEF defect is present. It stops accumulating when the SEF defect is terminated. It is reset to zero when the SEF defect is absent continuously for 3 ms (i.e., when the in-frame timer reaches 3 ms).
3. An LOF defect is detected when the accumulated SEF timer reaches the 3-ms threshold. Once detected, the LOF defect is terminated when the in-frame timer reaches 3 ms.

-
- R6-63 [423]** If the optional integration timer described above is provided for LOF monitoring, the supplier shall clearly describe its use in the product documentation.
- R6-64 [424]** The SONET NE shall terminate an LOF defect 1 ms to 3 ms after terminating the SEF defect on the incoming SONET signal, if the SEF defect is not (re)detected before the LOF defect is terminated.
- O6-65 [425]** The SONET NE should terminate an LOF defect 1 ms after terminating the SEF defect on the incoming SONET signal, if the SEF defect is not detected within the 1-ms time period. (This objective is not applicable if the optional 3-ms integration timer described above is used.)
- R6-66 [426v3]** Unless an LOS defect is also present (see **R6-58 [418v3]**), a SONET NE shall declare an LOF failure when an LOF defect persists for 2.5 (± 0.5) seconds. In addition, an NE shall declare an LOF failure if it would have previously declared an LOF failure except that an LOS defect was present (so it declared an LOS failure instead), the conditions for clearing the LOS failure are met, and the LOF defect is (still) present when the LOS failure is cleared. Upon declaring an LOF failure, the NE shall set an LOF failure indication and send an alarm message to an OS unless the condition in **R6-312 [626v3]** applies (see [Section 6.2.1.8](#)).
- R6-67 [427]** For the purposes of trunk conditioning, SONET NEs that contain DS0 PTE or VT PTE that supports the rearrangement of the DS0 channels in byte-synchronously mapped DS1s shall use the integration technique described in **R6-59 [419]** to declare LOF failures. Upon declaring an LOF failure, the NE shall perform the actions listed in **R6-66 [426v3]**.
- R6-68 [428v2]** A SONET NE shall clear an LOF failure when the LOF defect is absent for 10 (± 0.5) seconds. Upon clearing the LOF failure, the SONET NE shall clear the LOF failure indication and send a clear message to an OS (if the failure was reported to the OS).
-

For SONET signals, detection of the LOF defect initiates certain maintenance-related actions (e.g., generation of AIS and RDI), which are discussed in [Sections 6.2.1.2](#), [6.2.1.3](#) and [6.2.1.4](#). For DS_n signals, it is the detection of an OOF defect that initiates those maintenance-related actions. Requirements on DS_n OOF

defect detection and termination and maintenance-related actions for DS1, DS1C, DS2 and DS3 signals are contained in GR-499-CORE. Framing criteria for other transport signal interfaces are FFS.

-
- R6-69 [429v2]** An NE shall monitor for DS_n OOF defects on DS_n paths that are terminated by the NE.
-

While a SONET NE may terminate the DS_n path in some applications, in many other applications it is expected to provide clear-channel transport of DS_n payloads.⁴ In clear-channel DS_n transport, incoming DS_n signals are monitored for DS_n LOS defects, incoming SONET signals are monitored for a variety of defects (e.g., LOS, LOF), and the appropriate AIS is inserted downstream when any of those defects are detected. In some cases an NE providing clear-channel DS_n transport may also non-intrusively monitor the framing of the DS_n signals (e.g., to support the accumulation of intermediate DS_n path PM, see [Section 6.2.2.9](#)); however, in those cases AIS is not inserted when a DS_n OOF defect is detected. If an NE inserts AIS when it detects a DS_n OOF defect, then it is not providing clear-channel transport. As indicated in [Section 3.4](#), the DS_n asynchronous mappings were defined specifically for the clear-channel transport of DS_n signals in STS and VT SPEs. Therefore, in most situations where the asynchronous mappings are used, a SONET NE is not expected to insert AIS upon detecting a DS_n OOF defect.

-
- R6-70 [941]** If an NE supports an asynchronous DS_n mapping as defined in [Section 3.4](#), then it shall be capable of providing clear-channel transport of DS_n signals using that mapping.
-

In addition, an NE that supports an asynchronous DS_n mapping could support one or more options to provide non-clear-channel transport of DS_n signals using that mapping. In non-clear-channel transport, the NE monitors for DS_n OOF defects (as if it were terminating the DS_n path) and inserts the appropriate AIS downstream when an OOF defect is detected. If such a feature is provided, the criteria in GR-499-CORE on detecting/terminating DS_n OOF defects and declaring/clearing DS_n LOF failures are applicable.

-
- CR6-71 [942v2]** An NE that supports an asynchronous DS_n mapping may be required to provide non-clear-channel transport where the incoming DS_n signal is monitored for DS_n OOF defects.

4. Note that in this context, the term “clear-channel transport” refers to the transport of DS_n payloads through the SONET network. The same term is also used in other contexts to refer to the use of various line codes or consecutive-zero suppression codes which allow DS_n signals with arbitrary numbers of consecutive logical zeros to be transported on DS_n facilities (e.g., B8ZS).

CR6-72 [943v2] An NE that supports an asynchronous DS_n mapping may be required to provide non-clear-channel transport where the outgoing DS_n signal (i.e., the DS_n signal that is demultiplexed from the STS or VT SPE) is monitored for DS_n OOF defects.

A SONET NE may also provide an option to monitor any error detection bits provided in the DS_n signal (e.g., the P-bits in a DS3 signal) and correct those bits when errors are detected. That option, and the support of non-clear-channel DS_n transport in general, may be useful in situations where the DS_n signals are subsequently transported on certain older transmission systems.⁵ Note however, that clear-channel DS_n transport is the preferred mode for NEs supporting the asynchronous DS_n mappings, and that **CR6-71 [942v2]** and **CR6-72 [943v2]** are not expected to be changed to requirements in future issues of this document.

6.2.1.1.3 Loss of Pointer

To detect a failure related to the pointer processing mechanism (e.g., for trouble isolation purposes), an NE that processes pointers also monitors incoming STS and VT paths for LOP. SONET equipment detects an LOP defect when a valid STS or VT pointer that it is able to accommodate (see the discussion following **R6-74 [944v2]**) cannot be obtained by using the pointer interpretation rules described in [Section 3.5](#) and the pointer word does not contain all-ones. An LOP defect is also detected when a number of consecutive pointers with the NDF set to '1001' but not indicating concatenation is received. [Under normal operation, the NDF would be set only once to indicate a change in pointer value (or VT size). Except when set continuously in a concatenation indicator, consecutive NDFs would indicate a pointer processor failure (e.g., stuck bits).]

R6-73 [430v4] STS PTE and LTE that processes the STS pointers shall monitor for LOP-P. An LOP-P defect shall be detected if a valid pointer that is able to be accommodated is not found in N consecutive frames (where $8 \leq N \leq 10$), or if N consecutive NDFs (other than in a concatenation indicator, see [Section 3.5.1.3](#)) are detected. An LOP-P defect shall not be detected when LTE is receiving and relaying an all-ones STS pointer, or when STS PTE is receiving pointers that qualify as those necessary to cause the detection of an AIS-P defect (i.e., three or more consecutive all-ones pointers).

R6-74 [944v2] VT PTE and STS PTE that processes VT pointers shall monitor for LOP-V. An LOP-V defect shall be detected if a valid pointer is not found in N consecutive VT superframes (where $8 \leq N \leq 10$), or if N consecutive NDFs are detected. An LOP-V

5. For example, some older asynchronous fiber optic transport systems required incoming DS3 signals to have valid framing and (in some cases) parity. In those cases, the SONET NE would need to insert DS3 AIS when it detects a DS3 OOF, and may also need to correct the P-bits in the DS3 signal.

defect shall not be detected when STS PTE is receiving and relaying an all-ones VT pointer, or when VT PTE is receiving pointers that qualify as those necessary to cause the detection of an AIS-V defect.

Depending on the functionality of the receiving equipment, some pointers that would be considered valid in other situations may not be able to be accommodated, and therefore would need to cause the detection of LOP defects. This would be the case, for example, if LTE that has been provisioned to perform an STS-Nc cross-connection (covering a particular set of N STS-1 timeslots) instead receives some combination of STS-1 and/or STS-Mc ($M < N$) path signals, or LTE that has been provisioned to perform some combination of STS-1 and/or STS-Mc cross-connections instead receives an STS-Nc ($N > M$) path signal. In the former case, **R3-8 [11]** and **R6-73 [430v4]** together indicate that the N sets of pointer bytes associated with the provisioned STS-Nc path need to be processed as a single entity, and the reception of “normal” pointers in one or more of the $N - 1$ sets that are supposed to contain concatenation indications would need to result in the detection of a single LOP-P defect. In the latter case, the desired result would be for separate LOP-P defects to be detected for each of the provisioned STS-1 or STS-Mc paths. While **R6-73 [430v4]** is sufficient to achieve this result for the STS-1 and STS-Mc paths whose pointer processors receive all concatenation indications (rather than the expected normal pointers, or normal pointers and $M - 1$ concatenation indications), a problem occurs for the STS-1 or STS-Mc path that is supposed to be utilizing the timeslots occupied by the first or first M STS-1s of the STS-Nc path. In particular, the contents of the pointers associated with the first or first M STS-1s of an STS-Nc path are identical to the expected contents for an STS-1 or STS-Mc path. Therefore, the following objectives are applicable.

- O6-75 [1189]** STS PTE and LTE that processes STS pointers and is provisioned to (or can only) process an STS-1 path signal that occupies timeslots that could be occupied by the first STS-1 of an STS-Nc path signal should detect an LOP-P defect within 100 ms if the pointers received by the NE containing that equipment indicate the persistent presence of an STS-Nc path signal in those timeslots.
- O6-76 [1190]** STS PTE and LTE that processes STS pointers and is provisioned to (or can only) process an STS-Mc path signal that occupies timeslots that could be occupied by the first M STS-1s of an STS-Nc ($N > M$) path signal should detect an LOP-P defect within 100 ms if the pointers received by the NE containing that equipment indicate the persistent presence of an STS-Nc path signal in those timeslots.

Note that if **O6-75 [1189]** or **O6-76 [1190]** is not met, then the defect that will be detected when the STS path is terminated will depend on several factors. For example, if the STS PTE is provisioned to detect TIM-P defects (see [Section 6.2.1.1.9](#)), then that is the highest priority defect that would be expected to be detected. If TIM-P defect detection is not enabled and the STS-Nc SPE happens to contain a different type of payload than the STS-1 or STS-Mc SPE is supposed to be carrying, then an STS Path Payload Label Mismatch (PLM-P) defect would be detected (see [Section 6.2.1.1.8.A](#)). Finally, if no PLM-P defect is detected, then it is possible that a payload-layer defect (e.g., loss of cell delineation) would be detected.⁶

Also, several definitions of “valid pointer” are possible, and the particular definition that a supplier uses will have a large impact on how that supplier interprets the above requirement and implements an LOP detection algorithm. However, since LOP is not expected to be a common defect in the network, these different definitions and interpretations are not expected to have a significant impact on network performance.

The intended definition of “valid pointer” is:

1. A pointer with an in-range value, the N-bits set to their normal value and, for VT pointers, constant or correct size bits (see [Section 3.5.2.3](#)) that is received identically in three consecutive frames (VT superframes)
2. A pointer containing the concatenation indicator (if the NE is able to accommodate contiguously concatenated signals, see [Section 3.5.1.4](#)) that is received identically in three consecutive frames.

Therefore, the intended interpretation of the LOP defect detection requirement is that a pointer processor detects an LOP defect at the end of any N frame (VT superframe) window ($8 \leq N \leq 10$) that does not contain at least one set of three consecutive frames (VT superframes) with either a valid pointer as defined above, or all-ones pointer words (in which case the all-ones pointer would be relayed or an AIS defect would be detected instead).

Other possible definitions of a “valid” pointer could include a single pointer that exactly matches the current “valid” pointer value, or a pointer that causes the receiving NE to locate the start of the STS or VT SPE at a new location according to the pointer interpretation rules in [Sections 3.5.1.6](#) and [3.5.2.6](#). For example, a received pointer that is interpreted as an increment or a decrement from the previous “valid” pointer value could establish a new “valid” pointer value.

Although various definitions and interpretations of the requirement are acceptable, some limitations are needed. Therefore, as a minimum, a pointer processor must detect an LOP defect if it receives 10 consecutive frames (VT superframes) in which all of the following are true:

1. None of the pointers contain the same pointer value as the current “valid” pointer value, along with the N-bits set to their normal value and, for VT pointers, constant or correct size bits
2. None of the pointers can be interpreted as indicating an increment or decrement from the current “valid” pointer value (thus establishing a new current valid pointer value)
3. No three consecutive frames (VT superframes) contain a constant, in-range pointer value, the N-bits set to their normal value and, for VT pointers, constant or correct size bits
4. None of the pointers have the NDF set (along with an in-range pointer value and, for VT pointers, constant or correct size bits), or all of the pointers have the NDF set

-
6. In addition to these possible defects, the STS PTE should also detect STS path BIP-8 (B3) errors at a very high rate, and this would be expected to eventually lead to the generation of an STS Path Unavailable Second Threshold Crossing Alert (UAS-P TCA, see [Section 6.2.2](#)).

5. For LTE that processes STS pointers (STS PTE that processes VT pointers) and meets the all-ones pointer relay objective (see Sections 3.5.1.5 and 3.5.2.5), none of the pointers is all-ones
6. For STS (VT) PTE, or LTE that processes STS pointers (STS PTE that processes VT pointers) and does not meet the all-ones pointer relay objective, no three consecutive frames (VT superframes) contain all-ones pointers.

-
- R6-77 [431v3]** STS PTE and LTE that processes the STS pointers shall terminate an LOP-P defect when the STS has a valid pointer with a normal NDF that it can accommodate, or a valid concatenation indicator that it can accommodate, in three consecutive frames.
- R6-78 [432v2]** STS PTE shall terminate an LOP-P defect when it detects an AIS-P defect.
- R6-79 [945]** LTE that processes STS pointers shall terminate an LOP-P defect when it relays an all-ones STS pointer.
- R6-80 [946]** VT PTE and STS PTE that processes VT pointers shall terminate an LOP-V defect when the VT has a valid pointer with a normal NDF in three consecutive VT superframes.
- R6-81 [947]** VT PTE shall terminate an LOP-V defect when it detects an AIS-V defect.
- R6-82 [948]** STS PTE that processes VT pointers shall terminate an LOP-V defect when it relays an all-ones VT pointer.
- R6-83 [433v2]** A SONET NE shall declare an LOP-P failure when an LOP-P defect persists for 2.5 (± 0.5) seconds. Upon declaring an LOP-P failure, the NE shall set an LOP-P failure indication and send an alarm message to an OS unless the condition in **R6-312 [626v3]** applies.
- R6-84 [949]** A SONET NE shall declare an LOP-V failure when an LOP-V defect persists for 2.5 (± 0.5) seconds. Upon declaring an LOP-V failure, the NE shall set an LOP-V failure indication and send an alarm message to an OS unless the condition in **R6-312 [626v3]** applies.
- R6-85 [434]** For the purposes of trunk conditioning, SONET NEs that contain DS0 PTE or VT PTE that supports the rearrangement of the DS0 channels in byte-synchronously mapped DS1s shall use the integration technique described in **R6-59 [419]** to declare LOP-P and LOP-V failures. Upon declaring an LOP-P or LOP-V failure, the NE shall perform the actions listed in **R6-83 [433v2]** or **R6-84 [949]**.
- R6-86 [435v2]** A SONET NE shall clear an LOP-P failure when an LOP-P defect is absent for 10.0 (± 0.5) seconds. Upon clearing the LOP-P failure, the SONET NE shall clear the LOP-P failure indication and send a clear message to an OS (if the failure was reported to the OS).

R6-87 [950] A SONET NE shall clear an LOP-V failure when an LOP-V defect is absent for 10.0 (± 0.5) seconds. Upon clearing the LOP-V failure, the SONET NE shall clear the LOP-V failure indication and send a clear message to an OS (if the failure was reported to the OS).

6.2.1.1.4 *Equipment Failures*

This GR does not define equipment failure states, since they are largely implementation-dependent. However, it does list a minimum set of conditions to be reported as alarms.

R6-88 [436] Equipment failures shall be classified as either Service-Affecting (SA) or Non-Service-Affecting (NSA), depending on whether they affect the services that the equipment transports.

R6-89 [437] Equipment failures shall be classified as critical, major or minor.

R6-90 [438] Because hardware designs vary, the report of the equipment failure shall describe the failure condition.

R6-91 [439v2] The NE shall be able to declare the following equipment failures (as a minimum):

- Fuse or power circuit failures
- Synchronization equipment failures
- Protection switching equipment failures
- Central Processor Unit (CPU) failures
- Local nonvolatile backup memory failures
- SONET signal origination and termination equipment failures
- Receiver failures (e.g., optical detector failures)
- Transmitter failures (e.g., light source failure, including laser failures)
- OFA or OFA pump failures (if a BA or PA is supported)
- Non-SONET signal (e.g., DS_n) origination and termination equipment failures
- Switching matrix module failures (if cross-connect functionality is provided)
- DCC hardware failures (also see [Section 6.2.1.1.7](#))
- Manual removal of in-service (i.e., active) equipment.

R6-92 [440] Upon declaring an equipment failure, a SONET NE shall:

- Switch to duplex or standby equipment, if available
- Set a local indication

- Send an alarm message to an OS.

In addition, certain equipment failures cause AIS to be generated (see [Section 6.2.1.2](#)), and an NE may declare and report equipment failures that are not specifically listed in this section.

R6-93 [441] Upon clearing an equipment failure, a SONET NE shall clear the equipment failure indication and send a clear message to the OS.

CR6-94 [442] A SONET NE may be required to detect and report certain environmental conditions in some applications [e.g., NEs in a Controlled Environmental Vault (CEV)].

The CPU of a SONET NE may be provided in a redundant manner so that a standby can be automatically switched into service upon failure of the main CPU. GRs and TRs covering specific types of NEs contain requirements regarding this hardware redundancy feature for the CPU and other equipment (e.g., DCC or LCN termination).

O6-95 [443] A SONET NE should detect operating system or other software errors, and report them to an OS, independently of CPU hardware failures.

6.2.1.1.5 Loss of Synchronization

This section is concerned with Loss of Synchronization failures, which result when a provisioned timing reference or derived DS1 source signal fails or becomes unavailable.

R6-96 [444v3] A SONET NE shall declare a Loss of Synchronization failure for a provisioned reference or derived DS1 source when that reference or source is determined to be failed or unavailable for timing purposes (see [Section 5.4.6.1](#)) and, in the case of an LOS, OOF, LOF or AIS, the corresponding defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, the NE shall set a Loss of Synchronization failure indication and send a message to an OS. The message shall include an indication of the type of synchronization signal (i.e., timing reference or derived DS1 source), reference switches, and the reason for failure (e.g., LOS, LOF or OOF, AIS, invalid synchronization status message).

R6-97 [445v2] A SONET NE shall clear a Loss of Synchronization failure when the condition or defect that caused it to be declared is absent for 10.0 (± 0.5) seconds. Upon clearing the Loss of Synchronization failure, a SONET NE shall clear the Loss of Synchronization failure indication and send a clear message to the appropriate OS.

Synchronization equipment failures (as opposed to synchronization signal failures) are listed with other equipment failures in [Section 6.2.1.1.4](#). In addition, Section 8 of GR-1244-CORE provides related maintenance criteria applicable to all NEs that contain synchronized clocks. Those criteria cover such issues as default alarm levels for various types of failures, synchronization-related events that should be autonomously reported, and information that can be retrieved on demand.

R6-98 [1125] A SONET NE shall meet the criteria in Section 8 of GR-1244-CORE related to alarms, reports and control commands.

6.2.1.1.6 APS Troubles

Five types of failures related to the operation of the APS channel are defined for LTE that supports linear APS. These are Protection Switching Byte failures ([Section 6.2.1.1.6.A](#)), Channel Mismatch failures ([Section 6.2.1.1.6.B](#)), APS Mode Mismatch failures ([Section 6.2.1.1.6.C](#)), Invalid APS Mode failures ([Section 6.2.1.1.6.D](#)), and Far-End Protection Line failures ([Section 6.2.1.1.6.E](#)). In addition, [Section 6.2.1.1.6.F](#) contains criteria related to the alarm generated by an NE when it receives an AIS-L and cannot perform a protection switch, and the messages that are generated when BER-based SF and SD conditions are detected.

The criteria in [Sections 6.2.1.1.6.A](#), [6.2.1.1.6.B](#) and [6.2.1.1.6.E](#) apply to LTE that is *operating* in a linear APS mode other than the 1+1 unidirectional mode, while the criteria in [Section 6.2.1.1.6.C](#) apply to LTE that is *provisioned to operate* in a linear APS mode other than the 1+1 unidirectional mode and the criteria in [Sections 6.2.1.1.6.D](#) and [6.2.1.1.6.F](#) apply to LTE that is provisioned to operate in any linear APS mode. For example, LTE that is provisioned to operate in the 1:1 bidirectional mode, but that is actually operating in the 1+1 unidirectional mode (because that is the mode indicated by the far-end LTE), must meet the criteria in [Sections 6.2.1.1.6.C](#), [6.2.1.1.6.D](#) and [6.2.1.1.6.F](#), but does not need to meet those in [Sections 6.2.1.1.6.A](#), [6.2.1.1.6.B](#) and [6.2.1.1.6.E](#).

Since all LTE that uses linear APS must transmit the appropriate codes in the APS channel, LTE that is operating in the 1+1 unidirectional mode can expect to receive those codes, and may (but is not required to) use them to detect and declare the defects and failures discussed in [Sections 6.2.1.1.6.A](#), [6.2.1.1.6.B](#) and [6.2.1.1.6.E](#). Note however, that those defects and failures must not affect the operation of the 1+1 unidirectional system, and that the LTE must meet **R5-103 [207]** and **O5-104 [208]** (even though it is not operating in the bidirectional mode) to avoid declaring extraneous alarms in some situations.

6.2.1.1.6.A Protection Switching Byte Failure

Unless it is operating in the 1+1 unidirectional mode, LTE must monitor the incoming K1 byte for Protection Switching Byte failures. A Protection Switching Byte defect occurs when either an inconsistent APS byte or an invalid code is detected. An inconsistent APS byte occurs when no three consecutive K1 bytes of the last 12 successive frames are identical, starting with the last frame containing a previously consistent byte. An invalid code occurs when the incoming K1 byte

contains an unused code (see [Table 5-6](#)), or a code irrelevant for the specific switching operation (e.g., Reverse Request while no switching request is outstanding) in three consecutive frames. An invalid code also occurs when the incoming K1 byte contains an invalid channel number in three consecutive frames. (Also see the discussion in [Section 5.3.5.5](#).)

R6-99 [446v2] LTE operating in a linear APS mode other than the 1+1 unidirectional mode shall detect a Protection Switching Byte defect within 50 ms of the occurrence of either an inconsistent APS byte or an invalid code, unless the condition for terminating the defect occurs before the defect is detected.

R6-100 [447] LTE shall not detect a Protection Switching Byte defect when it has detected an AIS-L defect.

R6-101 [448v2] LTE shall terminate the Protection Switching Byte defect within 50 ms of the occurrence of three consecutive, identical and valid APS codes, unless the condition for detecting the defect occurs before the defect is terminated.

R6-102 [449] LTE shall not terminate a Protection Switching Byte defect when it has detected the AIS-L defect.

R6-103 [450] An NE shall declare a Protection Switching Byte failure when a Protection Switching Byte defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall perform the following actions:

1. A Protection Switching Byte failure indication shall be set and an alarm message shall be sent to an OS
2. If a working channel that was being selected from the protection line is switched back to the working line (see [Section 5.3.5.5](#)), a message shall be sent to an OS indicating the switch back to the working line.

R6-104 [451] An NE shall clear the Protection Switching Byte failure when the Protection Switching Byte defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, it shall clear the Protection Switching Byte failure indication and send an alarm clear message to an OS.

6.2.1.1.6.B Channel Mismatch Failure

Unless it is operating in the 1+1 unidirectional mode, LTE must monitor the channel numbers in its transmitted K1 byte (bits 5 through 8) and the received K2 byte (bits 1 through 4) for Channel Mismatch. (In the 1+1 unidirectional mode, each end operates independently and the LTE does not need to monitor for Channel Mismatch.) If the channel number in the received K2 byte is not identical to the channel number transmitted in the K1 byte, there is a mismatch.

Under normal conditions, a mismatch will occur each time the LTE changes the channel number on its transmitted K1 byte (i.e., when the channel with the highest priority switch request changes). A mismatch could also occur if the channel number on the incoming K2 byte changes (e.g., due to a failure at the far-end LTE).

R6-105 [452] LTE operating in a linear APS mode other than the 1+1 unidirectional mode shall detect a Channel Mismatch defect if the channel numbers in the transmitted K1 byte and the received K2 byte do not match for 50 ms.

O6-106 [453] LTE operating in a linear APS mode other than the 1+1 unidirectional mode should detect a Channel Mismatch defect if the channel numbers in the transmitted K1 byte and the received K2 byte are mismatched in three consecutive frames.

[Section 5.3.5.4](#) discusses the actions that LTE is required to take when it detects or terminates a Channel Mismatch defect.

R6-107 [454] LTE shall not detect a Channel Mismatch defect when it has detected an AIS-L defect.

R6-108 [455] LTE shall terminate the Channel Mismatch defect if the channel numbers in the transmitted K1 byte and the received K2 byte match in three consecutive frames.

R6-109 [456] LTE shall not terminate a Channel Mismatch defect when it has detected an AIS-L defect.

R6-110 [457] An NE shall declare a Channel Mismatch failure if the Channel Mismatch defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set a Channel Mismatch failure indication and send an alarm message to an OS.

R6-111 [458] An NE shall clear the Channel Mismatch failure if the Channel Mismatch defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, it shall clear the Channel Mismatch failure indication and send an alarm clear message to the OS.

6.2.1.1.6.C APS Mode Mismatch Failure

Unless it is provisioned to operate in the 1+1 unidirectional mode, LTE must monitor the mode of operation indicators in the incoming K2 byte for APS Mode Mismatch failures. As discussed in [Section 5.3.5.2.3](#), the APS mode information in bits 5 through 8 of the received K2 byte is mismatched when either:

1. LTE provisioned for 1+1 protection switching receives an indication from the far-end LTE (in bit 5 of K2) that it is provisioned for 1:n (or vice versa), or
2. LTE provisioned for bidirectional switching receives an indication from the far-end LTE (in bits 6 through 8 of K2) that it is provisioned for unidirectional switching (or vice versa).

Codes other than '101' and '100' in the incoming K2 bits 6 through 8 are ignored for the purposes of monitoring for the APS Mode Mismatch defect (although also see [Section 6.2.1.1.6.D](#)).

-
- R6-112 [459]** LTE provisioned to operate in a linear APS mode other than the 1+1 unidirectional mode shall detect an APS Mode Mismatch defect within 100 ms of receiving the first of five consecutive samples (which may or may not be consecutive frames) with identical mode information (either in bit 5 of K2 or bits 6 through 8 of K2) that is mismatched, as defined above, unless the condition for terminating the defect occurs before the defect is detected.
- R6-113 [460]** LTE shall not detect an APS Mode Mismatch defect when it has detected an AIS-L defect.
- R6-114 [461v2]** LTE shall terminate an APS Mode Mismatch defect within 50 ms of receiving the first of five consecutive samples (which may or may not be consecutive frames) with identical mode information that is not mismatched as defined above, unless the condition for detecting the defect occurs before the defect is terminated.
- R6-115 [462]** LTE shall not terminate an APS Mode Mismatch defect when it has detected an AIS-L defect.
- R6-116 [463]** An NE shall declare an APS Mode Mismatch failure when an APS Mode Mismatch defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set the APS Mode Mismatch failure indication and send an alarm message to an OS.
- R6-117 [464]** An NE shall clear the APS Mode Mismatch failure if the APS Mode Mismatch defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, it shall clear the APS Mode Mismatch failure indication and send an alarm clear message to the OS.
-

6.2.1.1.6.D Invalid APS Mode Failure

As discussed in [Section 5.3.5.2](#), the sustained reception of '000', '001', '010' or '011' in K2 bits 6 through 8 is likely to be indicative of a serious provisioning or connection problem. Therefore, LTE that is provisioned to operate in any of the possible linear APS modes needs to monitor those bits for Invalid APS Mode failures.

-
- R6-118 [1191]** LTE provisioned to operate in any linear APS mode shall detect an Invalid APS Mode defect within 100 ms of receiving the first of five consecutive samples (which may or may not be consecutive frames) with '000', '001', '010' or '011' in bits 6 through 8 of K2, unless the condition for terminating the defect occurs before the defect is detected.
- R6-119 [1192]** LTE shall terminate an Invalid APS Mode defect within 50 ms of receiving the first of five consecutive samples (which may or may not be consecutive frames) with '100' or '101' in bits 6 through 8 of K2, unless the condition for detecting the defect occurs before the defect is terminated.

- R6-120 [1193]** LTE shall not terminate an Invalid APS Mode defect when it has detected an AIS-L or RDI-L defect.
- R6-121 [1194]** An NE shall declare an Invalid APS Mode failure when an Invalid APS Mode defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set the Invalid APS Mode failure indication and send an alarm message to an OS.
- R6-122 [1195]** An NE shall clear the Invalid APS Mode failure if the Invalid APS Mode defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, it shall clear the Invalid APS Mode failure indication and send an alarm clear message to the OS.

Note that in some cases LTE in a properly provisioned linear APS system may insert the '000' code in K2 bits 6 through 8 for as long as 200 ms after deactivating RDI-L (see [Section 6.2.1.3](#)), and thus it could be considered undesirable to detect an Invalid APS Mode defect in less than 200 ms. On the other hand, the specification of a 100-ms maximum detection time in **R6-118 [1191]** provides consistency between that requirement and **R6-112 [459]** on APS Mode Mismatch defect detection. In addition, the only immediate action that an NE is supposed to perform upon detecting an Invalid APS Mode defect is the activation of a timer to determine if it needs to declare the corresponding failure. Therefore, in cases where a valid '000' code is temporarily inserted, the receiving LTE should simply detect and (almost immediately) terminate the Invalid APS Mode defect.

6.2.1.1.6.E *Far-End Protection-Line Failure*

Unless it is operating in the 1+1 unidirectional mode, LTE must monitor the K1 byte for Far-End Protection-Line failures. When LTE receives an SF code for the protection line, it knows that the far-end LTE is no longer receiving its request, or (in bidirectional operation) that the far-end LTE considers the near-end LTE's request to be invalid. In unidirectional operation, the near-end LTE can maintain any existing switch (e.g., continue to select a working or extra traffic channel from the protection line) if the far-end LTE maintains and indicates the appropriate bridge in K2 bits 1 through 4 (see **R5-93 [198]**). In bidirectional operation, any working channel that had been switched to the protection line is switched back to its working line (since the SF on protection request is higher priority than any request that would cause the working channel to be switched to the protection line).

-
- R6-123 [465]** LTE operating in a linear APS mode other than the 1+1 unidirectional mode shall detect a Far-End Protection-Line defect when it receives three consecutive K1 bytes with the code indicating "SF on the protection line."
- R6-124 [466]** LTE shall terminate the Far-End Protection-Line defect when it receives three consecutive, identical and valid K1 bytes with any code other than "SF on the protection line."

R6-125 [467v2] An NE shall declare a Far-End Protection-Line failure when a Far-End Protection-Line defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, the NE shall set a Far-End Protection-Line failure indication and (if it is a Reported failure) send an alarm message to an OS.

R6-126 [468] An NE shall clear the Far-End Protection-Line failure if the Far-End Protection-Line defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, it shall clear the Far-End Protection-Line failure indication and send an alarm clear message to the OS (if the failure was reported to the OS).

6.2.1.1.6.F Other APS Conditions

Although they are not defined as defects/failures (see [Section 6.2.1](#)), BER-based SF and SD conditions that are used for protection switching purposes can have a direct impact on the transport of traffic, and therefore the user needs to be alerted when those conditions exist. The following requirements are intended to support this. Note that to avoid the generation of multiple autonomous messages in situations where a BER-based SF or SD condition is repetitively detected and cleared over a relatively short time period, the following requirements include “soaking times” similar to those used in most of the failure declaration and clearing criteria in the various subsections of [Section 6.2.1](#).

R6-127 [1064v2] When a BER-based SF condition persists for 2.5 (± 0.5) seconds, an NE shall set an SF BER indication and, as a default, send a message to an OS.

R6-128 [1065v2] When a BER-based SD condition persists for 2.5 (± 0.5) seconds, an NE shall set an SD BER indication and, as a default, send a message to an OS.

R6-129 [1066v2] An NE shall clear an SF BER indication 10 (± 0.5) seconds after detecting that the BER is less than the SF clearing threshold, assuming that it does not detect that the BER is greater than the SF threshold during that time period (see [Section 5.3.4](#)). Upon clearing an SF BER indication, the NE shall send a clear message to the OS (if the SF BER was reported to the OS).

R6-130 [1067v2] An NE shall clear an SD BER indication 10 (± 0.5) seconds after detecting that the BER is less than the SD clearing threshold, assuming that it does not detect that the BER is greater than the SD threshold during that time period (see [Section 5.3.4](#)). Upon clearing an SD BER indication, the NE shall send a clear message to the OS (if the SD BER was reported to the OS).

R6-131 [1068v2] As a default, SF BER and SD BER indications shall be Not Alarmed.

R6-132 [469] If LTE that provides the linear APS feature is unable to perform automatic protection switching upon receiving an AIS-L signal, it shall send an SA alarm to an OS, indicating the inability to perform a protection switch.

6.2.1.1.7 DCC Failure

DCC hardware failures are listed with other equipment failures in [Section 6.2.1.1.4](#), which covers the declaration and clearing of all equipment failures. This section describes the DCC failure in more detail.

A DCC failure is defined as either a DCC hardware failure (as in [Section 6.2.1.1.4](#)) or a failure of the line carrying the DCC. In the first case, standby DCC equipment can be provided to protect against a service DCC equipment failure. In the second case, the DCC protection scheme described in [Section 8.3.1.3](#) is used to recover from a failure on the line carrying the primary DCC. End System-Intermediate System (ES-IS) and IS-IS routing protocols can be used as a protection scheme when alternate routes are available to route messages around failures where both primary and secondary DCCs either have been lost or are unavailable.

R6-133 [470] A SONET NE shall be able to declare a DCC failure. Upon declaring the failure, it shall set a DCC failure indication and send a message to an OS.

R6-134 [471] Upon clearing a DCC failure, a SONET NE shall clear the DCC failure indication and send a clear message to an OS.

6.2.1.1.8 Signal Label Mismatch

A received STS or VT signal label (the C2 byte, or V5 bits 5 through 7 and possibly Z7 bit 1, respectively) is considered mismatched if it does not equal either a label value corresponding to the locally provisioned PTE functionality or the label value corresponding to the equipped, non-specific code (see [Tables 3-2, 3-3 and 3-4](#)). Two types of defects, PLM defects and UNEQ defects, are currently defined at each path layer. [Tables 6-2 and 6-3](#) identify the conditions corresponding to these two types of defects. In those tables, the “Received Payload Label” corresponds to the signal label in the STS and VT paths received on the incoming signal, while the “Provisioned Functionality” corresponds to the mapping used by the STS or VT PTE. Only in-service, provisioned PTE can detect these defects. PTE is considered provisioned if it has been configured for a mapping (or only supports one mapping) and has been assigned a timeslot (or is hardwired to a specific timeslot) in a SONET signal. When an PLM or UNEQ defect is detected, the appropriate AIS is sent to downstream equipment and an ERDI code (if supported) is sent to upstream equipment.

Note that if STS PTE has detected an AIS-P or LOP-P defect on its incoming signal, the C2 byte cannot be accessed to monitor for PLM-P or UNEQ-P defects. Therefore, it can neither detect nor terminate a PLM-P or UNEQ-P defect. Similarly, if VT PTE has detected an AIS-V or LOP-V defect, the V5 and Z7 bytes cannot be monitored for PLM-V or UNEQ-V defects (see [Section 6.2.1.8.2](#)).

6.2.1.1.8.A STS Payload Label Mismatch

R6-135 [472] STS PTE shall detect a PLM-P defect within 250 ms of the onset of at least five consecutive samples (which may or may not be consecutive frames) of mismatched STS signal labels (C2 byte), as specified in [Table 6-2](#).

For testing conformance to this requirement, it is sufficient to transmit a continuous string of identical, mismatched STS signal labels in the C2 byte for at least 250 ms.

O6-136 [473] STS PTE should detect a PLM-P defect immediately upon receipt of five contiguous frames with mismatched STS signal labels, as specified in [Table 6-2](#).

R6-137 [476] STS PTE shall terminate a PLM-P defect within 250 ms of detecting the onset of at least five consecutive samples (which may or may not be consecutive frames) of matched STS signal labels, as specified in [Table 6-2](#).

For testing conformance to this requirement, it is sufficient to transmit a continuous string of matched STS signal labels for at least 250 ms.

O6-138 [477] STS PTE should terminate a PLM-P defect immediately upon receipt of five contiguous frames with matched STS signal labels, as specified in [Table 6-2](#).

R6-139 [478v2] STS PTE shall terminate a PLM-P defect upon detecting an UNEQ-P defect.

R6-140 [479] An NE shall declare a PLM-P failure if a PLM-P defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set a PLM-P failure indication and send an alarm message to an OS unless the condition in **R6-312 [626v3]** applies.

R6-141 [480v2] An NE shall clear the PLM-P failure if the PLM-P defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the PLM-P failure indication and send a clear message to an OS (if the failure was reported to the OS).

6.2.1.1.8.B STS Path Unequipped

R6-142 [481] STS PTE shall detect an UNEQ-P defect within 10 ms of the onset of at least five consecutive samples (which may or may not be consecutive frames) of unequipped STS signal labels (C2 byte), as specified in [Table 6-2](#).

O6-143 [482] STS PTE should detect an UNEQ-P defect immediately upon receipt of five contiguous frames with unequipped STS signal labels, as specified in [Table 6-2](#).

R6-144 [485v2] STS PTE shall terminate an UNEQ-P defect within 10 ms of the onset of at least five consecutive samples (which may or may not be consecutive frames) of STS signal labels that are not unequipped or all-ones, as specified in [Table 6-2](#).

- O6-145 [486v2]** STS PTE should terminate an UNEQ-P defect immediately upon receipt of five contiguous frames with STS signal labels that are not unequipped or all-ones, as specified in [Table 6-2](#).
- R6-146 [488]** An NE shall declare an UNEQ-P failure if an UNEQ-P defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set an UNEQ-P failure indication and send an alarm message to an OS unless the condition in **R6-312 [626v3]** applies.
- R6-147 [489v2]** An NE shall clear the UNEQ-P failure if the UNEQ-P defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the UNEQ-P failure indication and send a clear message to an OS (if the failure was reported to the OS).
-

When a service is first activated, it is possible that STS PTE would detect an UNEQ-P defect until the service is properly and completely provisioned and activated. As an option, to avoid declaring and reporting UNEQ-P failures when a service is first activated, a SONET NE may be designed such that it does not detect an UNEQ-P defect until the service is activated.

Table 6-2 STS Signal Label Mismatch Defect Conditions

Provisioned STS PTE Functionality	Received Payload Label (C2 Byte, in hex format)	Defect
None ^a	Anything	None
Any Equipped functionality ^b	Unequipped or Unequipped Supervisory (00)	UNEQ-P
Any Equipped functionality ^b	Equipped – Non Specific (01)	None (Matched)
Equipped – Non Specific	A value corresponding to any Payload Specific functionality (02 to E0, FD, or FE)	None (Matched)
Any Payload Specific functionality ^b	A value corresponding to the same Payload Specific functionality as the provisioned functionality (02 to E0, FD, or FE)	None (Matched)
Any Payload Specific functionality ^b	A value corresponding to a different Payload Specific functionality than the provisioned functionality (02 to E0, FD, or FE)	PLM-P
Equipped – Non Specific, or VT-Structured STS-1	PDI, 1 to 27 VTn defects (E1 to FB)	None (Matched) or PDI-P ^d
Any Payload Specific functionality except VT-Structured STS-1 ^b	PDI, 1 to 27 VTn defects (E1 to FB)	PLM-P
Any Equipped functionality ^b	PDI, 28 VT1.5 defects or 1 non-VT payload defect (FC)	None (Matched) or PDI-P ^e
Any Equipped functionality ^b	Reserved ^c (FF)	No change ^f

Notes:

- a. This entry corresponds to the case when an NE is not provisioned to expect a signal. See [Section 3.3.2.4](#).
- b. See [Table 3-2](#) for the currently assigned STS signal label values.
- c. The all-ones signal label is reserved in ITU-T G.707 for a VC AIS function associated with the higher order tandem connection feature (and is not currently defined for any use in this document). It is not expected to be assigned to a specific payload.
- d. See [Section 6.2.1.4.1](#).
- e. This entry [i.e., None (Matched) or PDI-P] applies even if no PDI-P generation function has been defined for the particular functionality/mapping supported by the STS PTE (see [Section 6.2.1.4.1](#)).
- f. That is, the detection of ‘FF’ in the STS signal label shall not cause a new UNEQ-P or PLM-P defect to be detected for that path, or cause an existing UNEQ-P or PLM-P defect to be terminated.

6.2.1.1.8.C VT Payload Label Mismatch

- R6-148 [491v2]** VT PTE shall detect a PLM-V defect within 250 ms of the onset of at least five consecutive samples (which may or may not be consecutive VT superframes or extended VT signal label multiframes) of mismatched VT signal labels, as specified in [Table 6-3](#).

For testing conformance to this requirement, it is sufficient to transmit a continuous string of identical, mismatched VT signal labels in bits 5 through 7 of the V5 byte or Z7 bit 1 for at least 250 ms.

- O6-149 [492v2]** VT PTE should detect a PLM-V defect immediately upon receipt of five contiguous VT superframes or extended VT signal label multiframes with mismatched VT signal labels, as specified in [Table 6-3](#).

Note that an NE that meets **O6-149 [492v2]** would have a nominal detection time of 2.5 ms for a PLM-V defect detected as a result of a mismatch in V5 bits 5 through 7, or 80 ms for a defect detected as a result of a mismatch of an extended VT signal label.

- R6-150 [495v2]** VT PTE shall terminate a PLM-V defect within 250 ms of detecting the onset of at least five consecutive samples (which may or may not be consecutive VT superframes or extended VT signal label multiframes) of matched VT signal labels, as specified in [Table 6-3](#).

For testing conformance to this requirement, it is sufficient to transmit a continuous string of matched VT signal labels for at least 250 ms.

- O6-151 [496v2]** VT PTE should terminate a PLM-V defect immediately upon receipt of five contiguous VT superframes or extended VT signal label multiframes with matched VT signal labels, as specified in [Table 6-3](#).

- R6-152 [497v2]** VT PTE shall terminate the PLM-V defect immediately upon detecting an UNEQ-V defect on the incoming signal.

- R6-153 [498]** An NE shall declare a PLM-V failure when a PLM-V defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set a PLM-V failure indication and send an alarm message to an OS unless the condition in **R6-312 [626v3]** applies.

- R6-154 [499]** An NE shall clear a PLM-V failure when the PLM-V defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the PLM-V failure indication and send a clear message to the OS (if the failure was reported to the OS).

6.2.1.1.8.D VT Path Unequipped

-
- R6-155 [501]** VT PTE shall detect an UNEQ-V defect within 10 ms of the onset of at least five consecutive samples (which may or may not be consecutive VT superframes) of unequipped VT signal labels, as specified in [Table 6-3](#).
- O6-156 [502]** VT PTE should detect an UNEQ-V defect immediately upon receipt of five contiguous VT superframes with unequipped VT signal labels, as specified in [Table 6-3](#).
- R6-157 [505v2]** VT PTE shall terminate an UNEQ-V defect within 10 ms of the onset of at least five consecutive samples (which may or may not be consecutive VT superframes) of VT signal labels that are not unequipped or all-ones, as specified in [Table 6-3](#).
- O6-158 [506v2]** VT PTE should terminate an UNEQ-V defect immediately upon receipt of five contiguous VT superframes with VT signal labels that are not unequipped or all-ones, as specified in [Table 6-3](#).
- R6-159 [508]** An NE shall declare an UNEQ-V failure when an UNEQ-V defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set an UNEQ-V failure indication and send an alarm message to an OS unless the condition in **R6-312 [626v3]** applies.
- R6-160 [509]** An NE shall clear an UNEQ-V failure when the UNEQ-V defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the UNEQ-V failure indication and send a clear message to the OS (if the failure was reported to the OS).
-

When a service is first activated, it is possible that VT PTE would detect an UNEQ-V defect until the service is properly and completely provisioned and activated. As an option, to avoid declaring and reporting UNEQ-V failures when a service is first activated, a SONET NE may be designed such that it does not detect an UNEQ-V defect until the service is activated.

Table 6-3 VT Signal Label Mismatch Defect Conditions

Provisioned VT PTE Functionality	Received Payload Label (V5 Bits 5 Through 7)	Defect
None (zero) ^a	Anything	None
Any Equipped functionality ^b	Unequipped or Unequipped Supervisory (000)	UNEQ-V
Any Equipped functionality ^b	Equipped – Non Specific (001)	None (Matched)
Equipped – Non Specific	A value corresponding to any Payload Specific functionality (010 to 110)	None (Matched)
Any Payload Specific functionality ^b	A value corresponding to the same Payload Specific functionality as the provisioned functionality (010 to 110, and if 101, an extended VT signal label of 08 to FE)	None (Matched)
Any Payload Specific functionality ^b	A value corresponding to a different Payload Specific functionality than the provisioned functionality (010 to 110, and if 101, an extended VT signal label of 08 to FE)	PLM-V
Any Payload Specific functionality ^b	An invalid or reserved extended VT signal label (101 and either no MFAS ^c found, or an extended VT signal label of 00 to 07 or FF)	PLM-V
Any Equipped functionality ^b	Non-extended Reserved ^d (111)	No change ^e

Notes:

- a. This entry corresponds to the case when an NE is not provisioned to expect a signal. See [Section 3.3.3](#).
- b. See [Tables 3-4](#) and [3-5](#) for the currently assigned VT and extended VT signal label values.
- c. See the Z7 byte description in [Section 3.3.3](#).
- d. The all-ones (non-extended) signal label is reserved in ITU-T G.707 for a VC AIS function associated with the lower order tandem connection feature (but currently not defined in this document). It is not expected to be assigned to a specific payload.
- e. That is, the detection of ‘111’ in the VT signal label shall not cause a new UNEQ-V or PLM-V defect to be detected for that path, or cause an existing UNEQ-V or PLM-V defect to be terminated.

6.2.1.1.9 Trace Identifier Mismatch

This section provides information and criteria related to the use of the STS and VT path trace strings (carried in the J1 and J2 bytes) for TIM-P and TIM-V defect detection and failure declaration purposes. To reduce redundancies, much of the information that is common to both the STS and VT cases is listed below.

- Similar to an UNEQ defect, a TIM defect is considered a connectivity defect because it is expected to be caused by provisioning problems (e.g., incorrect cross-connections) within the network. In the case of a TIM defect, the incorrect provisioning causes PTE at one NE to be connected through the network to the wrong far-end PTE. At this time, TIM defects and failures have been defined in SONET only for STS and VT paths; however, they could possibly be defined at the Section layer in the future. In addition, in order to support certain applications, avoid unnecessary alarms during path setup and rearrangement procedures, and coexist with equipment that does not support the STS or VT path trace provisioning capabilities necessary for reliable TIM-P or TIM-V detection purposes (e.g., equipment that does not meet **R6-162 [997v3]** and **R6-164 [1069v2]**), an NE that supports the TIM-P or TIM-V detection feature must be capable of having that feature activated and deactivated on a per-path basis.
- In general, the possible impacts (both positive and negative) should be carefully considered before the detection of TIM defects is activated for a particular STS or VT path. For example, the user should be aware that due to interactions between various criteria (including the PM parameter definitions in [Section 6.2.2](#)), the activation of TIM defect detection for a path could result in discrepancies between different sets of PM data accumulated for that path.⁷ In addition, activation of TIM-P defect detection for a VT-accessed STS-1 path in a BLSR could cause an unnecessary loss of traffic in certain failure scenarios.
- Prior to GR-253-CORE, Issue 2, Revision 2, a received STS path trace string was expected to be monitored and used for diagnostics purposes only, while prior to Issue 4 the use of the VT path trace string was indicated to be FFS. When used for diagnostics purposes the received path trace string may be either monitored for changes or compared to an “expected” path trace string (as is done for TIM detection purposes), and in the latter case persistent mismatches may be reported to an OS. However, the detection of a mismatch on a path trace string that is being monitored for diagnostics purposes does not cause the generation

7. If ERDI is supported for a particular STS or VT path, then TIM defects/failures are required to affect the accumulation of certain PM parameters. For a particular path this effect is both direct (i.e., based on the presence or absence of the TIM defect or failure) at the PTE that is accumulating the near-end path PM data, and indirect (e.g., via the detection of ERDI Connectivity defects) at any upstream NEs that are accumulating the corresponding far-end path or intermediate-path PM data. On the other hand, any intermediate NEs that are accumulating near-end intermediate-path PM involving that path are currently not required to be able to monitor for TIM defects/failures. This could lead to significant discrepancies between the various sets of PM data accumulated for a path, and therefore it has been suggested that it may be desirable to revisit some of the relevant criteria (e.g., the criteria that define the TIM defect, the ERDI trigger criteria, the PM parameter definitions). However, at this time no changes are expected to be made.

of AIS downstream or RDI upstream, and does not affect the accumulation of any PM parameters. In general (and assuming ERDI is supported), all of those actions are required to be taken when a TIM defect is detected (e.g., see [Sections 6.2.1.2.3, 6.2.1.2.4, 6.2.1.3.2 and 6.2.2.5](#)).⁸

- In several cases, the criteria in the following subsections refer to path trace messages that meet the specifications in ANSI T1.269. Among the characteristics of such a message is that it begins with a header byte consisting of a '1' in bit 1 (as a frame start marker) and a static CRC-7 in bits 2 through 8. For a 16-byte message this leaves 15 bytes available to be provisioned by the user, while for a 64-byte message it leaves 63 bytes. These bytes are divided into a number of fixed-length data elements defined to carry specific information or codes, and if the provisioned information for a particular data element does not utilize all of the available bytes, the remaining bytes are padded with ASCII NULL or SPACE characters. In general, it is expected that the addition of the pad characters, as well as the calculation of the CRC-7 bits, will need to be able to be performed by the NE.
- Since TIM defects are expected to occur primarily as a result of provisioning errors (e.g., incorrect cross-connections in the network), they are generally expected to be of relatively long duration, and it is not considered essential that they be detected and terminated as quickly as other types of defects. Thus, T1.231.04 allows up to 30 seconds for the detection or termination of a TIM defect. This is reflected in the defect detection and termination time requirements in this document (e.g., **R6-167 [1071]** and **R6-172 [1075]**); however, significantly faster defect detection and termination times are both possible and desirable (see, for example, **O6-168 [1074]** and **O6-173 [1076]**).
- For the purpose of these criteria, the "TIM defect detection time" is the time that it would take the PTE to detect a TIM defect if the incoming path trace were to change to a new (non-matching) value. For example, this could be the path trace string sampling rate times the number of non-matching samples necessary to cause the detection of a TIM defect.
- In general, an NE's TIM defect detection algorithm needs to be able to tolerate both random bit errors and changes in the phase of the incoming path trace string. Such changes could be caused, for example, by protection switches at upstream NEs that result in one or more J1 or J2 bytes being dropped or repeated (e.g., due to transmission delay differences between the long and short paths around a UPSR).

8. According to the specifications in ANSI T1.105 and T1.231.04, *SONET Layer 1 In-Service Digital Transmission Performance Monitoring*, PTE is supposed to be able to be provisioned both to detect/not detect TIM defects, and to insert/not insert AIS downstream upon TIM defect detection. Conversely (and as previously discussed in detail in GR-253-ILR Issue ID 253-139), Telcordia believes that while monitoring for TIM defects needs to be a user-provisionable feature (e.g., see **R6-166 [1070v2]**), if that feature has been activated for a particular path then the detection of a TIM defect on that path should always cause AIS to be inserted downstream. Therefore, this document does not include any criteria that indicate that an NE should support a capability that would allow the user to turn off automatic AIS insertion. Such a capability could be supported by an NE (e.g., to meet a standard), but Telcordia recommends that it not be used.

- Given that the maximum TIM defect termination time of 30 seconds is much longer than the typical failure-declaration soaking time of approximately 2.5 seconds, there do not appear to be any situations in which an NE would be required to detect and terminate a TIM defect without also declaring a TIM failure. Similarly, since the maximum TIM defect detection time of 30 seconds is much longer than the typical failure-clearing soaking time of approximately 10 seconds, there do not appear to be any situations in which an NE would be required to terminate and (re)detect a TIM defect without also clearing a TIM failure. Therefore, it might not appear to be necessary for the criteria related to TIM failure declaration and clearing to specify the use of soaking intervals. However, for consistency with the requirements for other types of failures, the 2.5-second and 10-second soaking intervals are included in both the TIM failure declaration and clearing criteria in this document (e.g., see **R6-174 [1077]** and **R6-175 [1078]**) and in the corresponding ANSI T1.231.04 specifications. In addition, those soaking intervals could be of importance for NEs that detect and terminate TIM defects substantially faster than the maximum allowed detection and termination times of 30 seconds.
- If the PTE has detected an AIS or LOP defect on its incoming signal, the J1 or J2 byte cannot be accessed to monitor for TIM defects, and therefore a TIM defect cannot be either detected or terminated. Otherwise, during a 30-second interval the PTE can expect to receive 3750 copies of the incoming path trace string. To reduce the processing burden on the NE, the TIM defect detection criteria (e.g., **R6-167 [1071]**) indicate via the use of the word “sampled” that not all of these copies need to be compared to the provisioned expected value. While no specific value is provided for the number of consecutive non-matching samples that should trigger the detection of a TIM defect, that number needs to be large enough to avoid “false” TIM defect detection (i.e., the detection of TIM defects that are caused by random bit errors rather than an actual end-to-end connectivity problem).
- Unlike most other cases involving failures based on directly detected defects, it is very likely that the conditions for declaring both a TIM failure and either an UNEQ or PLM failure will be met in response to a single root-cause incoming signal problem. That is, an incorrect cross-connection in the network is likely to result in the PTE receiving either an unequipped path or a path signal containing the wrong payload mapping, and also a path trace that does not match the provisioned expected value. The TIM defect has not been defined such that its detection automatically causes the termination of either an UNEQ or PLM defect (or vice versa), and therefore the NE’s conformance to **R6-312 [626v3]** is particularly important to avoid the generation of extraneous messages to the OS. However, if the NE does not meet the objectives on detecting TIM defects within 2 seconds (i.e., **O6-168 [1074]** and **O6-182 [1204]**) then it is also very likely that the TIM defect will not yet have been detected at the time when the condition for declaring the UNEQ or PLM failure is met. For an UNEQ failure this is not an issue since that failure is considered higher priority than a TIM failure (see [Table 6-6](#)). That is, when the TIM failure is subsequently declared it would simply not be autonomously reported to the OS. On the other hand, it is an issue for a PLM failure since that failure is considered lower priority than a TIM failure. In that case the NE may or may not clear the PLM failure when the TIM failure is subsequently declared.

In addition, to this common information, [Section 6.2.1.1.9.A](#) contains (primarily) STS-specific information and criteria, while [Section 6.2.1.1.9.B](#) contains (primarily) VT-specific information and criteria.

6.2.1.1.9.A STS Path Trace Identifier Mismatch

R6-161 [725v2] A SONET NE that contains STS PTE shall allow the user to provision (on a per-STs-path basis) the contents of the STS path trace carried in the J1 byte of the STS POH originated by the PTE. The transmitted STS path trace string shall be 64 bytes in length.

Prior to the release of Issue 3 of this document, it was generally assumed that the user would provision STS PTE in a SONET NE to transmit path trace strings that consisted of ASCII printable characters [i.e., ‘20’ through ‘7E’ (hex)], and the criteria that appear in this section are primarily based on that assumption. However ANSI T1.269 describes a 64-byte trace message format that includes non-ASCII printable characters, and several of the criteria in this section refer to that standard.

It should also be noted that in the ITU-T SDH standards the J1 byte may carry a 16-byte identifier rather than a 64-byte string. The criteria in this section are not intended to preclude a SONET NE from supporting the capability to transmit or detect STS path trace strings that meet the SDH 16-byte format (e.g., for use at international boundaries).

R6-162 [997v3] A SONET NE that contains STS PTE shall support a feature that allows the contents of the STS path traces to be provisioned as a single string of ASCII printable characters. In addition, the following apply:

- The feature shall allow the user to enter a string of up to 62 characters
- The feature shall place no restriction on the content of the string, except that the characters shall be ASCII printable characters
- The NE shall automatically pad the string entered by the user to 62 characters using ASCII NULL characters, and then add <CR> and <LF> characters (i.e., ‘0D’ and ‘0A’) for a total of 64 characters.
- Each 8-bit ASCII character shall be loaded into one J1 byte.

CR6-163 [1196] A SONET NE that contains STS PTE may be required to support a feature that allows the contents of an STS path trace to be provisioned such that it meets the specifications in ANSI T1.269 related to 64-byte trace messages.

R6-164 [1069v2] A SONET NE shall support a feature to allow the user to provision the “expected” ASCII-based path trace for each STS path that it terminates and for which TIM-P detection has been or can be activated (see below). In addition, the following apply:

- The feature shall allow the user to enter a single string of up to 62 characters
- The feature shall place no restriction on the contents of the string, except that the characters shall be ASCII printable characters.

CR6-165 [1197] A SONET NE may be required to support a feature that allows the user to provision an “expected” path trace that meets the specifications in ANSI T1.269 related to 64-byte trace messages for each STS path that it terminates and for which TIM-P detection has been or can be activated.

The capability to provision an expected STS path trace may also need to be provided for STS path trace diagnostics purposes as described in [Section 6.2.3.2.3.A](#) (see **CR6-443 [999v2]**). Therefore, it is likely that an NE will need to provide the capability to activate and deactivate TIM-P detection (see below) separately from the provisioning of the expected trace. In addition, this capability (i.e., separate expected STS path trace and TIM-P activation provisioning) can help reduce the chance that traffic will be interrupted as a result of provisioning errors.

R6-166 [1070v2] A SONET NE that contains STS PTE shall support a feature that, when activated, monitors the received STS path trace string for TIM-P detection purposes. That feature shall be provisionable on a per-STs-path basis and shall have a default of “not active”. In addition, the feature shall not be automatically activated by the act of the user entering an expected STS path trace.

Note that **R6-166 [1070v2]** is not meant to imply that an NE cannot be designed to allow simultaneous provisioning of an expected STS path trace and activation of the TIM-P detection feature. For example, an NE could meet that requirement by supporting a single command with which the user would enter an expected path trace string and indicate whether that string is to be used for TIM-P detection or diagnostics purposes (with the default being that it be used for diagnostic purposes).

R6-167 [1071] STS PTE shall detect a TIM-P defect within 30 seconds (or less) when none of the sampled 64-byte STS path trace strings match the provisioned expected value.

O6-168 [1074] A SONET NE’s STS path trace sampling rate and TIM-P defect detection algorithm should be such that a TIM-P defect is detected within 2 seconds after the NE’s STS PTE stops receiving the provisioned expected path trace string.

R6-169 [1072] A SONET NE’s TIM-P defect detection algorithm shall be such that, given an incoming signal with a BER of 10^{-3} or less, the probability that the STS PTE will detect a false TIM-P defect during the TIM-P defect detection time is less than 1×10^{-n} .

The current value for “n” in **R6-169 [1072]** is ‘6’; however, this value is under study and may need to be changed in the future.

R6-170 [1073] A change in the phase of an incoming STS path trace string shall cause the STS PTE to consider, at most, one sample to be mismatched for the purpose of detecting and terminating TIM-P defects.

O6-171 [1001v3] If a SONET NE is comparing an incoming STS path trace string to an expected path trace (for either TIM-P defect detection/termination or diagnostics purposes) and the expected path trace consists of a single string of ASCII printable

characters, then the comparison should ignore any trailing ASCII NULL, <CR> and <LF> characters contained in the incoming path trace. Similarly, if the expected path trace meets the specifications in ANSI T1.269 for a 64-byte message that contains a frame start marker and CRC-7 in the first byte, then the comparison should ignore whether the pads for any of the data elements that require padding (to bring them to their specified lengths) consist of ASCII NULL or SPACE characters.

An NE that meets **O6-171 [1001v3]** could accept (and consider matched), for example, an incoming path trace generated by STS PTE that meets **R6-162 [997v3]** (i.e., that pads with NULLs to 62 characters and then adds <CR> and <LF>), or by older STS PTE that might not meet **R6-162 [997v3]** (e.g., that pads with NULLs to 64 characters, or that adds the <CR> and <LF> characters before padding with NULLs).

- R6-172 [1075]** STS PTE shall terminate a TIM-P defect within 30 seconds (or less) when four-fifths (or more) of the sampled STS path trace strings match the provisioned expected value.

Note that in order to meet **R6-172 [1075]** it would be sufficient for STS PTE to sample the incoming STS path trace string as infrequently as once every six seconds (i.e., 5 samples in any 30-second window). However, such a sampling rate would not be sufficient for detecting TIM-P defects as it would result in an unacceptably high probability of false TIM-P defect detection.

- O6-173 [1076]** A SONET NE's STS path trace sampling rate and TIM-P defect termination algorithm should be such that a TIM-P defect is terminated within 2 seconds after the STS PTE begins receiving the provisioned expected path trace string.
- R6-174 [1077]** An NE shall declare a TIM-P failure if a TIM-P defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set a TIM-P failure indication and send an alarm message to an OS unless the condition in **R6-312 [626v3]** applies.
- R6-175 [1078]** An NE shall clear the TIM-P failure if the TIM-P defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the TIM-P failure indication and send a clear message to an OS (if the failure was reported to the OS).

6.2.1.1.9.B VT Path Trace Identifier Mismatch

- CR6-176 [1198]** A SONET NE that contains VT PTE may be required to support a feature that allows the contents of the VT path traces to be provisioned (on a per-VT-path basis) by the user.
- R6-177 [1199]** If a SONET NE supports a feature that allows the contents of the VT path traces to be provisioned, then that feature shall allow the user to enter the contents of the path trace such that it meets the specifications in ANSI T1.269 related to 16-byte trace messages.

R6-178 [1200] A SONET NE shall support a feature to allow the user to provision the expected path trace for each VT path that it terminates and for which TIM-V detection has been or can be activated (see below). In addition, that feature shall allow the user to enter the contents of the path trace such that it meets the specifications in ANSI T1.269 related to 16-byte trace messages.

The capability to provision an expected VT path trace may also need to be provided for VT path trace diagnostics purposes as described in [Section 6.2.3.2.3.A](#) (see **CR6-446 [1226]**). Therefore, it is likely that an NE will need to provide the capability to activate and deactivate TIM-V detection (see below) separately from the provisioning of the expected trace. In addition, this (separate expected VT path trace and TIM-V activation provisioning) can help reduce the chance that traffic will be interrupted as a result of provisioning errors.

CR6-179 [1201] A SONET NE that contains VT PTE may be required to support a feature that, when activated, monitors the received VT path trace string for TIM-V detection purposes.

R6-180 [1202] If a SONET NE supports a TIM-V detection feature, that feature shall be provisionable on a per-VT-path basis and shall have a default of “not active”. In addition, the feature shall not be automatically activated by the act of the user entering an expected VT path trace.

Note that **R6-180 [1202]** is not meant to imply that an NE cannot be designed to allow simultaneous provisioning of an expected VT path trace and activation of the TIM-V detection feature. For example, an NE could meet that requirement by supporting a single command with which the user would enter an expected path trace string and indicate whether that string is to be used for TIM-V detection or diagnostics purposes (with the default being that it be used for diagnostic purposes).

R6-181 [1203] VT PTE shall detect a TIM-V defect within 30 seconds (or less) when none of the sampled 16-byte VT path trace strings match the provisioned expected value.

O6-182 [1204] A SONET NE’s VT path trace sampling rate and TIM-V defect detection algorithm should be such that a TIM-V defect is detected within 2 seconds after the NE’s VT PTE stops receiving the provisioned expected path trace string.

R6-183 [1205] A SONET NE’s TIM-V defect detection algorithm shall be such that, given an incoming signal with a BER of 10^{-3} or less, the probability that the VT PTE will detect a false TIM-V defect during the TIM-V defect detection time is less than 1×10^{-n} .

The current value for “n” in **R6-183 [1205]** is 6; however, this value is under study and may need to be changed in the future.

R6-184 [1206] A change in the phase of an incoming VT path trace string shall cause the VT PTE to consider, at most, one sample to be mismatched for the purpose of detecting and terminating TIM-V defects.

O6-185 [1207] If a SONET NE is comparing an incoming VT path trace string to an expected path trace that meets the specifications in ANSI T1.269 for a 16-byte message, then the comparison should ignore whether the pads for any of the data elements that require padding (to bring them to their specified lengths) consist of ASCII NULL or SPACE characters.

R6-186 [1208] VT PTE shall terminate a TIM-V defect within 30 seconds (or less) when four-fifths (or more) of the sampled VT path trace strings match the provisioned expected value.

Note that in order to meet **R6-186 [1208]** it would be sufficient for VT PTE to sample the incoming VT path trace string as infrequently as once every six seconds (i.e., 5 samples in any 30-second window). However, such a sampling rate would not be sufficient for detecting TIM-V defects as it would result in an unacceptably high probability of false TIM-V defect detection.

O6-187 [1209] A SONET NE's VT path trace sampling rate and TIM-V defect termination algorithm should be such that a TIM-V defect is terminated within 2 seconds after the VT PTE begins receiving the provisioned expected path trace string.

R6-188 [1210] An NE shall declare a TIM-V failure if a TIM-V defect persists for 2.5 (± 0.5) seconds. Upon declaring the failure, it shall set a TIM-V failure indication and send an alarm message to an OS unless the condition in **R6-312 [626v3]** applies.

R6-189 [1211] An NE shall clear the TIM-V failure if the TIM-V defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the TIM-V failure indication and send a clear message to an OS (if the failure was reported to the OS).

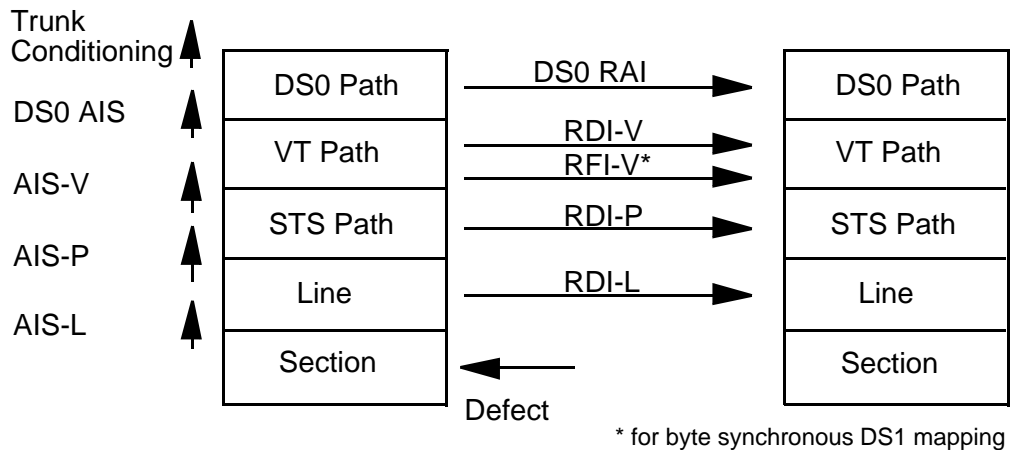
6.2.1.2 Alarm Indication Signal

An AIS is a maintenance signal used in the digital network to alert downstream equipment that a defect or equipment failure has been detected. The SONET signal format provides AISs for the Line (AIS-L), STS path (AIS-P) and VT path (AIS-V) layers. In addition, SONET NEs may also deal with DS_n (e.g., DS0, DS1, DS3) AIS and trunk conditioning.

In most cases, equipment generates the AIS defined for the next higher layer when it detects a defect on the signal that it is terminating (e.g., STE generates AIS-L when it detects an LOS or LOF defect on its incoming signal), as depicted in [Figure 6-2](#). [Figure 6-2](#) also illustrates the upstream maintenance signals discussed in [Section 6.2.1.3](#), and shows that AIS is sent downstream in both the physical sense (i.e., as an externally discernible signal) and the logical sense (i.e., as an internally generated signal communicated between logical SONET layers). Thus, a defect detected at a given layer functionally causes an internal AIS to be generated to all applicable higher layers (e.g., the Section layer generates AIS-L).

In addition to being generated by equipment that detects a defect on the signal that it is terminating, in some cases AIS is generated by equipment that originates a signal. This occurs when equipment that is originating a signal detects that the circuitry supporting provisioned higher layer origination functions has failed or been removed without having been “unprovisioned”, and it continues until standby circuitry (if available) is switched in or the failure is cleared. For example, LTE generates AIS-P (for the affected STS paths) when it detects that STS PTE supporting provisioned STS path origination functions has failed.

Figure 6-2 Maintenance Signals for SONET Layers



6.2.1.2.1 Line AIS

STE (and physical layer regenerators, see TR-NWT-000917) send AIS-L to alert the downstream LTE that a defect has been detected on the incoming SONET signal, or that LTE supporting provisioned line origination functions has failed. If line-level APS is provided as a feature, the downstream LTE can initiate protection switching upon detection of the AIS-L defect, as specified in [Section 5.3.3.1](#) for linear APS.

R6-190 [512v2] STE shall generate AIS-L downstream within 125 μs of detecting an LOS or LOF defect on the incoming signal or the failure of LTE supporting provisioned line origination functions. The AIS-L shall be generated as an OC-N or STS-N electrical signal that contains valid SOH and a scrambled all-ones pattern for the remainder of the signal.

Figures 6-3 and 6-13 (in [Section 6.2.1.7](#)) illustrate when STE generates AIS-L in response to a defect detected on an incoming signal or an equipment failure. Note that the AIS-L generated as described above automatically provides convenient generation of AIS for higher layers (e.g., STS and VT path AIS). Also note that although ANSI T1.105 defines an alternate format for an AIS-L generated by STE that does not support the full signal scrambling function defined in [Section 5.1.3](#), Telcordia does not recommend the use of that format.

- R6-191 [513v2]** STE shall deactivate AIS-L within 125 μ s of terminating the defect that caused it to be sent, or in the case of a local equipment failure, within 125 μ s of clearing the failure or determining that standby equipment has been switched in.
- R6-192 [514]** LTE shall detect an AIS-L defect on the incoming signal when bits 6, 7 and 8 of the K2 byte contain the '111' pattern in five consecutive frames.
- R6-193 [515]** LTE shall terminate the AIS-L defect on the incoming signal when bits 6, 7 and 8 of the K2 byte have any pattern other than '111' in five consecutive frames.

For testing conformance to **R6-193 [515]**, it is sufficient to use a consistent non-'111' pattern.

- R6-194 [516]** An NE shall declare an AIS-L failure if an AIS-L defect persists for 2.5 (± 0.5) seconds. Upon declaring an AIS-L failure, the NE shall set an AIS-L failure indication for the line and (if AIS-L is a Reported failure for the line) send a message to an OS unless the condition in **R6-312 [626v3]** applies.
- R6-195 [517]** For the purposes of trunk conditioning, SONET NEs that contain DS0 PTE or VT PTE that supports the rearrangement of the DS0 channels in byte-synchronously mapped DS1s shall use the integration technique described in **R6-59 [419]** to declare AIS-L failures. Upon declaring an AIS-L failure, the NE shall perform the actions listed in **R6-194 [516]**.
- R6-196 [518]** An NE shall clear the AIS-L failure when the AIS-L defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the AIS-L failure indication and send a clear message to an OS (if the failure was reported to the OS).

6.2.1.2.2 STS Path AIS

LTE sends AIS-P to alert the downstream STS PTE that it has detected a defect on its incoming line signal, or that STS PTE supporting provisioned path origination functions has failed. Figures 6-4 and 6-13 (in Section 6.2.1.7) illustrate when AIS-P is generated in response to a defect detected on an incoming signal or an equipment failure.

- R6-197 [519v3]** LTE shall generate AIS-P downstream for the affected STS paths within 125 μ s of detecting an AIS-L defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2) or (if the STS pointer is processed) an LOP-P defect on the incoming signal, or the failure of STS PTE supporting provisioned path origination functions. The AIS-P shall be generated as all-ones in the H1, H2 and H3 bytes, and the entire STS SPE.

Note that for an STS-Mc path, all-ones is generated in all M of the H1, H2 and H3 bytes. Also note that in the functional model used in this document, incoming defects that are detected at STE (i.e., LOS, LOF) cause AIS-L to be generated downstream. This AIS-L is then detected by the LTE and causes it to generate AIS-P.

R6-198 [521v2] LTE shall deactivate the AIS-P within 125 μ s of terminating the defect that caused it to be sent, or in the case of a local equipment failure, within 125 μ s of clearing the failure or determining that standby equipment has been switched in. LTE that performs STS pointer processing shall deactivate AIS-P by constructing a correct STS pointer with a set NDF, followed by normal pointer operations, as well as ceasing to insert the all-ones pattern in the STS SPE. LTE that does not perform STS pointer processing shall deactivate AIS-P by ceasing the insertion of all-ones in the H1, H2 and H3 bytes, and the STS SPE.

R6-199 [523] STS PTE shall detect an AIS-P defect when the H1 and H2 bytes for an STS path contain an all-ones pattern in three consecutive frames. For an STS-Nc path, only the H1 and H2 bytes of the first STS-1 need to be observed.

Note that in the functional model used in this document, incoming AIS-P defects are detected by STS PTE, not by LTE. If LTE that processes STS pointers receives all-ones H1 and H2 bytes, it is required to relay those all-ones values downstream (see [Section 3.5.1.5](#)). No “all-ones pointer relay” defect has been defined (and therefore an all-ones pointer relay failure is not declared if the LTE continues to relay the all-ones pointer); however, the act of performing all-ones pointer relay does have several effects on the operation of LTE. For example, it affects the LTE’s detection of LOP-P defects (see [Section 6.2.1.1.3](#)). In addition it should be noted that since LTE cannot perform increment or decrement operations while it is performing all-ones pointer relay, the incoming STS SPE associated with the all-ones pointers cannot be expected to be passed through unaltered. As a minimum, it is likely that the SPE sent downstream will be affected by buffer overflows or underflows that occur at the LTE. Alternatively, some LTE may discard the incoming SPE and generate a new all-ones SPE (synchronized to the local clock) for transmission downstream.

R6-200 [524] STS PTE shall terminate an AIS-P defect when the H1 and H2 bytes for the STS path contain a valid STS pointer with a set NDF, or when they contain valid, identical STS pointers with normal NDFs for three consecutive frames. For an STS-Nc path, the concatenation indicators must also be valid.

R6-201 [1079v2] STS PTE shall terminate an AIS-P defect when it detects an LOP-P defect.

Note that prior to Issue 2, Revision 2 of this document, STS PTE was required to terminate an AIS-P defect only in the cases shown in **R6-200 [524]**, and to *not* detect an LOP-P defect if an AIS-P defect had already been detected. Therefore, NEs designed to earlier issues of this document are not expected to conform to **R6-201 [1079v2]**. Also note that the current requirements are consistent with the specifications that appear in ITU-T Rec. G.783 and ANSI T1.231.04, while the previous criteria were consistent with those that appeared in the 1997 version of ANSI T1.231.

R6-202 [525] An NE shall declare an AIS-P failure if an AIS-P defect persists for 2.5 (± 0.5) seconds. Upon declaring an AIS-P failure, the NE shall set an AIS-P failure indication for that path and (if AIS-P is a Reported failure for the path) send a message to an OS unless the condition in **R6-312 [626v3]** applies.

R6-203 [526] For the purposes of trunk conditioning, SONET NEs that contain DS0 PTE or VT PTE that supports the rearrangement of the DS0 channels in byte-synchronously mapped DS1s shall use the integration technique described in **R6-59 [419]** to declare AIS-P failures. Upon declaring an AIS-P failure, the NE shall perform the actions listed in **R6-202 [525]**.

R6-204 [527] An NE shall clear an AIS-P failure when the AIS-P defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the AIS-P failure indication and send a clear message to an OS (if the failure was reported to the OS).

6.2.1.2.3 VT Path AIS

STS PTE sends AIS-V to alert the downstream VT PTE that it has detected a defect on its incoming path signal, or that VT PTE supporting provisioned path origination functions has failed. In addition, VT PTE that byte-synchronously maps an incoming DS1 into a VT1.5 sends AIS-V when certain DS1 defects are detected. Figures 6-6, 6-9 and 6-13 (in Section 6.2.1.7) illustrate when AIS-V is generated in response to a defect detected on an incoming signal or an equipment failure.

R6-205 [528v5] STS PTE shall generate AIS-V downstream for the affected VT paths within 500 μ s of detecting an AIS-P defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2), an LOP-P defect, an UNEQ-P defect, a TIM-P defect (if activated), a PLM-P defect, or (if the VT pointer is processed) an LOP-V defect on the incoming signal, or the failure of VT PTE supporting provisioned path origination functions. The AIS-V shall be generated as an all-ones code in the entire VT, including the V1 through V4 bytes.

Note that in the functional model used in this document, incoming defects that are detected at STE or LTE (i.e., LOS, LOF, AIS-L) cause AIS-P to be generated downstream. This AIS-P is then detected by the STS PTE and causes it to generate AIS-V.

R6-206 [529] VT PTE shall generate AIS-V within 500 μ s of detecting a DS1 LOS, OOF or AIS defect on an incoming DS1 that it byte-synchronously maps into a single VT1.5.

R6-207 [531v2] STS PTE shall deactivate AIS-V within 500 μ s of terminating the defect that caused it to be sent, or in the case of a local equipment failure, within 500 μ s of clearing the failure or determining that standby equipment has been switched in. STS PTE that performs VT pointer processing shall deactivate AIS-V by constructing a correct VT pointer with valid VT size and a set NDF, followed by normal pointer operations, as well as ceasing to insert the all-ones pattern in the rest of the VT. STS PTE that does not perform VT pointer processing shall deactivate AIS-V by ceasing the insertion of the all-ones pattern in the entire VT.

R6-208 [533] VT PTE shall deactivate AIS-V (as described in **R6-207 [531v2]**) within 500 μ s of terminating the defect on the incoming DS1 that caused it to be generated.

- R6-209 [534]** VT PTE shall detect an AIS-V defect for the VT path upon receiving an all-ones pattern in the V1 and V2 bytes in three consecutive VT superframes.
-

Note that in the functional model used in this document, incoming AIS-V defects are detected by VT PTE, not by STS PTE. If STS PTE that processes VT pointers receives all-ones V1 and V2 bytes, it is required to relay those all-ones values downstream (see [Section 3.5.2.5](#)). No “all-ones pointer relay” defect has been defined (and therefore an all-ones pointer relay failure is not declared if the STS PTE continues to relay the all-ones pointer); however, the act of performing all-ones pointer relay does have several effects on the operation of STS PTE (see [Section 6.2.1.1.3](#)). In addition it should be noted that since STS PTE cannot perform increment or decrement operations while it is performing all-ones pointer relay, the incoming VT SPE associated with the all-ones pointers cannot be expected to be passed through unaltered. As a minimum, it is likely that the SPE sent downstream will be affected by buffer overflows or underflows that occur at the STS PTE. Alternatively, some STS PTE may discard the incoming SPE and generate a new all-ones SPE (synchronized to the local clock) for transmission downstream.

- R6-210 [535]** VT PTE shall terminate an AIS-V defect upon receiving a valid VT pointer with valid VT size and a set NDF, or upon receiving three consecutive VT superframes containing valid, identical VT pointers with a valid VT size and normal NDFs.

- R6-211 [1080v2]** VT PTE shall terminate an AIS-V defect when it detects an LOP-V defect.
-

Note that prior to Issue 2, Revision 2 of this document, VT PTE was required to terminate an AIS-V defect only in the cases shown in **R6-210 [535]**, and to *not* detect an LOP-V defect if an AIS-V defect had already been detected. Therefore, NEs designed to earlier issues of this document are not expected to conform to **R6-211 [1080v2]**. Also note that the current requirements are consistent with the specifications that appear in ITU-T Rec. G.783 and ANSI T1.231.04, while the previous criteria were consistent with those that appeared in the 1997 version of ANSI T1.231.

- R6-212 [536]** An NE shall declare an AIS-V failure if an AIS-V defect persists for 2.5 (± 0.5) seconds. Upon declaring the AIS-V failure, the NE shall set an AIS-V failure indication for the path and (if AIS-V is a Reported failure for the path) send a message to an OS unless the condition in **R6-312 [626v3]** applies.

- R6-213 [537]** For the purposes of trunk conditioning, SONET NEs that contain DS0 PTE or VT PTE that supports the rearrangement of the DS0 channels in byte-synchronously mapped DS1s shall use the integration technique described in **R6-59 [419]** to declare AIS-V failures. Upon declaring an AIS-V failure, the NE shall perform the actions listed in **R6-212 [536]**.

- R6-214 [538]** An NE shall clear an AIS-V failure when the AIS-V defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the AIS-V failure indication and send a clear message to an OS (if the failure was reported to the OS).
-

6.2.1.2.4 DS_n AIS

In many applications, SONET is the transport medium for lower-speed digital signals such as DS1s and DS3s, and in those applications the SONET NEs may need to generate or detect DS_n AIS, including DS1 AIS-Customer Installation (DS1 AIS-CI). Figures 6-5 through 6-11 (in Section 6.2.1.7) illustrate the use of DS_n AIS for STS and VT PTE terminating SPEs of various compositions. Criteria on the construction of DS1, DS1C, DS2 and DS3 AIS are contained in GR-499-CORE, along with a conditional requirement to support the generation and detection of DS1 AIS-CI and a reference to ANSI T1.403, *Network and Customer Installation Interfaces – DS1 Electrical Interface*, for additional specifications related to that type of signal.

DS0 AIS (also sometimes referred to as “UNICODE”) is defined in SONET as a specific pattern in the ABCD signaling bits (i.e., A = 0, B = 0, C = 1, D = 0). Note that DS0 AIS may not be applicable outside of SONET networks, so an interface to a non-SONET NE may require the application of service-specific trunk conditioning codes (i.e., one or more provisioned ABCD signaling codes and an 8-bit insertion word). Also note that DS0 AIS is defined to include an all-ones pattern in the DS0 payload (in addition to the ‘0010’ code in the signaling bits) in GR-303-CORE. For defects detected on an incoming SONET signal, the generation of a lower-layer AIS (e.g., AIS-P, AIS-V) could be sufficient to ensure all-ones DS0 payloads. However, for defects detected on an incoming DS1, it may be necessary for the NE to actively insert all-ones in the affected DS0s.

Criteria on generating AIS at other transport signal interfaces is FFS.

-
- R6-215 [539v2]** STS or VT PTE shall generate DS1, DS1C, DS2 or DS3 AIS downstream within 125 μs of the detection of certain defects, as shown in Figures 6-5 through 6-9.
- R6-216 [951]** VT PTE with DS0 rearrangement capabilities shall generate DS0 AIS downstream within 3 ms of the detection of certain defects (unless it is provisioned to apply a service-specific trunk conditioning code, see Section 6.2.1.6), as shown in Figures 6-10 and 6-11.
-

Since DS0 AIS is carried in the ABCD signaling bits, which have a cycle time of 3 ms, VT PTE could meet the preceding requirement by (for example) sending the DS0 AIS code in the next complete set of ABCD signaling bits after the defect is detected.

-
- R6-217 [540]** STS or VT PTE shall remove a downstream DS1, DS1C, DS2 or DS3 AIS within 125 μs of terminating the defect that caused it to be sent.
- R6-218 [541v2]** If a defect that causes VT PTE to insert DS0 AIS is terminated before a failure is declared, then the VT PTE shall remove the DS0 AIS within 3 ms of terminating that defect. If a failure was declared, then the VT PTE shall remove the DS0 AIS within 3 ms of clearing the failure.
-

In addition to generating and removing DS_n AIS, in some applications a SONET NE may also need to be able to (non-intrusively) monitor for DS_n AIS even if it does not terminate the DS_n path. For example, STS PTE that supports the generation of PDI-P signals and uses the asynchronous DS3 mapping would need to be capable of

detecting DS3 AIS on the incoming DS3 signal in order to insert the appropriate code in the STS signal label. Similarly, an NE that performs intermediate DS_n path PM would need to monitor for DS_n AIS to correctly accumulate certain parameters.

R6-219 [542] A SONET NE shall monitor for DS_n AIS and declare DS_n AIS failures on DS_n paths that it terminates.

Requirements on DS_n AIS defect detection and failure declaration for DS1, DS1C, DS2 and DS3 signals are contained in GR-499-CORE. For DS0 AIS, see **R6-223 [547]** through **R6-226 [955]**.

CR6-220 [544v2] A SONET NE with DS_n interfaces may be required to monitor for DS_n AIS (as if it were terminating the DS_n path) on incoming DS_n signals where the DS_n path is not terminated.

CR6-221 [952] A SONET NE with DS_n interfaces may be required to monitor for DS_n AIS (as if it were terminating the DS_n path) on DS_n signals that are asynchronously demultiplexed from STS or VT SPEs (i.e., on outgoing DS_n signals).

See [Section 6.2.1.8.7](#) for criteria on the declaration of DS_n AIS failures when the DS_n path is not terminated.

O6-222 [953] If the DS_n path is not terminated, then the capability to monitor for DS_n AIS should be provided independently of any options to provide non-clear-channel transport of DS_n signals using the asynchronous DS_n mappings (see [Section 6.2.1.1.2](#)).

An NE that meets **O6-222 [953]** would be capable of simultaneously providing (for example) both clear-channel DS3 transport and PDI-P generation for the STS-1 SPE containing that DS3. Note that in order to detect DS_n AIS, an NE may need to monitor the framing bits in the DS_n (e.g., to detect DS1 AIS an NE must determine that valid DS1 framing is not present, and to detect DS3 AIS it must determine that valid DS3 framing is present). However, an NE should be able to perform that monitoring without also inserting DS_n AIS downstream when a DS_n OOF is detected.

R6-223 [547] A SONET NE that monitors for DS0 AIS shall detect a DS0 AIS defect if it receives two consecutive sets of ABCD signaling bits set to the code '0010' (i.e., the '0010' code persisting for 6 ms).

R6-224 [548] A SONET NE shall terminate a DS0 AIS defect if it receives four consecutive sets of ABCD signaling bits set to any code other than the '0010' code (i.e., any code other than '0010' persisting for 12 ms).

R6-225 [954] A SONET NE that is provisioned to insert a service-specific trunk conditioning code on a particular DS0 path shall declare a DS0 AIS failure when a DS0 AIS defect persists for 3.25 (±0.25) seconds.

- R6-226 [955]** A SONET NE that is provisioned to insert a service-specific trunk conditioning code on a particular DS0 path shall clear a DS0 AIS failure when the DS0 AIS defect is terminated.
-

See [Section 6.2.1.6](#) for criteria related to the actions that an NE provisioned to insert a service-specific trunk conditioning code takes when it declares or clears a DS0 AIS failure. Criteria on the declaration and clearing of DS0 AIS failures by other NEs that detect DS0 AIS are provided in the GRs or TRs specific to those types of NEs (e.g., GR-303-CORE for IDLC NEs).

6.2.1.3 Remote Defect Indication and Remote Failure Indication

Remote Alarm Indication (RAI) signals have historically been used in the digital network to alert upstream terminals of a downstream failure so that they can initiate trunk conditioning on the failed circuit. The layered maintenance strategy incorporated in the SONET signal format (as shown in [Figure 6-2](#)) expands the applications of these signals in the network. In SONET, RDI signals occur at the Line, STS path and VT path layers, and a persistent RDI defect results in the derivation of an RFI failure. In addition, for byte-synchronously mapped DS1s, a VT path layer RFI signal is also defined, both to provide backward compatibility with equipment using the traditional RAI signal timing, and for translation to and from DS1 RAI signals.

6.2.1.3.1 Line Remote Defect Indication (RDI-L) and Remote Failure Indication (RFI-L)

The RDI-L signal (formerly called Line FERF) indicates to LTE that its peer LTE has detected an AIS-L (or lower-layer) defect on the signal that it (the first LTE) originated. An incoming RDI-L defect is used to derive an RFI-L failure.

-
- R6-227 [549v2]** LTE shall generate RDI-L within 125 μ s of detecting an AIS-L defect (or a lower-layer, traffic-related, near-end defect, see [Section 6.2.1.8.2](#) and [Figures 6-4 through 6-12](#)). The LTE shall generate RDI-L by inserting the code '110' in bits 6, 7 and 8 of the K2 byte.
- R6-228 [550v2]** If bits 6 through 8 of the K2 byte are not used for other purposes (e.g., the linear APS mode indication), the LTE shall deactivate RDI-L by inserting the code '000' in bits 6, 7 and 8 of the K2 byte within 125 μ s of terminating the defect that caused it to be sent (assuming it has been sent for any minimum RDI-L assertion time supported by the NE, see below).
- R6-229 [551v2]** If bits 6 through 8 of the K2 byte are used for other purposes, LTE shall deactivate RDI-L by inserting an "appropriate code" (see below) in those bits within 125 μ s of terminating the defect that caused it to be sent (assuming it has been sent for the minimum assertion time).
-

For LTE that uses bits 6 through 8 of the K2 byte to indicate its linear APS mode, the “appropriate code” referred to in **R6-229 [551v2]** can be either the mode indication, or the code ‘000’ for 50 to 200 ms, followed by the mode indication. In some (but not all) situations, inserting the code ‘000’ for 50 to 200 ms will allow new LTE to interoperate with older LTE that was built to early issues of the Telcordia SONET criteria documents. (In those documents, an RDI-L defect was required to be terminated only if the ‘000’ code was received.)

Note that the RDI-L generation and deactivation times in **R6-227 [549v2]**, **R6-228 [550v2]** and **R6-229 [551v2]** are more stringent than the RDI-L specifications that appear in ANSI T1.105 (which states only that RDI-L must be generated or removed within 100 ms of detecting or terminating an incoming defect). However, it is not expected that these requirements will be changed. Three reasons for this are:

- The 125 μ s RDI-L generation and removal times in these requirements are consistent with the times that appeared in previous issues of this document
- Many existing implementations conform to these time requirements
- LTE that conforms to these requirements can also meet the standard (if it also conforms to **O6-230 [956]**).

With the definition and increasing support of far-end Line PM parameters whose accumulation is affected by the presence or absence of an RDI-L defect (see [Section 6.2.2.4.2](#)), it is important for LTE that generates an RDI-L signal to continue to generate it long enough to assure that its peer LTE will detect an RDI-L defect. ANSI T1.105 specifies that LTE must detect an RDI-L defect if the ‘110’ code is received in 10 consecutive frames, and therefore specifies an RDI-L minimum assertion time of 20 frames (i.e., the defect will still be detected even if transmission errors occur that affect the ‘110’ code in any one frame). To better align with T1.105 in this area, the following criteria were added (**O6-230 [956]**) and modified (**R6-231 [552v3]** and **R6-232 [553v3]**) in Issue 2 of this document.

O6-230 [956] When LTE generates RDI-L, it should generate it for at least 20 frames.

R6-231 [552v3] LTE shall detect an RDI-L defect when bits 6, 7 and 8 of the K2 byte contain the ‘110’ pattern in x consecutive frames, where x equals 5 or 10.

R6-232 [553v3] LTE shall terminate the RDI-L defect when bits 6, 7 and 8 of the K2 byte contain any pattern other than the code ‘110’ in x consecutive frames, where x equals 5 or 10.

To test conformance to **R6-231 [552v3]** and **R6-232 [553v3]**, it is sufficient to apply a consistent non-‘110’ code.

R6-233 [554] An NE shall declare an RFI-L failure when an RDI-L defect persists for 2.5 (± 0.5) seconds. Upon declaring an RFI-L failure, the NE shall set an RFI-L failure indication for the line and (if RFI-L is a Reported failure, see [Section 6.2.1.8](#)) send a message to an OS.

- R6-234 [555]** An NE shall clear the RFI-L failure when the RDI-L defect is absent for 10 (± 0.5) seconds. Upon clearing the failure, the NE shall clear the RFI-L failure indication and send a clear message to an OS (if the failure was reported to the OS).
-

6.2.1.3.2 STS Path Remote Defect Indication (RDI-P) and Remote Failure Indications (RFI-P)

An RDI-P signal indicates to STS PTE that its peer STS PTE has detected a defect on the signal that it (the first STS PTE) originated. Three bits are reserved for the RDI-P signal (as shown in [Table 6-4](#)), and detected RDI-P defects are used to derive RFI-P failures.

The details of the signal that STS PTE uses to notify upstream STS PTE that it has detected a defect underwent a number of changes as the SONET standards and Telcordia criteria evolved (i.e., from an “STS path yellow” signal using bit 5 of the G1 byte, to a one-bit RDI-P signal using bit 5 of G1, to the enhanced two-bit RDI-P signal using bits 5 and 6 of G1 that appeared in Issue 1 of this document). Currently an enhanced, three-bit RDI-P signal using bits 5 through 7 of G1 is defined in ANSI T1.105. The specifications in that standard are not expected to change, and therefore this section contains criteria that are relevant to that enhanced RDI-P signal. However, criteria for a one-bit RDI-P signal using bit 5 of G1 are included to cover products that have not implemented the enhanced version of RDI-P. Note that in this document, “RDI-P” is used generically to refer to both the one-bit and enhanced versions in text and criteria that are not specific to one particular version. In cases where it is necessary to distinguish between the two versions, the one-bit version is referred to as “one-bit RDI-P”, while the enhanced version is referred to as “ERDI-P”. Similarly, “RFI-P” is used to generically refer to a failure derived from a (generic) RDI-P defect, while “one-bit RFI-P” and “ERFI-P” are used in cases that are specific to a particular version.

Although ERDI-P was originally intended to be “backward compatible” with one-bit RDI-P, in the final version there are actually some significant differences. These differences are expected to primarily affect the accumulation of near-end and far-end STS path PM data. Therefore in applications where the STS path PM data is expected to be of importance, it is highly recommended that the STS PTE at both ends of the path (and any intermediate NEs provisioned to perform intermediate-path PM on that path) support the same version of RDI-P. In addition, it is recommended that the supported version be ERDI-P.

Finally, it should be noted that none of the criteria in this document are meant to imply that an NE that supports ERDI-P also needs to support one-bit RDI-P on a provisionable or, in the case of incoming signals, automatic-detection basis.⁹ Such capabilities may be provided by some NEs (either per-STS-path or NE-wide); however, it is also considered acceptable for an NE to support only one version of RDI-P.

9. For example, an NE that supports ERDI-P does not need to be able to be provisioned to transmit one-bit RDI-P, to provide different failure indications for incoming one-bit RDI-P and ERDI-P Server signals, or to change its near-end intermediate STS path PM parameter definitions based on the version of RDI-P being used by the end NEs (see [Sections 6.2.2.5.1](#) and [6.2.2.9](#)).

O6-235 [957v2] An NE should support ERDI-P generation and detection.

R6-236 [556v2] STS PTE shall generate an appropriate RDI-P signal, as specified in [Table 6-4](#), within 100 ms of detecting a listed defect. (Also see [Figures 6-5](#) through [6-12](#).)

O6-237 [557v3] With the possible exception noted following **O6-240 [959v2]**, STS PTE should generate an appropriate RDI-P signal as specified in [Table 6-4](#) within 125 μ s of detecting a listed defect.

Note that for the ERDI-P codes shown in [Table 6-4](#), bits 6 and 7 of the G1 byte are always set to opposite values. Conversely, in the one-bit RDI-P signal those two bits must be set to the same value (and should be set to '00', see [Section 3.2](#)). With RDI-P defined as described, STS PTE that supports ERDI-P can (but is not required to) determine which version of RDI-P the far-end STS PTE is sending from the incoming G1 bits 6 and 7. If the received bits 6 and 7 are set to either '01' or '10', then the far-end STS PTE is sending ERDI-P and the near-end can monitor G1 bits 5 through 7 to detect and terminate RDI-P defects. If the received bits 6 and 7 are set to either '00' or '11', then the far-end STS PTE is sending one-bit RDI-P and the near-end can monitor only G1 bit 5 to detect and terminate RDI-P defects.

R6-238 [958] If ERDI-P is supported and the STS PTE has detected two or more of the listed defects, it shall generate the higher priority ERDI-P code based on the priorities shown in [Table 6-4](#).

For example, if STS PTE that supports ERDI-P has detected an UNEQ-P defect (and is therefore sending the '110' code upstream) and then it detects an AIS-P defect, it must change its transmitted ERDI-P code to '101' within the 100 ms time period required by **R6-236 [556v2]**. (Also see **O6-237 [557v3]**, **R6-239 [558v2]** and **O6-240 [959v2]**.)

R6-239 [558v2] With the possible exception noted below, when STS PTE generates a particular type of RDI-P signal, it shall generate it for at least 10 frames.

O6-240 [959v2] With the possible exception noted below, when STS PTE generates a particular type of RDI-P signal, it should generate it for at least 20 frames.

Note that in cases where a higher priority incoming defect is detected before the ERDI-P code for a lower priority defect has been sent for the time specified in **R6-239 [558v2]** or **O6-240 [959v2]**, it is not necessary or recommended (but is allowed) for the STS PTE to delay changing its outgoing ERDI-P code in order to satisfy those criteria.

R6-241 [559v2] STS PTE shall deactivate (or change, as appropriate) an RDI-P signal within 100 ms of terminating the defect that caused it to be generated.

- O6-242 [560v2]** STS PTE should deactivate the RDI-P signal within 125 μ s of terminating the defect that caused it to be sent (assuming that the signal has been sent for the minimum assertion time supported by the NE, see **R6-239 [558v2]** and **O6-240 [959v2]**).
- R6-243 [561]** If **O6-237 [557v3]** and **O6-242 [560v2]** are not both met, then the delay time to generate an RDI-P signal (i.e., the time between detection of the defect and generation of the RDI-P signal) shall be within 500 μ s of the delay time to deactivate the RDI-P signal (i.e., the time between termination of the defect and deactivation of the RDI-P signal).
- R6-244 [960]** STS PTE that does not support ERDI-P shall detect a one-bit RDI-P defect when a '1' is received in bit 5 of G1 for 10 consecutive frames.
- R6-245 [562v3]** STS PTE that supports ERDI-P shall detect an RDI-P defect when one of the "RDI-P defect" codes shown in [Table 6-4](#) (one-bit or enhanced) is received for x consecutive frames, where x equals 5 or 10.
- R6-246 [961]** STS PTE that does not support ERDI-P shall terminate the one-bit RDI-P defect when a '0' is received in bit 5 of G1 for 10 consecutive frames.
- R6-247 [564v3]** STS PTE that supports ERDI-P shall terminate a particular type of RDI-P defect (one-bit or enhanced) when a code other than the code corresponding to that defect is received for x consecutive frames, where x equals 5 or 10.

Note that in some situations STS PTE that supports ERDI-P may simultaneously terminate one RDI-P defect and detect a different RDI-P defect.

- R6-248 [566v2]** An NE shall declare the corresponding one-bit RFI-P or ERFI-P failure when a particular type of RDI-P defect persists for 2.5 (± 0.5) seconds. Upon declaring an RFI-P failure, the NE shall set an RFI-P failure indication for the STS path and (if RFI-P is a Reported failure, see [Section 6.2.1.8](#)) send a message to an OS unless the condition in **R6-313 [629v3]** applies.
- R6-249 [962v3]** If ERDI-P is supported, the RFI-P failure indication shall indicate if the failure was derived from an incoming Server, Connectivity or Payload ERDI-P defect, or (if the capability to differentiate between one-bit RDI-P and ERDI-P Server defects is supported) a one-bit RDI-P defect (see [Table 6-4](#)).
- R6-250 [567v2]** An NE shall clear the corresponding RFI-P failure when the particular type of RDI-P defect that caused it to be declared is absent for 10 (± 0.5) seconds. Upon clearing the RFI-P failure, the NE shall clear the RFI-P failure indication and send a clear message to the OS (if the failure was reported to the OS).
-

Table 6-4 RDI-P Bit Settings and Interpretation

G1 Bits 5, 6 and 7	Priority of ERDI-P Codes	Trigger	Interpretation
0xx ^a	Not Applicable	No defects	No RDI-P defect
1xx ^a	Not Applicable	AIS-P ^b , LOP-P	One-bit RDI-P or ERDI-P Server defect ^c
001 ^d	4	No defects	No RDI-P defect
010 ^d	3	PLM-P, LCD-P ^e	ERDI-P Payload defect
101 ^d	1	AIS-P ^b , LOP-P	ERDI-P Server defect
110 ^d	2	UNEQ-P, TIM-P ^f	ERDI-P Connectivity defect

Notes:

- a. This code is transmitted by STS PTE that does not support ERDI-P. If ERDI-P is not supported, G1 bits 6 and 7 must be set to the same value, and should be set to '00' (see [Section 3.2](#)).
- b. In the functional model used in this document, incoming defects that are detected at STE or LTE (i.e., LOS, LOF, AIS-L) cause AIS-P to be generated downstream. This AIS-P is then detected by the STS PTE and causes it to generate RDI-P.
- c. Equipment that supports one-bit RDI-P will detect these codes as one-bit RDI-P defects. Equipment that supports ERDI-P may detect them as either one-bit RDI-P or ERDI-P Server defects.
- d. This code is transmitted by STS PTE that supports ERDI-P.
- e. STS Path Loss of Cell Delineation (LCD-P) is a Payload defect that is defined in ANSI T1.646 for ATM payloads carried in STS path signals using the ATM mapping shown in [Section 3.4.2.2.1](#).
- f. If TIM-P detection is activated for the terminating STS path.

6.2.1.3.3 VT Path Remote Defect Indication (RDI-V) and Remote Failure Indication (RFI-V)

An RDI-V signal is used in all VT-based applications to indicate to VT PTE that its peer VT PTE has detected a defect on the signal that it (the first VT PTE) originated. Four bits are reserved for the RDI-V signal (as shown in [Table 6-5](#)), and in most applications (i.e., applications other than those using the byte-synchronous DS1 mapping) detected RDI-V defects are used to derive RFI-V failures.

The details of the signals that VT PTE uses to notify upstream VT PTE that it has detected a defect underwent a number of changes as the SONET standards and Telcordia criteria evolved (i.e., from a "VT path yellow" signal using bit 8 of the V5 byte, to a one-bit RDI-V signal using bit 8 of V5 and a one-bit RFI-V signal using bit 4 of V5, to an enhanced two-bit RDI-V signal using bits 4 and 8 of V5 and a one-bit RFI-V signal using bit 8 of the Z7 byte). Currently an enhanced, four-bit RDI-V signal using bit 8 of V5 and bits 5 through 7 of Z7, and an RFI-V signal using bit 4 of V5 are defined in ANSI T1.105. The specifications in that standard are not expected to change, and therefore this section contains criteria that are relevant to those signals.

However, criteria for a one-bit RDI-V signal using bit 8 of V5 are included to cover products that have not yet implemented the enhanced version of RDI-V. Note that in this document, “RDI-V” is used to refer generically to both the one-bit and enhanced versions in text and criteria that are not specific to one particular version. In cases where it is necessary to distinguish between the two versions, the one-bit version is referred to as “one-bit RDI-V”, while the enhanced version is referred to as “ERDI-V”. Similarly, “RFI-V” is used to generically refer to a failure derived from a (generic) RDI-V defect, while “one-bit RFI-V” and “ERFI-V” are used in cases that are specific to the particular version.

Although ERDI-V was originally intended to be “backward compatible” with one-bit RDI-V, in the final version there are actually some significant differences. These differences are expected to primarily affect the accumulation of near-end and far-end VT path PM data. Therefore in applications where the VT path PM data is expected to be of importance, it is highly recommended that the VT PTE at both ends of the path (and any intermediate NEs provisioned to perform intermediate-path PM on that path) support the same version of RDI-V. In addition, it is recommended that the supported version be ERDI-V.

Finally, it should be noted that none of the criteria in this document are meant to imply that an NE that supports ERDI-V also needs to support one-bit RDI-V on a provisionable or, in the case of incoming signals, automatic-detection basis.¹⁰ Such capabilities may be provided by some NEs (either per-VT-path or NE-wide); however, it is also considered acceptable for an NE to support only one version of RDI-V.

-
- O6-251 [963v2]** An NE should support enhanced RDI-V generation and detection.
- R6-252 [568v2]** VT PTE shall generate an appropriate RDI-V signal, as specified in [Table 6-5](#), within 100 ms of detecting a listed defect. (Also see [Figures 6-7](#) through [6-12](#).)
- O6-253 [569v3]** With the possible exception noted following **O6-256 [965v2]**, VT PTE should generate an appropriate RDI-V signal as specified in [Table 6-5](#) within 500 μs of detecting a listed defect.
- R6-254 [964]** If ERDI-V is supported and the VT PTE has detected two or more of the listed defects, it shall generate the higher priority ERDI-V code based on the priorities shown in [Table 6-5](#).
-

10. For example, an NE that supports ERDI-V does not need to be able to be provisioned to transmit one-bit RDI-V, to provide different failure indications for incoming one-bit RDI-V and ERDI-V Server signals, or to change its near-end intermediate VT path PM parameter definitions based on the version of RDI-V being used by the end NEs (see [Sections 6.2.2.6.1](#) and [6.2.2.9](#)).

For example, if VT PTE that supports ERDI-V has detected an UNEQ-V defect (and is therefore sending the '1' and '110' codes upstream) and then it detects an AIS-V defect, it must change its transmitted ERDI-V code to '1' and '101' within the 100 ms time period required by **R6-252 [568v2]**. (Also see **O6-253 [569v3]**, **R6-255 [570v2]** and **O6-256 [965v2]**.)

R6-255 [570v2] With the possible exception noted below, when VT PTE generates a particular type of RDI-V signal, it shall generate that signal for at least 10 VT superframes.

O6-256 [965v2] With the possible exception noted below, when VT PTE generates a particular type of RDI-V signal, it should generate it for at least 20 VT superframes.

Note that in cases where a higher priority incoming defect is detected before the ERDI-V code for a lower priority defect has been sent for the time specified in **R6-255 [570v2]** or **O6-256 [965v2]**, it is not necessary or recommended (but is allowed) for the VT PTE to delay changing its outgoing ERDI-V code to satisfy those criteria.

R6-257 [571v2] VT PTE shall deactivate (or change, as appropriate) an RDI-V signal within 100 ms of terminating the defect that caused it to be sent.

O6-258 [572v2] VT PTE should deactivate the RDI-V signal within 500 μ s of terminating the defect that caused it to be sent (assuming that the signal has been sent for the minimum assertion time supported by the NE, see **R6-255 [570v2]** and **O6-256 [965v2]**).

R6-259 [573] If **O6-253 [569v3]** and **O6-258 [572v2]** are not both met, then the delay time to generate the RDI-V signal (i.e., the time between detection of the defect and generation of the RDI-V signal) shall be within 2 ms of the delay time to deactivate the RDI-V signal (i.e., the time between termination of the defect and deactivation of the RDI-V signal).

Note that for the ERDI-V codes shown in [Table 6-5](#), bits 6 and 7 of the Z7 byte are always set to opposite values. Conversely, in the one-bit RDI-V signal those two bits must be set to the same value (and should be set to '00'). With RDI-V defined as described, VT PTE that supports ERDI-V determines which version of RDI-V the far-end VT PTE is sending from the incoming Z7 bits 6 and 7. If the received bits 6 and 7 are set to either '01' or '10', then the far-end VT PTE is sending ERDI-V and the near-end monitors the Z7 bits to detect and terminate RDI-V defects. If the received bits 6 and 7 are set to either '00' or '11', then the far-end VT PTE is sending one-bit RDI-V and the near-end monitors V5 bit 8 to detect and terminate RDI-V defects.

R6-260 [966] VT PTE that does not support ERDI-V shall detect a one-bit RDI-V defect when a '1' is received in bit 8 of V5 for 10 consecutive VT superframes.

- R6-261 [574v3]** VT PTE that supports ERDI-V shall detect an RDI-V defect when one of the “RDI-V defect” codes shown in [Table 6-5](#) (one-bit or enhanced) is received for x consecutive VT superframes, where x equals 5 or 10.
- R6-262 [967]** VT PTE that does not support ERDI-V shall terminate the one-bit RDI-V defect when a ‘0’ is received in bit 8 of V5 for 10 consecutive VT superframes.
- R6-263 [576v3]** VT PTE that supports ERDI-V shall terminate a particular type of RDI-V defect (one-bit or enhanced) when a code other than the code corresponding to that defect is received for x consecutive VT superframes, where x equals 5 or 10.

Note that in some situations, VT PTE that supports ERDI-V may simultaneously terminate one RDI-V defect and detect a different RDI-V defect.

- R6-264 [578v2]** An NE that is not using the byte-synchronous DS1 mapping for a particular VT path shall declare the corresponding one-bit RFI-V or ERFI-V failure when a particular type of RDI-V defect persists for 2.5 (± 0.5) seconds. Upon declaring an RFI-V failure, the NE shall set an RFI-V failure indication for the VT path and (if RFI-V is a Reported failure, see [Section 6.2.1.8](#)) send a message to an OS unless the condition in **R6-313 [629v3]** applies.
- R6-265 [968v3]** If ERDI-V is supported, the RFI-V failure indication shall indicate if the failure was derived from an incoming Server, Connectivity or Payload ERDI-V defect, or (if the capability to differentiate between one-bit RDI-V and ERDI-V Server defects is supported) a one-bit RDI-V defect (see [Table 6-5](#)).
- R6-266 [579v2]** An NE that is not using the byte-synchronous DS1 mapping for a particular VT path shall clear the corresponding RFI-V failure when the particular type of RDI-V defect that caused it to be declared is absent for 10 (± 0.5) seconds. Upon clearing the RFI-V failure, the NE shall clear the RFI-V failure indication and send a clear message to the OS (if the failure was reported to the OS).
-

Table 6-5 RDI-V Bit Settings and Interpretation

Z7 Bits 5, 6, 7	V5 Bit 8	Priority of ERDI-V Codes	Trigger	Interpretation
yxx ^a	0	Not Applicable	No defects	No RDI-V defect
yxx ^a	1	Not Applicable	AIS-V ^b , LOP-V	One-bit RDI-V or ERDI-V Server defect ^c
001 ^d	0 ^e	4	No defects	No RDI-V defect
010 ^d	0 ^e	3	PLM-V	ERDI-V Payload defect
101 ^d	1 ^e	1	AIS-V ^b , LOP-V	ERDI-V Server defect
110 ^d	1 ^e	2	UNEQ-V, TIM-V ^f	ERDI-V Connectivity defect
0xx ^g	1 ^g	Not Applicable	Not Applicable	One-bit RDI-V defect or ERDI-V Server defect ^c
1xx ^g	0 ^g	Not Applicable	Not Applicable	No RDI-V defect

Notes:

- a. This code is transmitted by VT PTE that does not support ERDI-V. If ERDI-V is not supported, Z7 bits 6 and 7 must be set to the same value, and Z7 bits 5, 6 and 7 should be set to '000' (see [Section 3.2](#)).
- b. In the functional model used in this document, incoming defects that are detected at STE, LTE or STS PTE (e.g., LOS, LOF, AIS-L, AIS-P) cause AIS-V to be generated downstream. This AIS-V is then detected by the VT PTE and causes it to generate RDI-V.
- c. Equipment that supports one-bit RDI-V will detect these codes as one-bit RDI-V defects. Equipment that supports ERDI-V may detect them as either one-bit RDI-V or ERDI-V Server defects.
- d. This code is transmitted by VT PTE that supports ERDI-V.
- e. V5 bit 8 is set to the same value as Z7 bit 5 by the VT PTE that supports ERDI-V. At the receiving VT PTE, V5 bit 8 is ignored unless Z7 bits 6 and 7 are both set to '0' or both set to '1'.
- f. If TIM-V detection is activated for the terminating VT path.
- g. This code (in which V5 bit 8 and Z7 bit 5 are set to opposite values and Z7 bits 6 and 7 are also set to opposite values) will not be generated by any equipment that conforms to either the one-bit RDI-V or ERDI-V generation criteria contained in this document, and therefore is never expected to be present in the network. It is included here for completeness.

In addition to RDI-V signals, an RFI-V signal is also defined for use in applications where the byte-synchronous DS1 mapping is used. In those applications, it is the reception of the RFI-V signal that causes the declaration of the RFI-V failure for the path. This declaration can then be used to initiate trunk conditioning when required. Also, in some applications RFI-V signals are translated to and from DS1 RAI signals at a DS1 interface. Bit 4 of the V5 byte carries the RFI-V signal.

- R6-267 [580v4]** VT PTE that is using the byte-synchronous DS1 mapping for a particular VT path shall generate an RFI-V signal by setting bit 4 of the V5 byte to '1' within 500 μ s of declaring an AIS-V failure (or a lower-layer, traffic-related, near-end failure, see [Section 6.2.1.8.2](#)), an LOP-V failure, an UNEQ-V failure, a TIM-V failure, or a PLM-V failure (as shown in Figures 6-9 through 6-12).
- R6-268 [969]** VT PTE that is byte-synchronously mapping a DS1 into a single VT1.5 (i.e., no DS0 rearrangement) shall generate an RFI-V signal by setting bit 4 of the V5 byte to '1' within 500 μ s of declaring a DS1 RAI failure for an incoming DS1 (as shown in Figure 6-9).
- R6-269 [581v2]** VT PTE that is using the byte-synchronous DS1 mapping shall deactivate the RFI-V signal by setting bit 4 of V5 to '0' within 500 μ s of clearing the failure that caused the RFI-V signal to be sent.
- R6-270 [582v2]** An NE that is using the byte-synchronous DS1 mapping for a particular VT path shall declare an RFI-V failure upon receiving a '1' in bit 4 of V5 for 10 consecutive VT superframes (i.e., for 5 ms). Upon declaring the RFI-V failure, the NE shall set an RFI-V failure indication for the VT path and (if RFI-V is a Reported failure, see [Section 6.2.1.8](#)) send a message to an OS unless the condition in **R6-313 [629v3]** applies.
- R6-271 [583v2]** An NE shall clear the RFI-V failure within 50 ms of receiving a '0' in bit 4 of V5 for 10 consecutive VT superframes. Upon clearing the RFI-V failure, the NE shall clear the RFI-V failure indication and send a clear message to the OS (if the failure was reported to the OS).

For testing an NE's conformance to **R6-271 [583v2]**, it is sufficient to send '0' in bit 4 of V5 for less than 10 consecutive VT superframes (the RFI-V failure must not be cleared), and for more than 50 ms (the RFI-V failure must be cleared).

Finally, VT PTE that uses the byte-synchronous DS1 mapping for a particular VT path may optionally provide the ability to use the RDI-V signal to derive RFI-V failures, in addition to the requirement to use the RFI-V signal. However, only the RFI-V signal is used to initiate trunk conditioning (see [Section 6.2.1.6](#)).

6.2.1.3.4 DS_n RDI and RAI Signals

In some applications a SONET NE may need to generate or detect DS_n RDI or RAI signals, including DS1 RAI-Customer Installation (DS1 RAI-CI) signals. Figures 6-8 through 6-12 (in [Section 6.2.1.7](#)) illustrate the use of DS_n RDI and RAI for STS and VT PTE terminating SPEs of various compositions. GR-499-CORE contains the requirements on the construction of DS1, DS1C, DS2 and DS3 RAI, and DS3 RDI,¹¹

11. DS3 RDI uses the X-bits in the DS3 frame, and is defined for M23 and C-bit parity applications. DS3 RAI is defined only in C-bit parity applications, and consists of a subset of the defined Far End Alarm and Control (FEAC) signals. See GR-499-CORE for additional information.

along with a conditional requirement to support the generation and detection of DS1 RAI-CI and a reference to ANSI T1.403 for additional specifications related to that type of signal.

DS0 RAI is defined in SONET as a specific pattern in the ABCD signaling bits (i.e., A = 0, B = 1, C = 1, D = 1). Note that DS0 RAI may not be applicable outside of SONET networks, so an interface to a non-SONET NE may require the application of service-specific trunk conditioning codes.

-
- CR6-272 [970]** A SONET NE that terminates an M23 application DS3 path may be required to generate DS3 RDI upstream upon detecting certain defects, as shown in [Figure 6-8](#).
- R6-273 [971]** A SONET NE that terminates a C-bit parity application DS3 path shall generate DS3 RDI upstream upon detecting certain defects, as shown in [Figure 6-8](#).
- R6-274 [972]** A SONET NE shall remove a DS3 RDI upon terminating the defect that caused it to be generated.
- R6-275 [584v2]** A SONET NE that terminates a DS_n path other than an M23 application DS3 path shall generate DS_n RAI upstream (unless it is provisioned to apply a service-specific trunk conditioning code upstream) upon declaring certain failures, as shown in [Figures 6-8, 6-11 and 6-12](#).
- R6-276 [586v2]** VT PTE provisioned to byte-synchronously map a DS1 into a single VT1.5 (i.e., no DS0 rearrangement) shall generate DS1 RAI downstream upon declaring an RFI-V failure as a result of receiving an RFI-V signal, as shown in [Figure 6-9](#).

As discussed in [Section 6.2.1.3.3](#), VT PTE's generation of DS1 RAI is based on the detection of an incoming RFI-V signal and the subsequent declaration of an RFI-V failure, not on RFI-V failures derived from detected RDI-V defects.

-
- R6-277 [973]** A SONET NE that provides access to and processing of individual DS0 channels shall generate DS0 RAI upstream upon declaring certain failures that cause trunk conditioning to be applied downstream (unless it is provisioned to also apply a service-specific trunk conditioning code upstream), as shown in [Figures 6-10 and 6-11](#).

Note that DS0 RAI is not sent upstream if DS0 AIS is generated (or allowed to pass) downstream, or if the NE is also provisioned to apply trunk conditioning in the upstream direction. It is sent upstream only if the NE is provisioned to apply trunk conditioning only in the downstream direction.

-
- R6-278 [974]** A SONET NE that provides access to and processing of individual DS0 channels shall generate DS0 RAI downstream when it declares a DS1 RAI failure (unless it is provisioned to apply a service-specific trunk conditioning code), as shown in [Figure 6-11](#).

- R6-279 [975]** A SONET NE shall remove a DS_n RAI upon clearing the failure that caused it to be generated.

In addition to generating and removing DS_n RDI and RAI, in some applications a SONET NE may also need to be able to monitor for DS_n RDI and RAI (e.g., to detect incoming DS_n RAI signals, and declare and clear DS_n RAI failures).

- CR6-280 [976]** A SONET NE that terminates M23 application DS3 paths may be required to detect DS3 RDI for those paths.

- R6-281 [977]** A SONET NE that terminates C-bit parity application DS3 paths shall detect DS3 RDI for those paths.

- R6-282 [978]** A SONET NE that terminates DS_n paths shall detect DS_n RAI signals (if defined) for those paths.

- R6-283 [585v2]** VT PTE that byte-synchronously maps a DS1 into a single VT1.5 (i.e., no DS0 rearrangement) shall detect a DS1 RAI signal received at the DS1 interface (and insert RFI-V downstream as shown in [Figure 6-9](#)).

- R6-284 [979]** A SONET NE that provides access to and processing of individual DS0 channels, and that is provisioned to apply trunk conditioning in a particular direction shall detect DS0 RAI signals in that direction (and apply the trunk conditioning code).

Note that if the NE is not provisioned to apply trunk conditioning in a particular direction, then DS0 RAI is simply passed through in that direction and does not need to be detected by the NE.

If an NE detects DS0 RAI signals, the following requirements apply.

-
- R6-285 [587v2]** A SONET NE that detects DS0 RAI signals shall declare a DS0 RAI failure if it receives two consecutive sets of ABCD signaling bits set to the code '0111' (i.e., the code '0111' persisting for 6 ms).

- R6-286 [588v2]** A SONET NE shall clear a DS0 RAI failure if it receives four consecutive sets of ABCD signaling bits set to any code other than the '0111' code (i.e., any code other than '0111' persisting for 12 ms).
-

6.2.1.4 Payload Defect Indication

PDI is an application-specific code inserted by PTE to indicate to downstream equipment that there is a defect in one or more of its directly mapped embedded payloads. PDI is expected to be used primarily in ring interconnection applications, and possibly in some dual-homing applications. For example, in an application where DCSs are used to perform VT-level grooming between two dual-interconnected STS level UPSRs, the DCSs and all of the NEs on the rings

would have to support PDI generation and/or detection. Note that PDI is currently only defined at the STS path level (i.e., PDI-P). Several options have been proposed for a VT level PDI (i.e., PDI-V carried as a code in the VT signal label, or as AIS-V); however, those options have not been accepted by ANSI and are not included here.

6.2.1.4.1 STS Path Payload Defect Indication

PDI-P is a set of application-specific codes contained in the STS POH generated by STS PTE. It is used to indicate to downstream equipment (e.g., to the path selector in a downstream STS-level UPSR NE) that there is a defect in one or more of the directly mapped payloads contained in that STS SPE. The PDI-P codes appear in the STS signal label (C2 byte). Valid codes are shown in [Table 3-3](#).

CR6-287 [591v3] STS PTE that supports a VT-structured STS-1 SPE or the asynchronous DS3 mapping may be required to also support PDI-P signal generation.

CR6-288 [589v2] STS PTE that supports PDI-P generation may be required to be provisionable as to whether it sends PDI-P signals.

Currently, PDI-P generation by STS PTE that supports payload mappings/structures other than those listed in **CR6-287 [591v3]** is FFS. Also, the capability to disable PDI-P generation would be needed (for example) for backwards compatibility with older NEs that would consider any of the PDI-P codes to be mismatched to the locally provisioned STS PTE functionality.

R6-289 [980] If STS PTE that supports PDI-P generation and a VT-structured STS SPE does not process VT pointers, it shall (non-intrusively) detect and terminate LOP-V defects as if it were processing the those pointers (i.e., according to the criteria in [Section 6.2.1.1.3](#)).

R6-290 [981] STS PTE that supports PDI-P generation and a VT-structured STS SPE shall (non-intrusively) detect and terminate AIS-V defects as if it were VT PTE (i.e., according to the criteria in [Section 6.2.1.2.3](#)).

R6-291 [982] STS PTE that supports PDI-P generation and the asynchronous DS3 mapping shall (non-intrusively) detect and terminate DS3 AIS defects as if it were terminating the DS3 path (i.e., according to the criteria in [Section 6.2.1.2.4](#)).

CR6-292 [593v2] STS PTE that supports PDI-P generation and the asynchronous DS3 mapping may be required to be provisionable to (non-intrusively) detect and terminate DS3 OOF defects as if it were terminating the DS3 path (i.e., according to the criteria in [Section 6.2.1.1.2](#)).

See [Section 6.2.1.8.7](#) for criteria concerning the declaration and clearing of failures based on the defects detected and terminated according to the preceding criteria.

-
- R6-293 [983]** As a default STS PTE shall not monitor for DS3 OOF for the purposes of generating PDI-P.
- R6-294 [592v2]** STS PTE that supports PDI-P generation shall generate (or change, as appropriate) the PDI-P signal within 100 ms of detecting an LOP-V, AIS-V, DS3 AIS, DS3 LOS, or (if so provisioned) DS3 OOF defect on any VT or DS3 payload that it embeds into the STS SPE that it is originating. The PDI-P signal shall be generated by inserting the code indicating the number of defective payloads as specified in [Table 3-3](#).
- R6-295 [595v2]** STS PTE that supports PDI-P generation shall deactivate (or change, as appropriate) the PDI-P signal within 100 ms of terminating a defect on one or more of its payloads.
-

Note that the specifications in ANSI T1.105 allow the STS PTE to generate, change or deactivate its PDI-P signals anytime within 250 ms (rather than 100 ms) after a triggering defect is detected or terminated. Thus, equipment that conforms to these requirements will also meet those specifications.

-
- CR6-296 [984]** A SONET NE may be required to support the detection of PDI-P signals.
- CR6-297 [985]** A SONET NE that supports PDI-P detection may be required to be provisionable (on a per-STS-path basis) as to whether it detects PDI-P.
- R6-298 [596v2]** A SONET NE that supports the detection of PDI-P shall detect a PDI-P defect (or a change in the PDI-P defect, as appropriate) within 10 ms of the onset of at least five consecutive samples (which might not be consecutive frames) of STS signal labels (C2 bytes) containing a new PDI-P code.
- O6-299 [597v2]** A SONET NE that supports the detection of PDI-P should detect a PDI-P defect (or a change in the PDI-P defect, as appropriate) immediately upon receipt of five consecutive frames of STS signal labels containing a new PDI-P code.
- R6-300 [598]** A SONET NE that supports the detection of PDI-P shall terminate a PDI-P defect within 10 ms of the onset of at least five consecutive samples (which might not be consecutive frames) of STS signal labels that do not contain a PDI-P code.
- O6-301 [599]** A SONET NE that supports the detection of PDI-P should terminate a PDI-P defect immediately upon receipt of five consecutive frames of STS signal labels that do not contain a PDI-P code.
-

Unlike most of the other defects defined in SONET, an NE that detects a persistent PDI-P defect is not required to declare a PDI-P failure or send a PDI-P failure report to an OS.¹² The reason for this is that PDI-P defects are detected when there is a problem (i.e., a defect) with an embedded payload, not with the STS SPE itself. The root-cause defect will normally be reported to the OS (if it persists so that a failure

is declared) by some NE other than the one that *detects* the PDI-P defect. For example, an NE that detects a DS3 LOS and *inserts* the appropriate PDI-P code (and DS3 AIS) downstream will normally declare a DS3 LOS failure.

Since PDI-P defects are expected to be used for STS path protection switching purposes in some applications, they can have a direct impact on the traffic carried by a SONET system and the user may need to determine the particular codes that are being transmitted and received. [Section 6.2.3.2.3.B](#) contains the criteria on diagnostics to report the STS signal label.

6.2.1.4.2 VT Path Payload Defect Indication

PDI-V is currently not defined in SONET.

6.2.1.5 Maintenance Signals for Other Mappings

The criteria for DS4NA, FDDI, ATM and DQDB maintenance signals are FFS.

6.2.1.6 Trunk Conditioning

Trunk conditioning is used during carrier failures to release connections, terminate customer billing, and remove trunks from service by instructing the switch or customer Private Branch Exchange (PBX) to go idle and make the trunks busy. In a digital bit stream, trunk conditioning codes (i.e., one or more provisioned ABCD signaling codes and an 8-bit insertion word) are inserted into each DS0 channel affected by the carrier failure.¹³ For incoming signal defects that would otherwise cause DS0 AIS to be inserted or passed downstream, the signaling is frozen (i.e., the last values for the signaling bits are used) when the defect is initially detected. Then the defect is soaked for a time interval of approximately 2.5 seconds (or 3.25 seconds for DS0 AIS defects) to verify that a failure should be declared. In addition, in most cases an integration timer is used to account for intermittent defects. For other signals that can trigger the application of trunk conditioning (i.e., RFI-V, DS1 RAI and DS0 RAI signals), the detection of the signal causes a failure to

12. If an NE does declare PDI-P failures, it is expected that the process used will be consistent with the failure declaration processes required for other defects/failures defined in this document. In addition, PDI-P defects are generally expected to be used for STS path protection switching purposes and the reader should see the appropriate application-specific GRs (e.g., GR-1230-CORE, GR-1400-CORE) for any additional criteria related to the reports an NE generates when it performs a path protection switch.

13. The required bit pattern or patterns depend on the service being carried by the DS0 channel, and thus may differ for different services. Most services require a single (service-specific) code for trunk conditioning. Some require the application of one (service-specific) code followed by another to support a 2-stage (make-idle/make-busy) form of trunk conditioning. The trunk conditioning signaling codes are accompanied by an 8-bit insertion word. TR-NWT-000170 identifies signaling and 8-bit insertion codes for various types of access circuits.

be declared immediately (i.e., there is no soaking interval). In both cases, the trunk conditioning code or codes are applied to each affected DS0 when the failure is declared.

In SONET, trunk conditioning is applicable only when access to and processing of individual DS0 channels are provided in the SONET signal. Therefore, SONET NEs that support the byte-synchronous VT mapping in conjunction with DS0 rearrangement or termination functions need to support trunk conditioning functions. SONET NEs that support the byte-synchronous VT mapping but not DS0 rearrangement or termination functions do not need to support trunk conditioning functions. VT and DS1 level maintenance signaling is sufficient for those cases.

The maintenance signaling and trunk conditioning requirements related to byte-synchronously mapped DS1s at a SONET NE that performs DS0 rearrangement or DS0 termination are summarized in Figures 6-10 through 6-12. Note that Figure 6-9 applies to the same mapping at a DS1 interface where no DS0 rearrangement or termination functions are performed.

-
- R6-302 [610v4]** An NE that provides access to and processing of individual DS0 channels shall set a red alarm when it declares an LOS, LOF, LOP-P, UNEQ-P, TIM-P, PLM-P, LOP-V, UNEQ-V, TIM-V or PLM-V failure on the incoming signal.
- R6-303 [612]** An NE shall clear a red alarm when the failure that caused it to be set has cleared.
- R6-304 [986]** An NE that provides access to and processing of individual DS0 channels shall allow the user to provision it to apply trunk conditioning codes (for use in place of downstream DS0 AIS, downstream DS0 RAI, or upstream DS0 RAI, as applicable) on a per-DS0 basis. If the DS0 is not terminated, then the user shall be able to provision separate codes in each direction.
-

Note that trunk conditioning can only be applied in the upstream direction at DS0 PTE.

-
- R6-305 [609v2]** If it is provisioned to apply a trunk conditioning code on a particular DS0, the SONET NE shall freeze the DS0 signaling state upon detecting a defect that would otherwise cause DS0 AIS to be generated or passed downstream as shown in Figures 6-10 and 6-11.
- R6-306 [611v2]** If a defect that causes a SONET NE to freeze DS0 signaling states persists so that the NE declares the associated failure, the NE shall apply the provisioned downstream trunk conditioning codes.
- R6-307 [987]** If a defect that causes a SONET NE to freeze DS0 signaling states persists so that it declares a failure, the NE shall apply the provisioned upstream trunk conditioning codes (if any) instead of DS0 RAI.

R6-308 [617v2] If it is provisioned to apply a trunk conditioning code on a particular DS0, the SONET NE shall apply the provisioned code downstream upon declaring an RFI-V, DS1 RAI, or DS0 RAI failure as a result of an incoming RFI-V, DS1 RAI, or DS0 RAI signal (as shown in Figures 6-10 and 6-11).

R6-309 [618v2] A SONET NE shall remove trunk conditioning upon clearing the failure that had caused it to be applied.

6.2.1.7 Alarm-Related Events

Figures 6-3 through 6-13 illustrate the maintenance signals sent by SONET NEs when various defects are detected on incoming signals or certain equipment failures are declared. Those figures are followed by Figures 6-14 through 6-17, which illustrate the timing sequence for SONET NEs detecting defects, declaring failures and setting their indications. These timing sequences are summarized below. Note that timing requirements regarding DS0 AIS and trunk conditioning signals and functions are not depicted in these particular figures (see Sections 6.2.1.2.4 and 6.2.1.6).

As Figures 6-14 through 6-17 indicate, downstream AIS and upstream RDI signals are sent immediately upon detecting the appropriate defect. Declaring a failure, setting the indication and reporting these events to the OS are delayed for 2.5 (± 0.5) seconds to verify that the defect persists. Where the RFI-V signal is used, the signal is sent upstream when the failure is declared.

When a defect is terminated, downstream AIS and upstream RDI signals are removed immediately. When a failure is cleared on a signal using RFI-V, the upstream RFI-V signal is removed.

Figure 6-3 STE Maintenance Signals


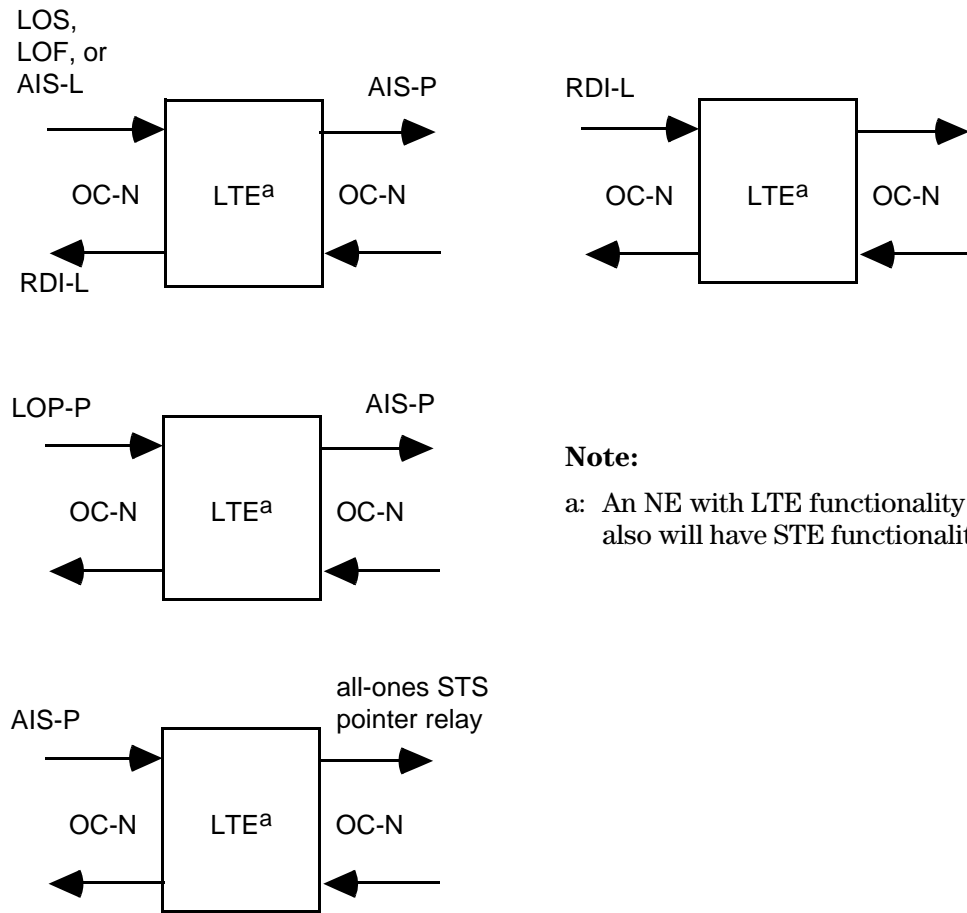


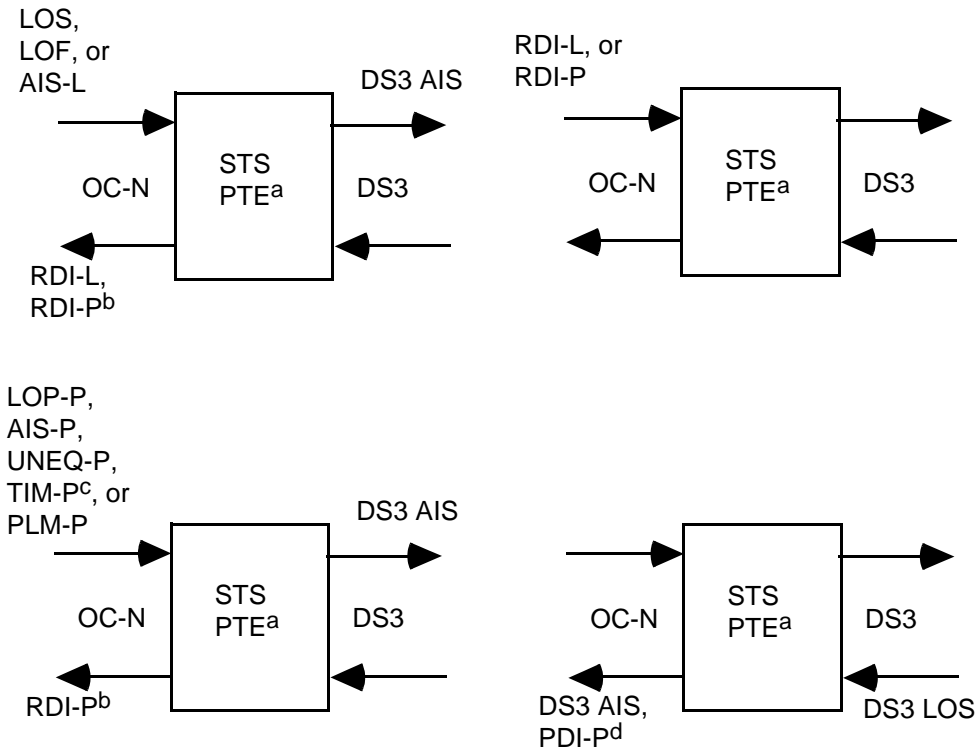
Figure 6-4 LTE Maintenance Signals



Note:

a: An NE with LTE functionality also will have STE functionality.

Figure 6-5 STS PTE (Asynchronous Mapping for DS3) Maintenance Signals



Notes:

- a: An NE with STS PTE functionality also will have STE and LTE functionality.
- b: See [Section 6.2.1.3.2](#) for the applicable RDI-P code for each type of incoming defect.
- c: If TIM-P detection is activated for the STS path.
- d: Provisionable.
- e: Provisionable to cause generation of PDI-P and/or DS3 AIS.


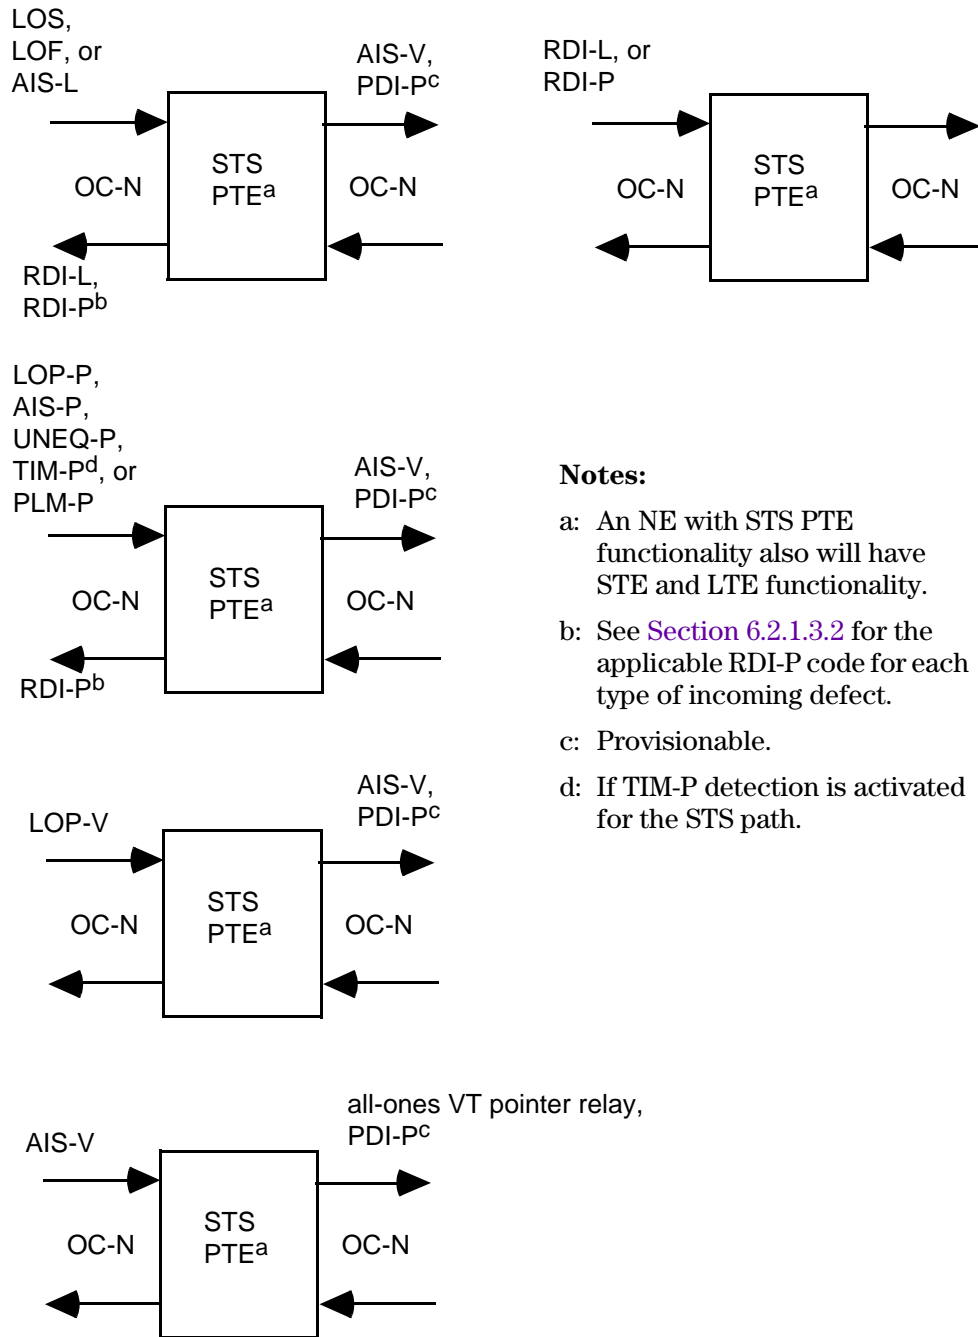


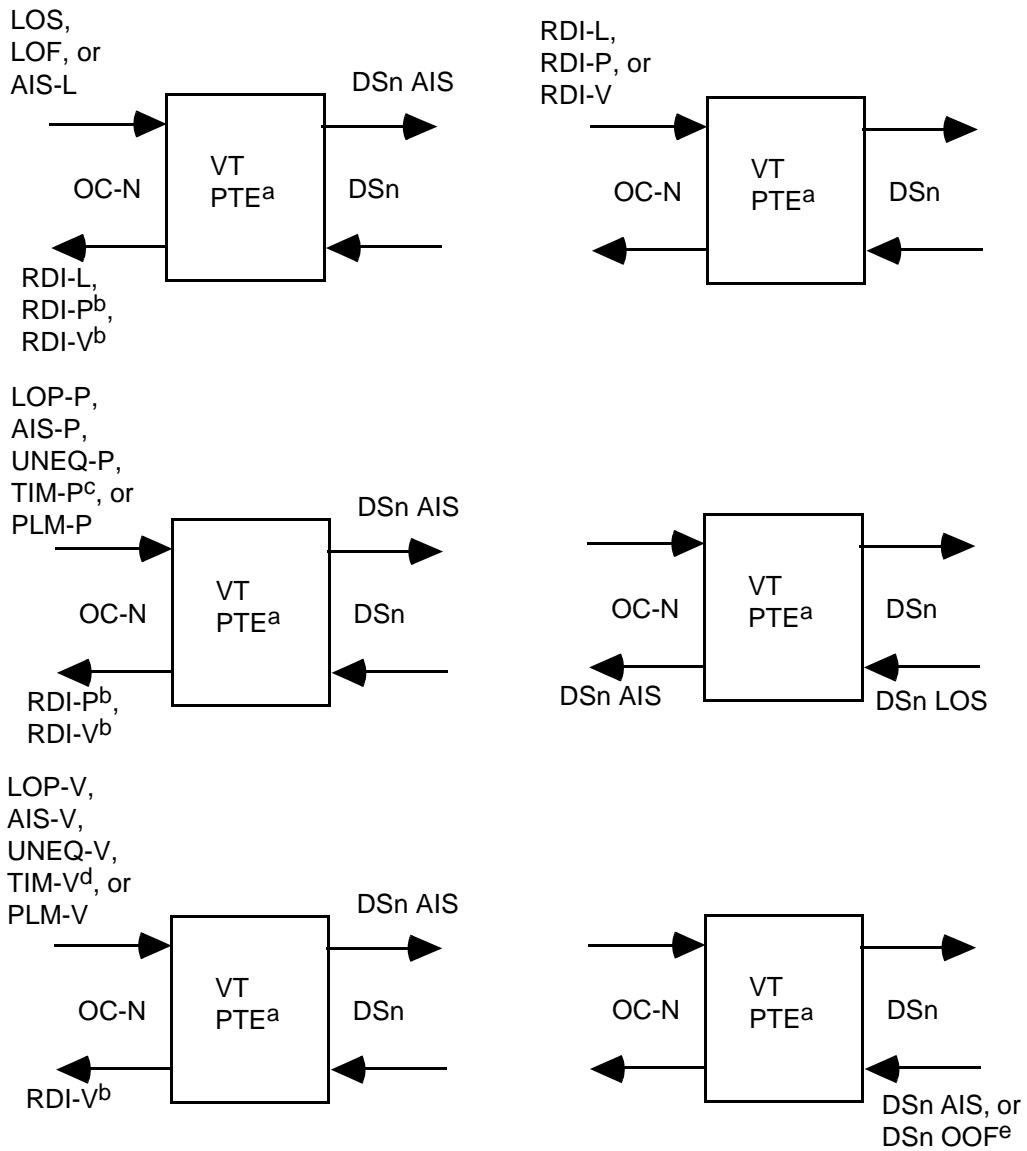
Figure 6-6 STS PTE (VT-Structured STS-1 SPE) Maintenance Signals



Notes:

- a: An NE with STS PTE functionality also will have STE and LTE functionality.
- b: See [Section 6.2.1.3.2](#) for the applicable RDI-P code for each type of incoming defect.
- c: Provisionable.
- d: If TIM-P detection is activated for the STS path.


Figure 6-7 VT PTE (Asynchronous Mapping for DS_n into VT_n) Maintenance Signals, DS_n Interface



Notes:

- a: An NE with VT PTE functionality also will have STE, LTE and STS PTE functionality.
- b: See [Sections 6.2.1.3.2](#) and [6.2.1.3.3](#) for the applicable RDI-P and RDI-V codes for each type of incoming defect.
- c: If TIM-P detection is activated for the STS path.
- d: If TIM-V detection is activated for the VT path.
- e: Provisionable to cause generation of DS_n AIS.

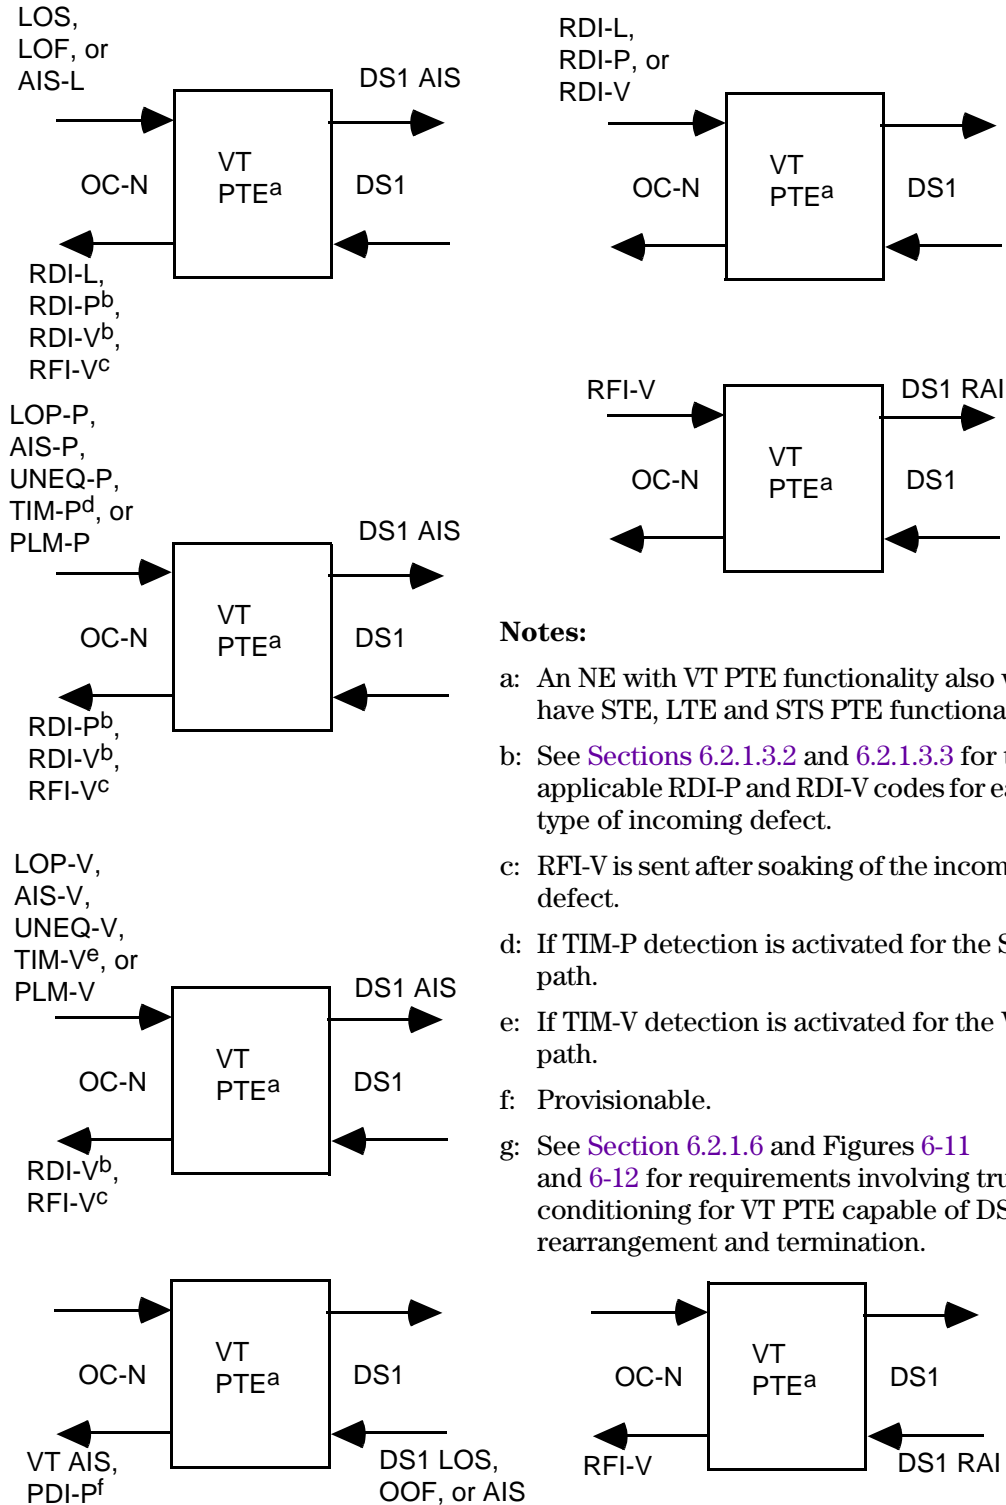
Figure 6-8 VT PTE (Asynchronous Mapping for DS1 into VT1.5) Maintenance Signals, DS3 Interface with Embedded M13 Multiplex



Notes:

- a: An NE with VT PTE functionality also will have STE, LTE and STS PTE functionality.
- b: DS3 RAI is sent after soaking of the incoming defect.

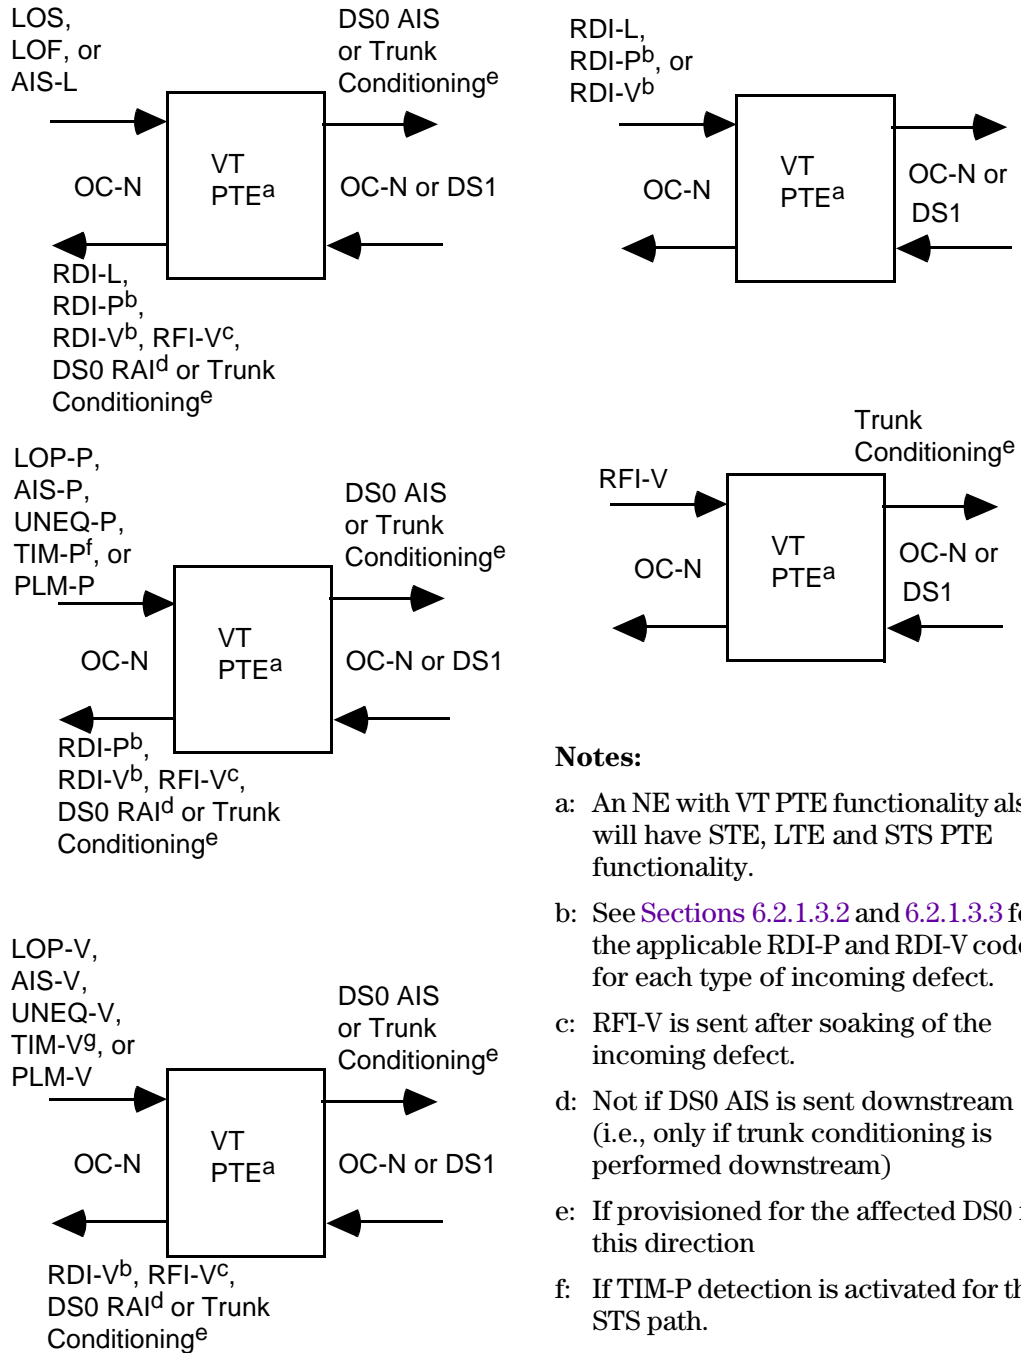
Figure 6-9 VT PTE (Byte-Synchronous Mapping for DS1 into a Single VT1.5) Maintenance Signals, DS1 Interface



Notes:

- a: An NE with VT PTE functionality also will have STE, LTE and STS PTE functionality.
- b: See [Sections 6.2.1.3.2](#) and [6.2.1.3.3](#) for the applicable RDI-P and RDI-V codes for each type of incoming defect.
- c: RFI-V is sent after soaking of the incoming defect.
- d: If TIM-P detection is activated for the STS path.
- e: If TIM-V detection is activated for the VT path.
- f: Provisionable.
- g: See [Section 6.2.1.6](#) and [Figures 6-11](#) and [6-12](#) for requirements involving trunk conditioning for VT PTE capable of DS0 rearrangement and termination.

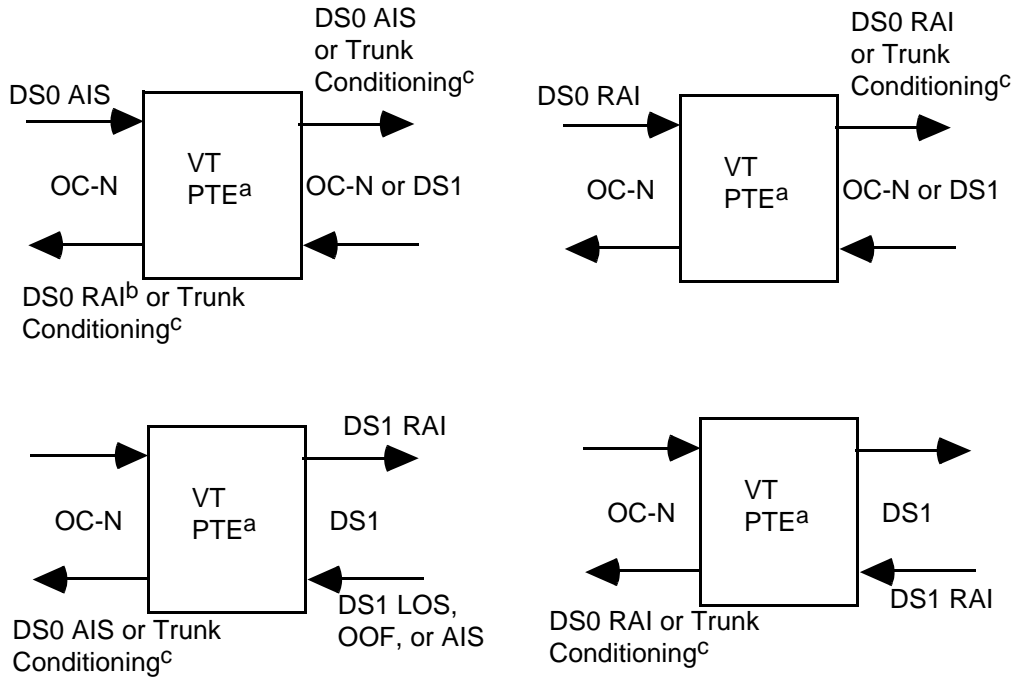
Figure 6-10 VT PTE with DS0 Rearrangement Functions (Byte-Synchronous Mapping for DS1 into VT1.5) Maintenance Signals and Trunk Conditioning, OC-N and DS1 Interfaces



Notes:

- a: An NE with VT PTE functionality also will have STE, LTE and STS PTE functionality.
- b: See Sections 6.2.1.3.2 and 6.2.1.3.3 for the applicable RDI-P and RDI-V codes for each type of incoming defect.
- c: RFI-V is sent after soaking of the incoming defect.
- d: Not if DS0 AIS is sent downstream (i.e., only if trunk conditioning is performed downstream)
- e: If provisioned for the affected DS0 in this direction
- f: If TIM-P detection is activated for the STS path.
- g: If TIM-V detection is activated for the VT path.

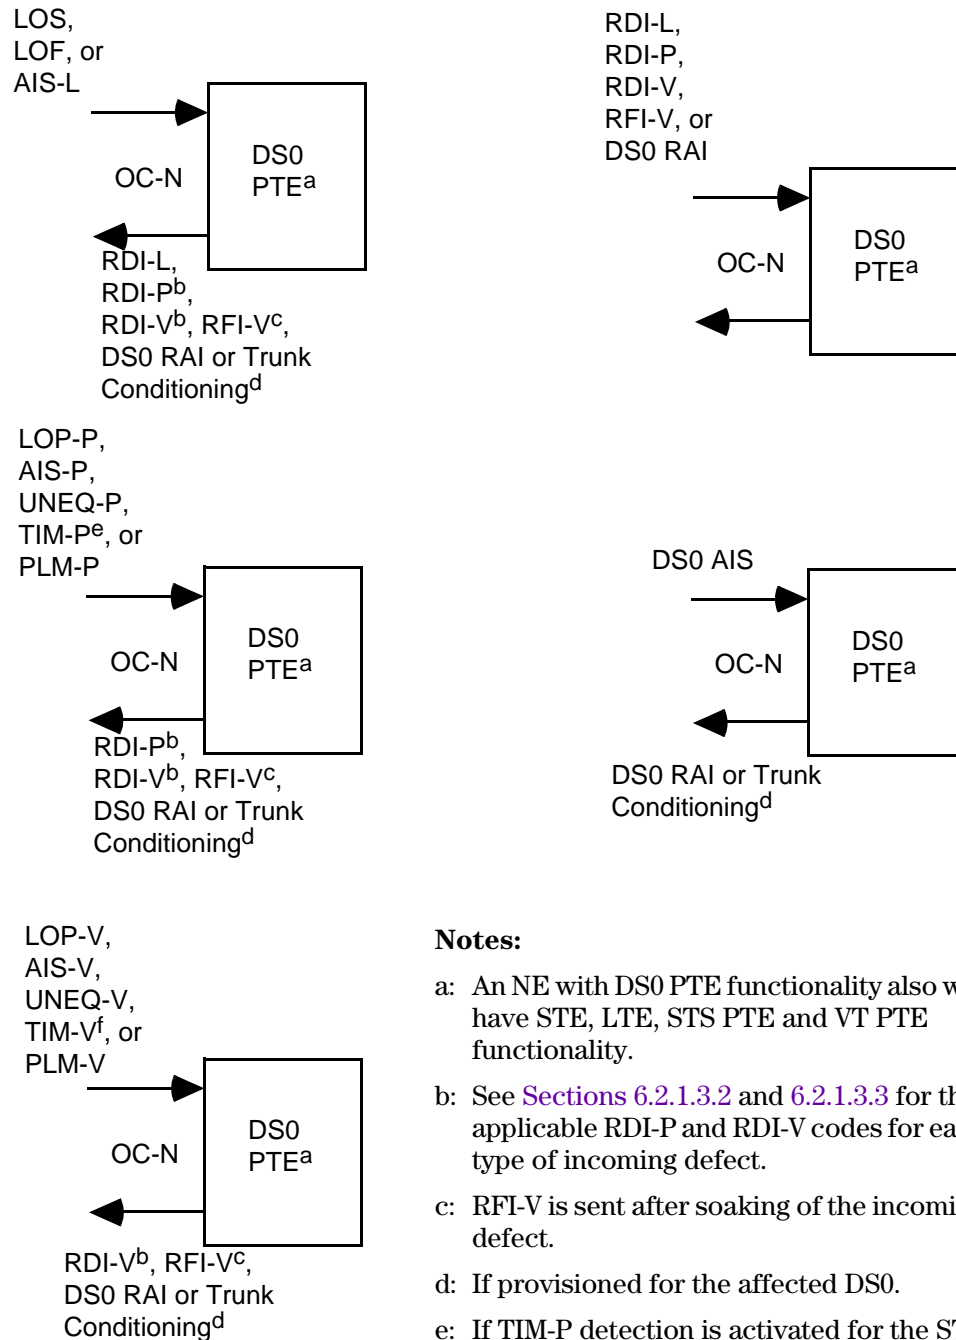
Figure 6-11 VT PTE with DS0 Rearrangement Functions (Byte-Synchronous Mapping for DS1 into VT1.5) Maintenance Signals and Trunk Conditioning, OC-N and DS1 Interfaces (Continued)



Notes:

- a: An NE with VT PTE functionality also will have STE, LTE and STS PTE functionality.
- b: Not if DS0 AIS is sent downstream (i.e., only if trunk conditioning is performed downstream).
- c: If provisioned for the affected DS0 in this direction.

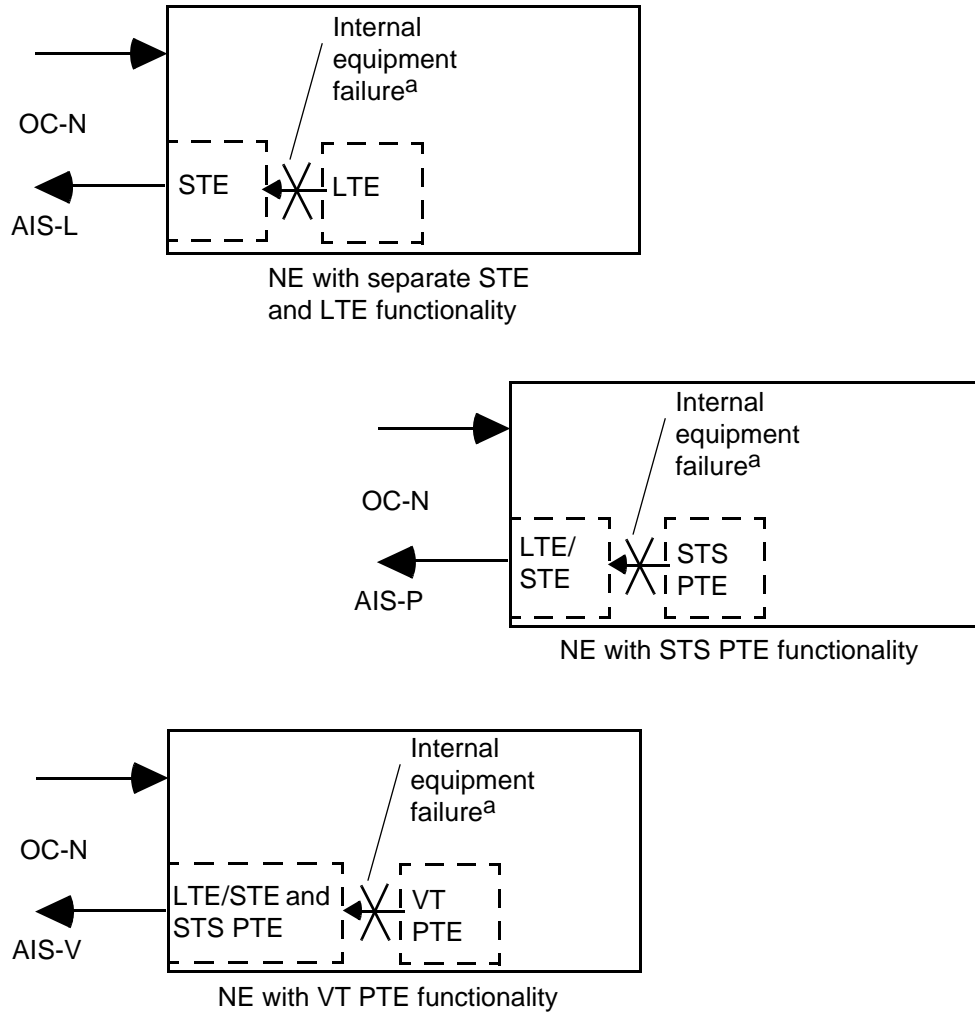
Figure 6-12 DS0 PTE (Byte-Synchronous Mapping for DS1 into VT1.5) Maintenance Signals and Trunk Conditioning



Notes:

- a: An NE with DS0 PTE functionality also will have STE, LTE, STS PTE and VT PTE functionality.
- b: See [Sections 6.2.1.3.2](#) and [6.2.1.3.3](#) for the applicable RDI-P and RDI-V codes for each type of incoming defect.
- c: RFI-V is sent after soaking of the incoming defect.
- d: If provisioned for the affected DS0.
- e: If TIM-P detection is activated for the STS path.
- f: If TIM-V detection is activated for the VT path.

Figure 6-13 SONET Maintenance Signals for Internal Equipment Failures



Note:

- a: In each case, the AIS is inserted when equipment that is originating a signal detects that the circuitry supporting provisioned higher layer origination functions has failed or been removed without having been “unprovisioned”, and it continues until standby circuitry, if available, is switched in or the failure is cleared (see [Section 6.2.1.2](#)).

Figure 6-14 Alarm Timing Requirements for Directly Detected Defects and Failures

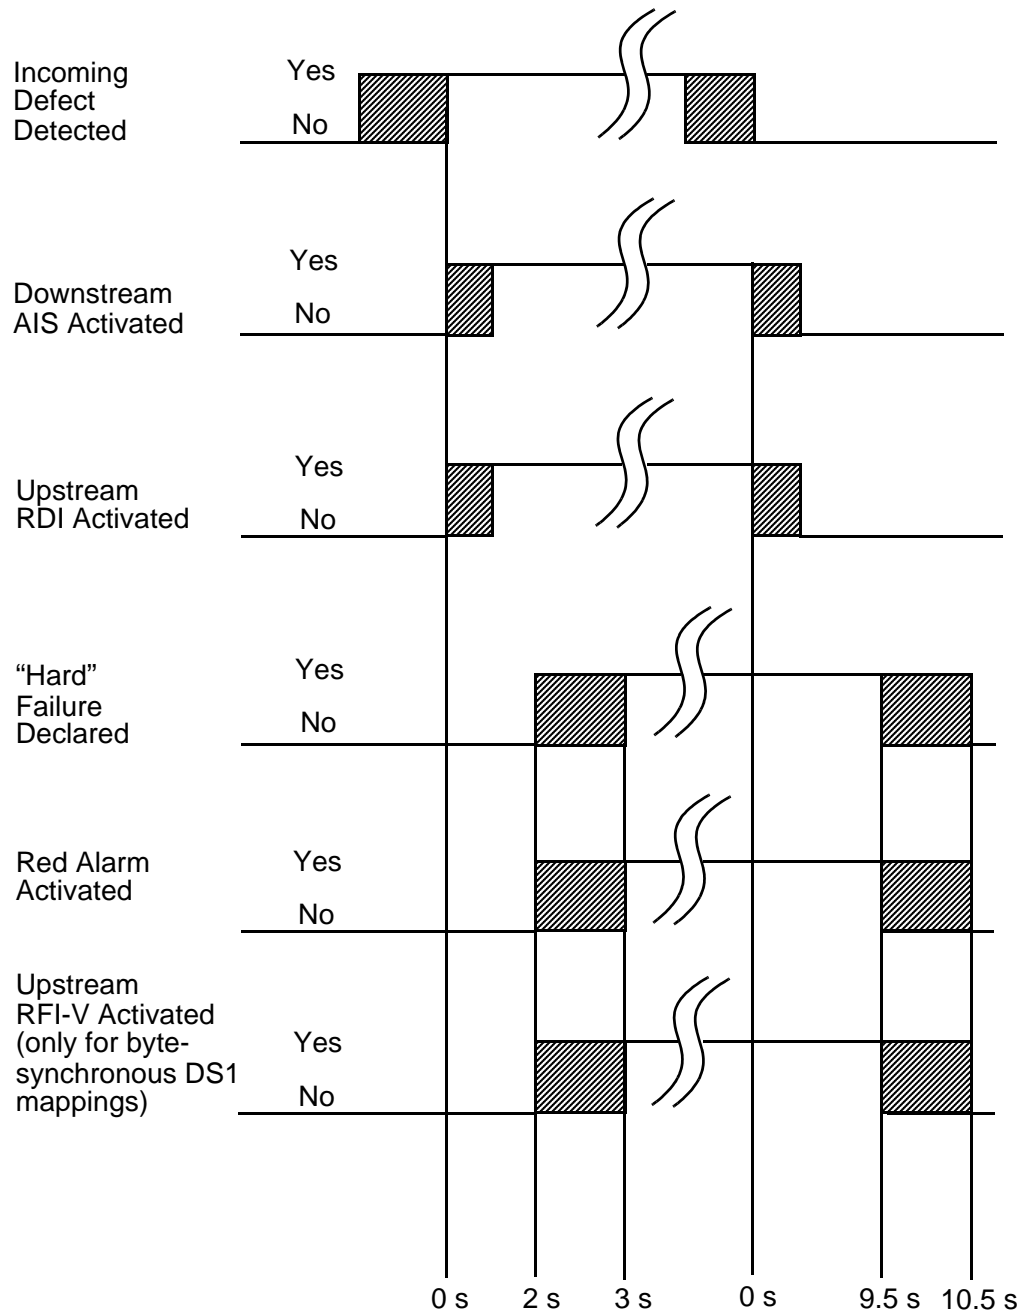


Figure 6-15 AIS Timing Requirements

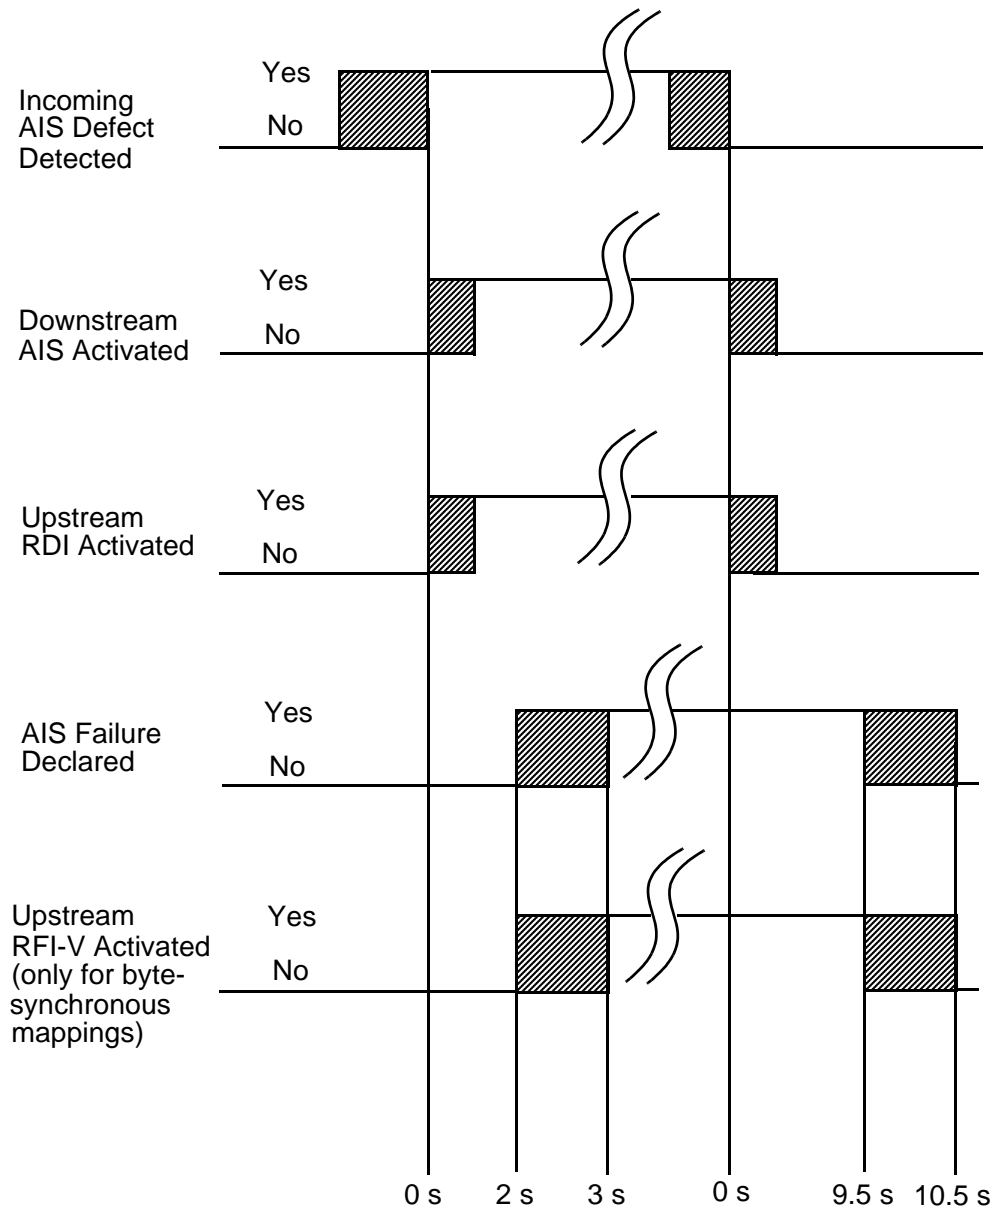


Figure 6-16 Derived RFI Timing Requirements

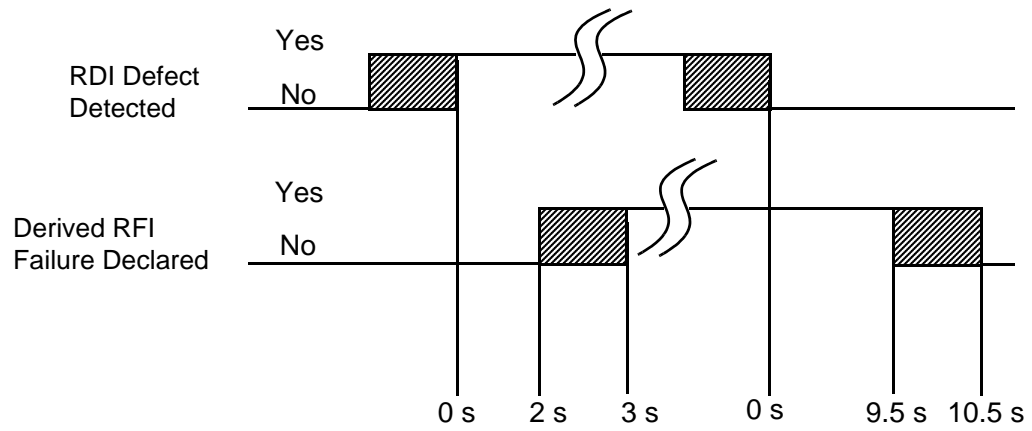
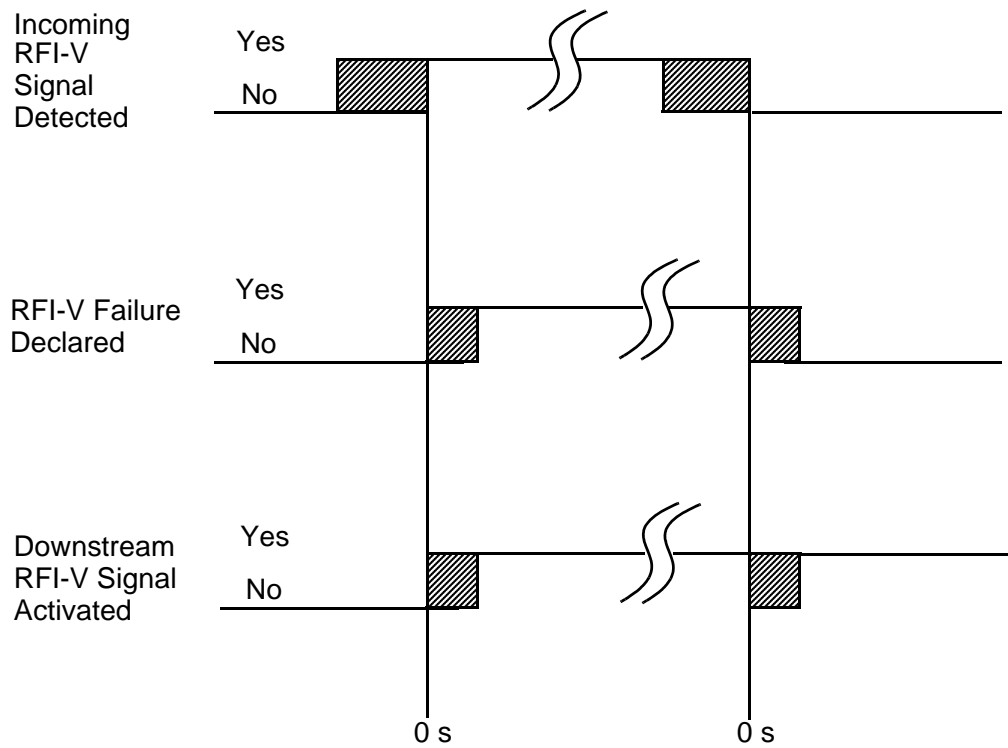


Figure 6-17 RFI-V Signal Timing Requirements (for Byte-Synchronous DS1 Mapping)



6.2.1.8 Control of Alarm Processing

GR-474-CORE contains criteria related to alarm processing that are applicable to transport and switching NEs in general. Although there are some differences in terminology and assumptions about the functionality of the NEs, many of the criteria in GR-474-CORE can be directly applied to SONET NEs. This section reviews and reiterates the criteria contained in GR-474-CORE that are applicable to transport NEs, and indicates where there are specific differences for SONET NEs.

In GR-474-CORE, the term “trouble notifications” is used to refer to the various ways that the user is alerted to the existence, location and severity of a trouble. Trouble notifications include output messages (e.g., to an OS), visual indications at the NE, and audible/visible Central Office (CO) frame indications. The following sections are primarily concerned with the messages that an NE autonomously sends to an OS when it declares a failure, and the information that it sends in response to a user’s request for the current condition of the NE.

Another document that contains information and criteria relevant to the area of alarm processing is GR-1093-CORE, *Generic State Requirements for Network Elements (NEs)*, which contains information about the possible service states that may be supported by various NEs (including SONET NEs). Among the states discussed in that document is an “Automatic In-Service Provisioning” (AISP) state. That state is considered to be an Alarm Reporting Control feature and can be used during pre-service provisioning and testing of any type of termination point (e.g., DS1, OC-3). In particular, when a termination point is in the AISP state the NE is required (according to GR-1093-CORE) to perform most of the functions normally associated with that point (e.g., defect detection and termination, failure declaration and clearing, accumulation of PM data). However, it does not autonomously generate alarm and TCA messages for that termination point (for transmission to an OS). In addition, if a good signal is continuously received at the termination point for some provisionable time period, that point is required to automatically transition to the In-Service state.

6.2.1.8.1 Alarm Level Designations

According to GR-474-CORE, a trouble notification message must contain certain information, including (among other items) the type of trouble, the time the event occurred (e.g., the time the failure was declared), the signal level affected, a designation as Service-Affecting (SA) or Non-Service-Affecting (NSA), a designation as alarmed or Not Alarmed (NA), and (for alarmed troubles) a designation as Critical (CR), Major (MJ) or Minor (MN). GR-474-CORE also provides defaults and guidelines regarding whether a trouble should be alarmed or NA, and whether alarmed troubles should be designated as CR, MJ or MN. In general, failures resulting from problems on incoming signals are supposed to be alarmed, while failures resulting from incoming maintenance signals are supposed to be NA or Not Reported. In addition, GR-474-CORE indicates that a severe SA failure that requires immediate corrective action (independent of the time of day) is supposed to be designated as CR, a serious SA failure that requires immediate attention is supposed to be designated as MJ, and an alarmed NSA failure is supposed to be designated as MN. In general, all of these criteria are applicable to SONET NEs; however, some additional criteria and explanations are needed.

For the user, the designation of a trouble as either SA or NSA is a very important issue. In general, an SA failure on an incoming signal is a non-AIS failure where the corresponding defect causes the detecting NE to generate AIS (or apply trunk conditioning) downstream, and the generated AIS appears at the NE's external interfaces after any available protection switching has been completed. For example, at an NE using 1+1 linear APS, an STS LOP failure declared on an STS-1 received on the selected line would be an SA failure, while an OC-N LOS would be considered NSA if the traffic is restored by a successful protection switch. Note that incoming AIS failures are generally not considered to be SA because the (upstream) NE that inserted the AIS should have already reported the root-cause incoming signal (or possibly equipment) problem as an SA failure. In addition, GR-474-CORE indicates that a declared AIS failure should either cause an NA notification or should be Not Reported, depending on the criteria in the applicable technology-specific document (such as this document).

R6-310 [621] AIS and RFI failures shall have a default setting of Not Reported.

R6-311 [622v2] The Far-End Protection Line failure shall have a default setting of a MN alarm.

6.2.1.8.2 *Single Failure/Single Message*

GR-474-CORE also indicates that any single failure (or root-cause incoming signal problem) must result in only one alarmed output message. This has a significant effect on SONET NEs, which can terminate multiple layers of the incoming signal, and for which multiple defects and failures are defined at each layer. The priorities (e.g., for alarm reporting purposes) of the various traffic-related failures that a SONET NE can declare based on its incoming SONET signals are discussed in this section.

In addition to traffic-related failures, a SONET NE may also monitor an incoming SONET signal for other types of failures. For example, a SONET signal could be carrying the APS channel, it could be used as a synchronization source, or it could be carrying operations communications via the Section DCC bytes. In general, the user may want to be alerted (via an alarmed output message) when one of those functions fails, even if that failure is a direct result of an incoming signal problem that causes a traffic-related alarmed output message (or an incoming maintenance signal).¹⁴ Therefore, non-traffic-related failures may be exceptions to the "single failure/single message" criteria, and are not discussed in this section.

For the purposes of alarm reporting, traffic-related failures can be divided into near-end failures (i.e., LOS, LOF, LOP, AIS, UNEQ, TIM and PLM failures), and far-end failures (i.e., RFI failures). In general, a near-end defect detected at a lower layer disrupts the capability of the equipment terminating higher layers to monitor

14. Note however, that to avoid the declaration of unnecessary alarms when the LTE terminating the protection line in a linear APS system detects an AIS-L (or lower layer) defect, the APS-related defects are defined such that they will not be detected when an AIS-L defect is present.

for both near-end and far-end defects, unless a successful protection switch restores the signal to the higher layer. For example, an LOF defect detected by an NE's STE disrupts the LTE's ability to monitor for an externally generated AIS-L, RDI-L and (if the LTE processes STS pointers) LOP-P defects, but does not disrupt the STS PTE's ability to monitor for STS path layer defects if a line-level protection switch is successfully completed. Therefore, the failure resulting from the lowest-layer defect detected is the one reported to the OS.

For near-end defects that are detected at the same layer, one of the following applies:

- The terminating equipment monitors the same part of the incoming signal to detect and terminate two different defects, and therefore the defects are defined such that the detection of one of them is a condition for terminating the other. In some cases the defects have been defined to result in symmetric operation (e.g., the detection of an UNEQ-P defect terminates a PLM-P defect, and the detection of almost any "non-unequipped" signal label code, including one that results in a PLM-P defect, terminates an UNEQ-P defect), while others may be asymmetric.
- One of the incoming defects disrupts the ability of the terminating equipment to access its POH (i.e., one of the incoming defects is an LOP or AIS defect). In such cases, the defects that the PTE would normally detect or terminate based on the contents of the POH [i.e., UNEQ, TIM, PLM (and RDI)] cannot be detected or terminated.
- The defects are LOS and LOF defects. In that particular case the applicable failure declaration criteria (i.e., **R6-58 [418v3]** and **R6-66 [426v3]**) and corresponding standards (e.g., ANSI T1.231.04) have been written so only an LOS failure should be declared when a root-cause problem such as a fiber cut occurs (even though both defects would normally be detected and remain present).

Based on the preceding discussion, an NE that uses internal SONET signals between layers, that has an alarm level of Not Reported for AIS failures, and that meets the existing criteria on detecting defects and declaring failures would be expected to automatically generate a single alarmed (traffic-related) output message for almost any single incoming signal problem. However, for other cases (e.g., a TIM-P or PLM-P failure that is declared in response to the same incoming signal problem as an UNEQ-P or TIM-P failure, or cases where AIS failures are set to be alarmed) the following requirement is necessary.

-
- R6-312 [626v3]** A SONET NE shall not autonomously report a near-end (or far-end) failure that is the result of the same root-cause incoming signal problem or maintenance signal as a near-end failure reported by the NE, per the hierarchy in [Table 6-6](#). In addition, the SONET NE shall not autonomously report a near-end failure declared for equipment (e.g., STS PTE) that has been provisioned to a service state in which autonomous reporting is inhibited (see [Section 6.2.1.8](#)).
-

Table 6-6 Hierarchy of Near-End Failures

Priority	Failure
Highest	LOS
	LOF
	AIS-L
	AIS-P ^a
	LOP-P ^b
	UNEQ-P
	TIM-P
	PLM-P
	AIS-V ^a
	LOP-V ^b
	UNEQ-V
	TIM-V
	PLM-V
Lowest	DSn AIS (if Reported for outgoing DSn signals)

Notes:

- a. Although it is not defined as a defect/failure, all-ones STS pointer relay is also higher priority than LOP-P. Similarly, all-ones VT pointer relay is higher priority than LOP-V.
- b. LOP-P and all of the entries above it are also higher priority than the far-end failure RFI-P, which does not affect the detection of any of the near-end failures. Similarly, LOP-V and all of the entries above it are higher priority than RFI-V.

For far-end defects (i.e., RDI defects), the detection of a defect at a particular layer does not disrupt the capability to detect defects (either near-end or far-end) at any higher layer. Therefore, it is possible for an NE to detect multiple RDI defects and declare multiple RFI failures simultaneously. In addition, the detection of a single near-end defect by the far-end NE could cause that NE to generate multiple upstream RDI signals. If the near-end NE is set to report RFI failures, then that single root-cause incoming signal problem or maintenance signal (at the far-end) could result in multiple failure messages being sent to the OS by the near-end NE. To avoid this, the required default setting for RFI failures is Not Reported. In addition, for NEs that are set to report RFI failures, the following requirement applies.

R6-313 [629v3] An NE that is set to report RFI failures shall not autonomously report an RFI failure that is apparently caused by the same root-cause incoming signal problem or maintenance signal at the far-end that caused the NE to concurrently declare (and report) a higher-priority RFI failure, per the hierarchy in [Table 6-7](#). In addition, the SONET NE shall not autonomously report a far-end failure declared for

equipment (e.g., STS PTE) that has been provisioned to a service state in which autonomous reporting is inhibited (see [Section 6.2.1.8](#)).

Table 6-7 Hierarchy of Far-End Failures

Priority	Failure
Highest	RFI-L
	RFI-P
Lowest	RFI-V

Note:

- a. Note that Issue 1 of this document also included the Far-End Protection Line failure in the hierarchy of far-end failures, between RFI-L and RFI-P. That failure was removed because it is considered an independent, APS-related failure; however, an NE may include it in the hierarchy and still conform to the requirement.

For example, an NE that is using linear APS and that detects RDI-L defects on both the working and the protection lines from an adjacent NE, along with RDI-P and RDI-V defects on all of the terminating STS and VT paths received on those lines would report (if RFI failures are set to be reported) only the RFI-L failures. It would not report all of the lower priority RFI-P and RFI-V failures.

6.2.1.8.3 Independent Failures

[Section 6.2.1.8.2](#) is concerned with the failures reported to an OS due to a single root-cause incoming signal problem or maintenance signal, while this section is concerned with independent failures, as defined below.

- Independent failures - Failures (as defined in [Section 6.2.1](#)) that are declared for a single entity (i.e., a SONET Section, Line, STS path, VT path, or possibly an outgoing DS_n) in response to two or more root-cause incoming signal problems or maintenance signals, possibly requiring separate maintenance actions.

In general, it is assumed that for an NE to determine that two failures are independent, those failures would need to be declared at “separate times”. However no assumptions have been made regarding a definition of “separate times”, as this is viewed as an implementation detail. (One possible definition is that two failures are considered independent if the first failure has already been declared at the point in time when the defect leading to the second failure is detected.)

As an example of the above definition, if a fiber break in the optical path carrying an unprotected OC-3 signal causes an NE to (approximately simultaneously) declare an OC-3 LOS failure, 3 AIS-P failures and 84 AIS-V failures, those failures would not be considered independent (and the criteria in [Section 6.2.1.8.2](#) would apply).

However, if prior to the fiber break, the NE had declared an UNEQ-V failure for one of the VT paths, then that UNEQ-V failure and the subsequent AIS-V failure on that same path would be considered to be independent (and the criteria in this section would apply).

Issue 1 of this document contained several requirements that indicated that certain failures were supposed to be cleared when certain other failures were declared. It also contained requirements indicating that certain defects were supposed to be terminated when certain higher-priority defects were detected, even though the higher-priority defects disrupted the signal such that the bits or bytes that the NE needed to monitor to terminate the lower-priority defect could not be accessed. Those requirements covered only a few of the possible combinations of defects and failures, and could cause confusion for combinations that were not covered. In addition, failures that are declared at separate times are generally caused by separate problems, which need to be individually addressed. Therefore, those requirements were removed and **R6-314 [988]** (which was reviewed in detail in GR-253-ILR, Issue ID 253-121) was added. While the former requirements are no longer considered necessary, an NE that meets those former requirements should not be considered nonconforming to the current criteria. Therefore, the following are allowed exceptions to **R6-314 [988]** and the defect termination criteria in [Sections 6.2.1.1.8](#) and [6.2.1.3](#):

- An NE may clear an existing LOF failure when an LOS failure is declared
- An NE may terminate an UNEQ-P defect when an AIS-P or LOP-P defect is detected
- An NE may clear an UNEQ-P failure when an AIS-P or LOP-P failure is declared
- An NE may terminate a PLM-P defect when an AIS-P or LOP-P defect is detected
- An NE may clear a PLM-P failure when an AIS-P or LOP-P failure is declared
- An NE may terminate an RDI-P defect when an AIS-P defect is detected
- An NE may terminate an UNEQ-V defect when an AIS-V or LOP-V defect is detected
- An NE may clear an UNEQ-V failure when an AIS-V or LOP-V failure is declared
- An NE may terminate a PLM-V defect when an AIS-V or LOP-V defect is detected
- An NE may clear an PLM-V failure when an AIS-V or LOP-V failure is declared
- An NE may terminate an RDI-V defect when an AIS-V defect is detected.

R6-314 [988] The declaration of a “new” failure by a SONET NE shall not automatically cause the NE to clear any previously declared, independent failures.

In some cases, an NE that meets the preceding requirement would still be expected to clear an existing failure after a new higher-priority (or lower-layer) failure is declared; however, the existing criteria on defect detection and termination, and failure declaration and clearing should be sufficient. In other cases, the lower-priority failure would be expected to remain in place because the defect that caused it cannot be terminated while the defect associated with the higher-priority failure is present.

6.2.1.8.4 Retrieval of an NE's Condition

While Sections 6.2.1.8.2 and 6.2.1.8.3 were concerned with the autonomous messages that an NE sends to an OS, this section discusses the response of the NE to requests from the user for a report on the current condition of the NE.

In general, when an NE is responding to a user request to report all of the failures at the NE, it needs to report all of the independent failures (i.e., the same failures that it has autonomously reported). If the NE were to also report lower-priority failures that were caused by the same root-cause incoming signal problem or maintenance signal as a higher-priority failure, then the number of failures reported to the user could easily become excessive. On the other hand, if the user requests only the failures at a particular layer or for a particular entity (e.g., for the PTE supporting a particular path), then it would be appropriate for the NE to report all of the failures that it has declared for that layer or entity.

R6-315 [625] A SONET NE shall provide the capability to report to the user, on demand, the current failure indications (i.e., the current condition of the NE).

R6-316 [627v2] A SONET NE shall not report a failure that is the result of the same root-cause incoming signal problem or maintenance signal as another failure reported by the NE (per the hierarchy in Table 6-6) in response to a request to report all failures at the NE.

R6-317 [628v2] A SONET NE shall report all failures, including each failure that is the result of the same root-cause incoming signal problem or maintenance signal as another failure reported by the NE, in response to a request to report all failures at a given SONET layer, or for a given entity.

For example, if an AIS-L failure is declared concurrently with an AIS-P failure, indications for both failures would be kept at the NE, even though neither failure or (if the NE was provisioned to report AIS failures) only the AIS-L failure would be autonomously reported. In either case, if the NE were subsequently requested by the user to report all failures, it would report the AIS-L failure (along with any other independent failures). However, if a request was received to report STS path layer failures, it would report the AIS-P failure (along with any other STS path layer failures).

6.2.1.8.5 Provisioning of Alarm Levels

GR-474-CORE also contains an objective for the designation of a trouble as alarmed or NA to be provisionable, and a requirement that the user be able to inhibit any NA notification (e.g., by changing the designation to Not Reported). In SONET, the capability to provision certain failures as alarmed or NA is changed from an objective to the following requirement.

R6-318 [620] SONET equipment shall provide the capability of setting any AIS (including DS_n AIS), RFI (including DS_n RAI) or Far-End Protection Line failure as either Reported or Not Reported, and if Reported, as either alarmed or Not Alarmed. The

settings shall be provisionable on a per-layer, per-entity basis (e.g., for the line layer, the settings shall be provisionable on a per-line basis).

R6-319 [623v2] A SONET NE shall provide the capability of reporting (on demand) all provisionable attributes.

GR-474-CORE contains requirements for controlling local alarms (e.g., alarm cutoff) and remote alarms.

6.2.1.8.6 Clear Messages

Finally, GR-474-CORE indicates that a clear message must be generated when an alarmed trouble clears. However, it does not indicate that a clear message also needs to be generated when certain NA troubles that were reported to the OS are cleared [i.e., when NA troubles that were reported as standing conditions (as opposed to transient conditions) are cleared]. Therefore, the requirement on clearing failures is expanded to the following for SONET NEs.

R6-320 [632] A SONET NE shall individually clear (and send a clear message to the OS) any failure that is individually reported to an OS.

6.2.1.8.7 Non-Intrusive Detection of Defects and Declaration of Failures

In some applications, equipment within an NE may need to detect defects that, based on the criteria in Sections 6.2.1.1.1 through 6.2.1.6, are normally detected only by equipment with other (higher-level) functionality. For example, STS PTE that generates PDI-P for an outgoing VT-structured STS-1 must detect AIS-V defects even though it does not terminate the VT path. Based on wording of the AIS-V detection criteria in Section 6.2.1.2.3, those criteria are applicable only to VT PTE. Therefore, the PDI-P criteria in Section 6.2.1.4.1 indicate that the STS PTE “shall detect and terminate AIS-V defects as if it were VT PTE”. Similar statements are made in other criteria, and in general the equipment that detects and terminates the defects does so non-intrusively. The equipment may take some action based on the detection of the defect (e.g., change the PDI-P code that it is sending), but it does not alter the monitored signal by generating (rather than passing) AIS downstream.

When an NE non-intrusively monitors a signal, it is generally not appropriate for it to autonomously report (to an OS) any failures that it could declare based on the defects detected on that signal.¹⁵ The reason for this is that when the signal is terminated at a downstream NE, that NE will generally detect the same defect, declare the same failure, and report that failure to an OS (unless it is a Not Reported

15. One exception to this may be defects that an NE uses for path protection switching purposes. See the appropriate application-specific GRs (e.g., GR-1230-CORE, GR-1400-CORE) for any additional criteria related to the reports generated when an NE detects such a defect and performs a path protection switch.

failure). If a number of NEs non-intrusively monitor a signal and report failures on that signal to an OS, the OS could receive a large number of reports for a single failure.

Although in most applications it is not appropriate for an NE that non-intrusively monitors a signal to autonomously report failures related to that signal to an OS, some users may want that capability. In addition, it can be useful for the NE to declare and clear those failures so that they can be retrieved by the user (e.g., for troubleshooting purposes). Therefore, the following criteria are applicable.

R6-321 [695v2] An NE that is non-intrusively monitoring a signal shall declare and clear failures for that signal as if it were terminating the signal. Upon declaring or clearing a failure, the NE shall set or clear the appropriate failure indication for that signal.

O6-322 [703v3] As a default, an NE that is non-intrusively monitoring a signal should not autonomously report to an OS the declaration or clearing of a failure on that signal (i.e., the default or “fixed” setting for the failures should be Not Reported).

Note that although **O6-322 [703v3]** refers to a “default” setting for failures declared for non-intrusively monitored signals, this is not meant to imply that the levels for those failures are required to be provisionable. Instead, it is intended to indicate that an NE is allowed to support provisionable alarm levels for those failures, and that if an NE does happen to support them, then the default level should be Not Reported. Similarly, if the NE does not happen to support provisionable alarm levels for those failures, then the “fixed” (or only supported) level should be Not Reported.

6.2.2 Performance Monitoring

PM refers to the in-service, non-intrusive monitoring of transmission quality. A SONET NE is required to support PM according to the layer of functionality that it provides, and in SONET the capability is provided to accumulate PM data based on overhead bits (e.g., BIP-Ns) at the Section, Line, STS path and VT path layers. In addition, PM data is available at the SONET Physical layer using physical parameters (rather than overhead bits). For SONET NEs that interface to the previously existing digital network (e.g., that support DS1 interfaces), the accumulation of certain DS_n PM parameters may also be appropriate.

R6-323 [633] Except as specifically noted, SONET NEs shall meet the general PM requirements in GR-820-CORE.

GR-820-CORE contains generic PM strategies, discusses various types of PM registers (e.g., current period, previous period, recent period and threshold registers), and defines PM parameters for DS1 and DS3 signals. In general, the generic strategies discussed in GR-820-CORE are directly applicable for monitoring SONET signals, and therefore are not repeated here. However, additional PM parameters are needed to support SONET PM. The following sections describe the PM parameters defined for each SONET layer, and identify the accumulation and thresholding requirements for SONET NEs. At certain SONET layers, both near-end

and far-end performance can be monitored. For those layers, both near-end and far-end PM parameter definitions, accumulation requirements and thresholding requirements are provided.

In general, an NE accumulates various PM parameters based on performance “primitives” that it detects in the incoming digital bit stream. Primitives can be either “anomalies” or defects. An *anomaly* is defined to be a discrepancy between the actual and desired characteristics of an item (e.g., an error in a received BIP-8 code), and may or may not affect the ability of the item to perform a function. A *defect* is defined to be a limited interruption in the ability of an item to perform a required function. The persistence of a defect results in a *failure*, which is defined to be the termination of the ability of an item to perform a required function (see [Sections 5.5](#) and [6.2.1](#) for descriptions of the various defects and failures). Functionally, PM is performed at each layer, independent of the other layers. However as [Figure 6-2](#) illustrates, part of the functional model assumes that layers pass maintenance signals to higher layers. For example, a defect (e.g., LOF) that occurs at the Section layer causes an AIS-L signal to (functionally) be passed up to the Line layer, which causes the AIS-P signal to be passed up to the STS path layer, etc. Thus, an AIS defect can be detected at a particular layer either by receiving the appropriate AIS on the incoming signal, or by receiving it from a lower layer. Therefore, defects and failures that occur at a lower layer affect the PM parameters at higher layers, and the definitions of the assorted PM parameters were written with that consideration.

Thresholds are defined for most of the PM parameters supported by SONET NEs, and are used to detect when transmission degradations have reached unacceptable levels. When a threshold for a non-Physical layer parameter is reached or exceeded and there is no potential for the parameter to be adjusted to be less than the threshold because of entry into (or exit from) unavailable time,¹⁶ a TCA is sent to an OS. Similarly, when the recorded value of a Physical layer parameter is outside of its acceptable range, an Out-of-Range Alert (ORA) is sent to an OS. A TCA or ORA is reported as a transient condition (i.e., it is not cleared at some later time), and only one TCA or ORA can generally be generated based on the data accumulated or recorded in any particular current period register. However, it is possible for another TCA or ORA to be generated based on the data accumulated or recorded in the next accumulation period (e.g., if the threshold is again reached or exceeded).

Figures [6-18](#), [6-19](#) and [6-20](#) illustrate (in the form of flowcharts) some of the processes involved in the accumulation of non-Physical layer PM data. Each process runs concurrently with the others, although some are triggered by the same conditions. These figures are intended to clarify the requirements in the following sections, and therefore they should be viewed only as a logical representation of actions that could be taken to meet the requirements. They are not intended to imply or unnecessarily constrain an implementation. In fact, some steps are intentionally vague or omitted, because they are considered implementation details.

In addition to the types of registers discussed in GR-820-CORE, several other types of registers are mentioned in Figures [6-18](#), [6-19](#) and [6-20](#). Note that in all cases, “register” simply means a location where data is stored, and should not be taken to

16. See [Sections 6.2.2.4](#) through [6.2.2.6](#) for the definitions of unavailable time at the SONET Line, STS path and VT path layers, and [Section 6.2.2.8](#) for criteria on the accumulation of PM parameters during unavailable time (including entry into and exit from unavailable time).

imply any type of architectural design. In the flowcharts, the term “current second register” refers to a register that contains only the data for the current second. At the end of a second, the data in the current second register is normally moved to the “current period” registers, unless some other action is warranted. Each parameter at the Section, Line, STS path and VT path layers has both a current second register and current period registers. Also, “negative adjustment registers” and “positive adjustment registers” (referred to jointly as “inhibition registers” in Issue 1 of this document) refer to registers that are used to subtract from or add to a current period register (and in some cases, possibly the previous period register) when unavailable time is entered or exited. Each Coding Violation (CV), Errored Second (ES), Severely Errored Second (SES) and Unavailable Second (UAS) parameter at the Line, STS path and VT path layers has a negative adjustment register, a positive adjustment register, or both.

The following list describes the purpose of each of the flowcharts (which are labeled A through J):

- A: Illustrates how a current second register is incremented
- B: Illustrates how a current period register is accumulated
- C: Illustrates conditions under which a TCA message is sent
- D: Illustrates how PM registers behave as a push-down stack
- E and F: Illustrate cases where the negative and positive adjustment registers are incremented
- G and H: Illustrate cases when the negative and positive adjustment registers are cleared
- I: Illustrates how current period registers are adjusted on entering unavailable time
- J: Illustrates how current period registers are adjusted on exiting unavailable time

Note that not all of the flowcharts are applicable for all parameters, and that the flowcharts shown do not form a complete set. For example, flowchart C does not cover several cases where a TCA may need to be sent during the first 10 seconds of a period based on the data accumulated for the previous period (see [Section 6.2.2.1](#)). In addition, flowcharts C and E through J are not applicable for UAS parameters (for which TCAs would normally be sent during unavailable time rather than available time, and the roles of the negative and positive adjustment registers are reversed), flowcharts E through J are not applicable for the parameters at the Section layer (at which no UAS parameter has been defined), and flowcharts E, G and I are not applicable for Line or path CV parameters (the accumulation of which is already inhibited during the first 10 seconds of unavailable time because those seconds are SESs).

Figure 6-18 SONET PM Accumulation and Thresholding Model

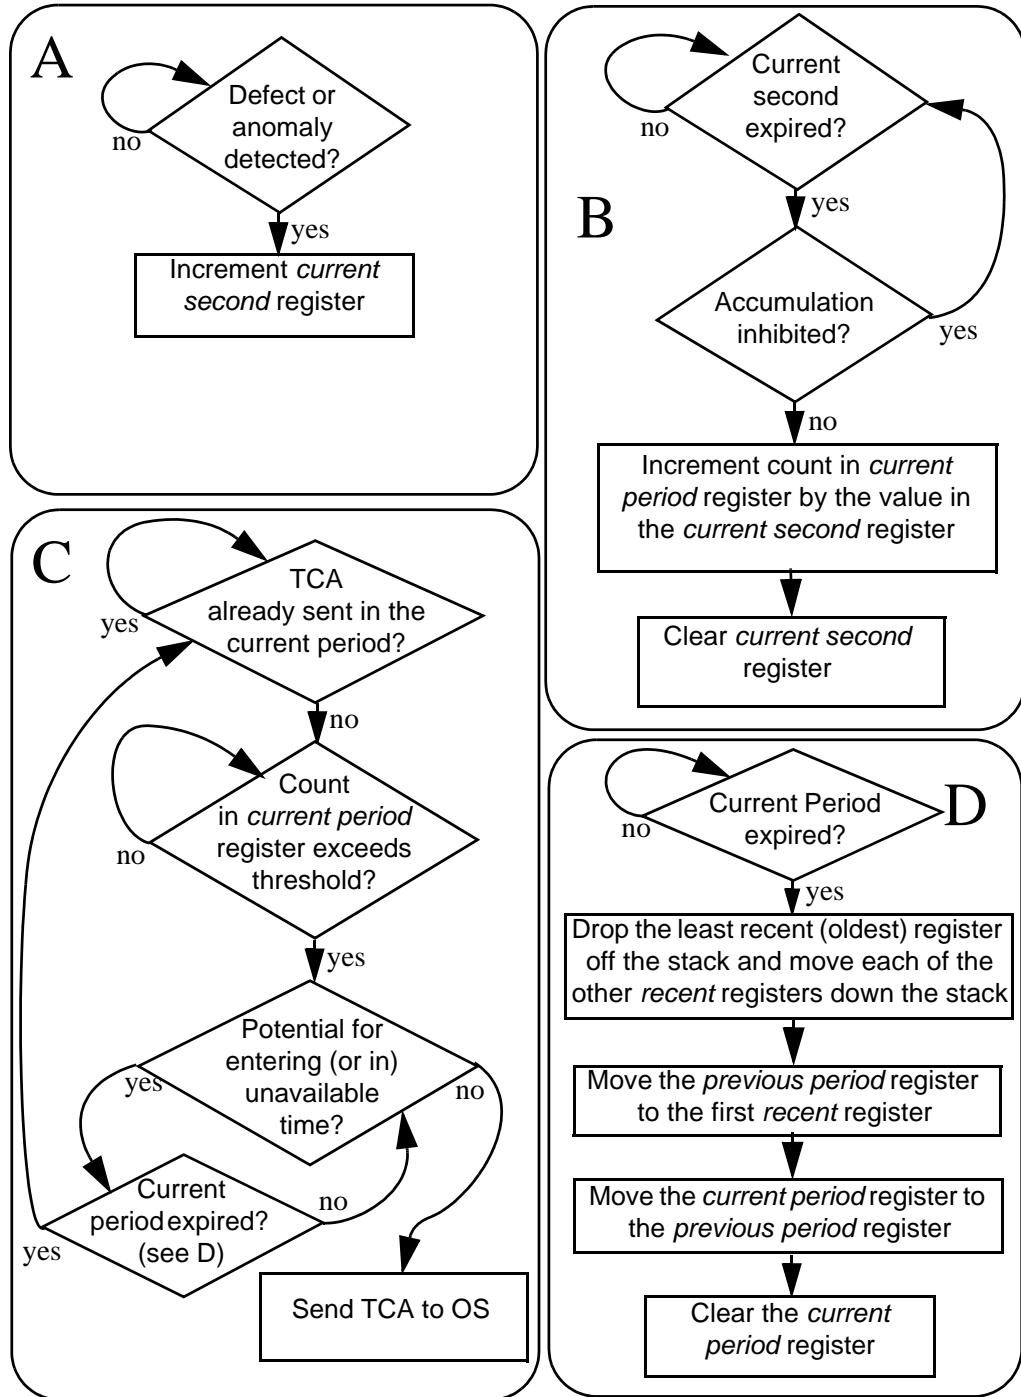


Figure 6-19 SONET PM Accumulation and Thresholding Model (Continued)

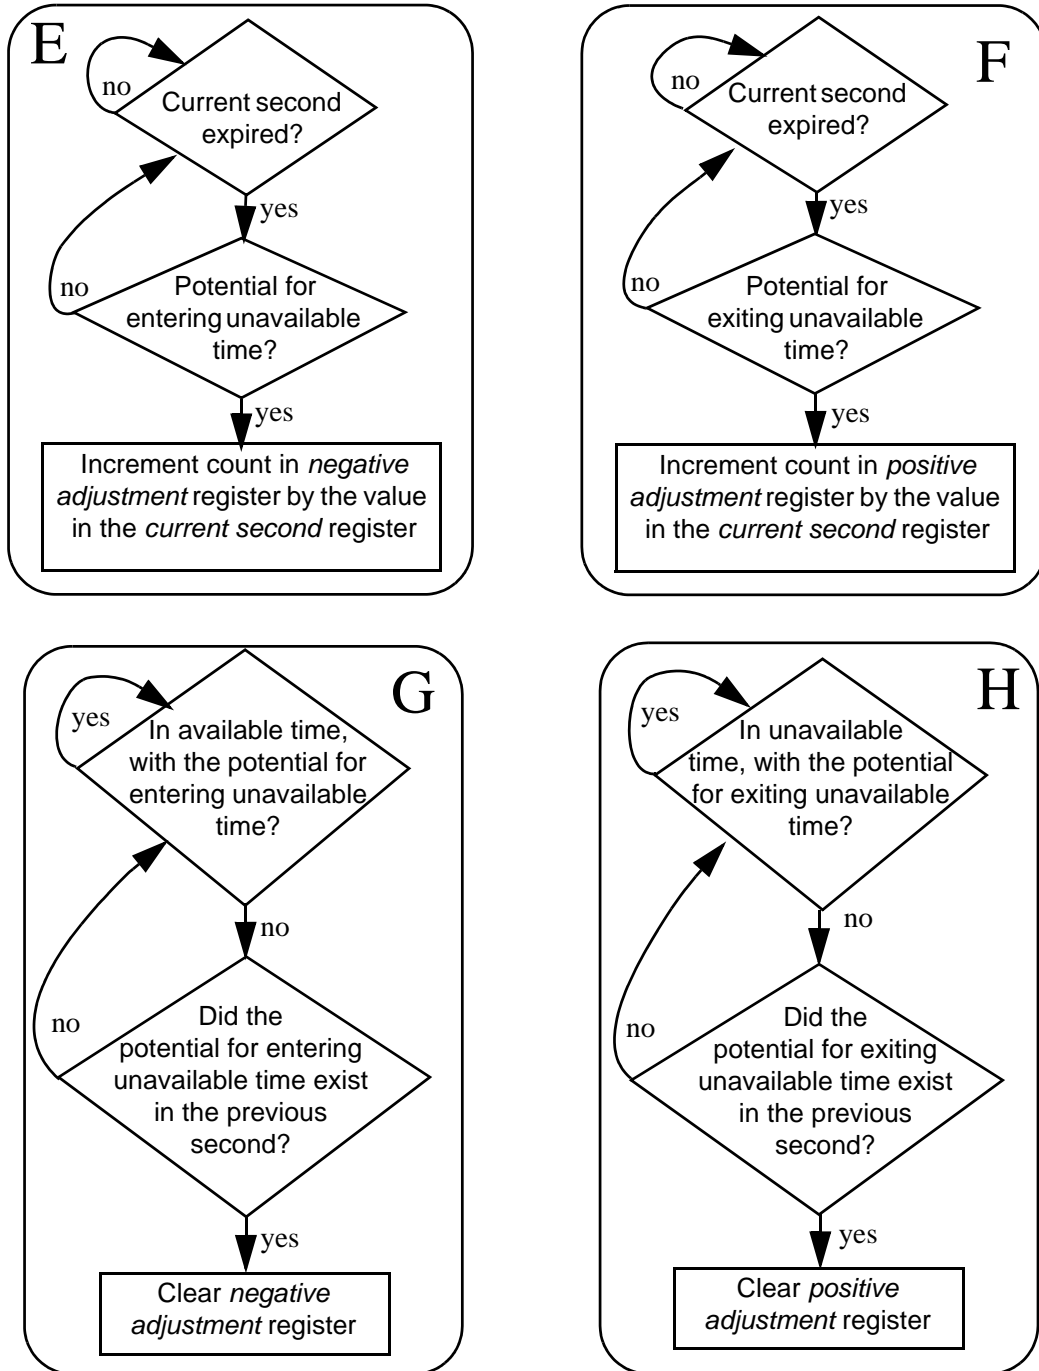
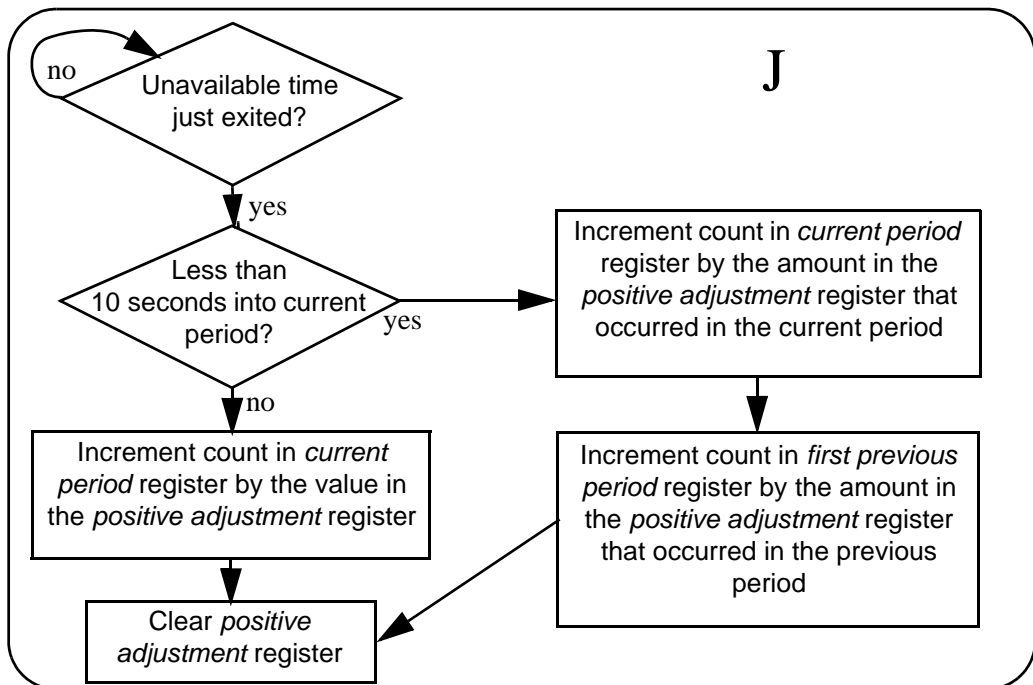
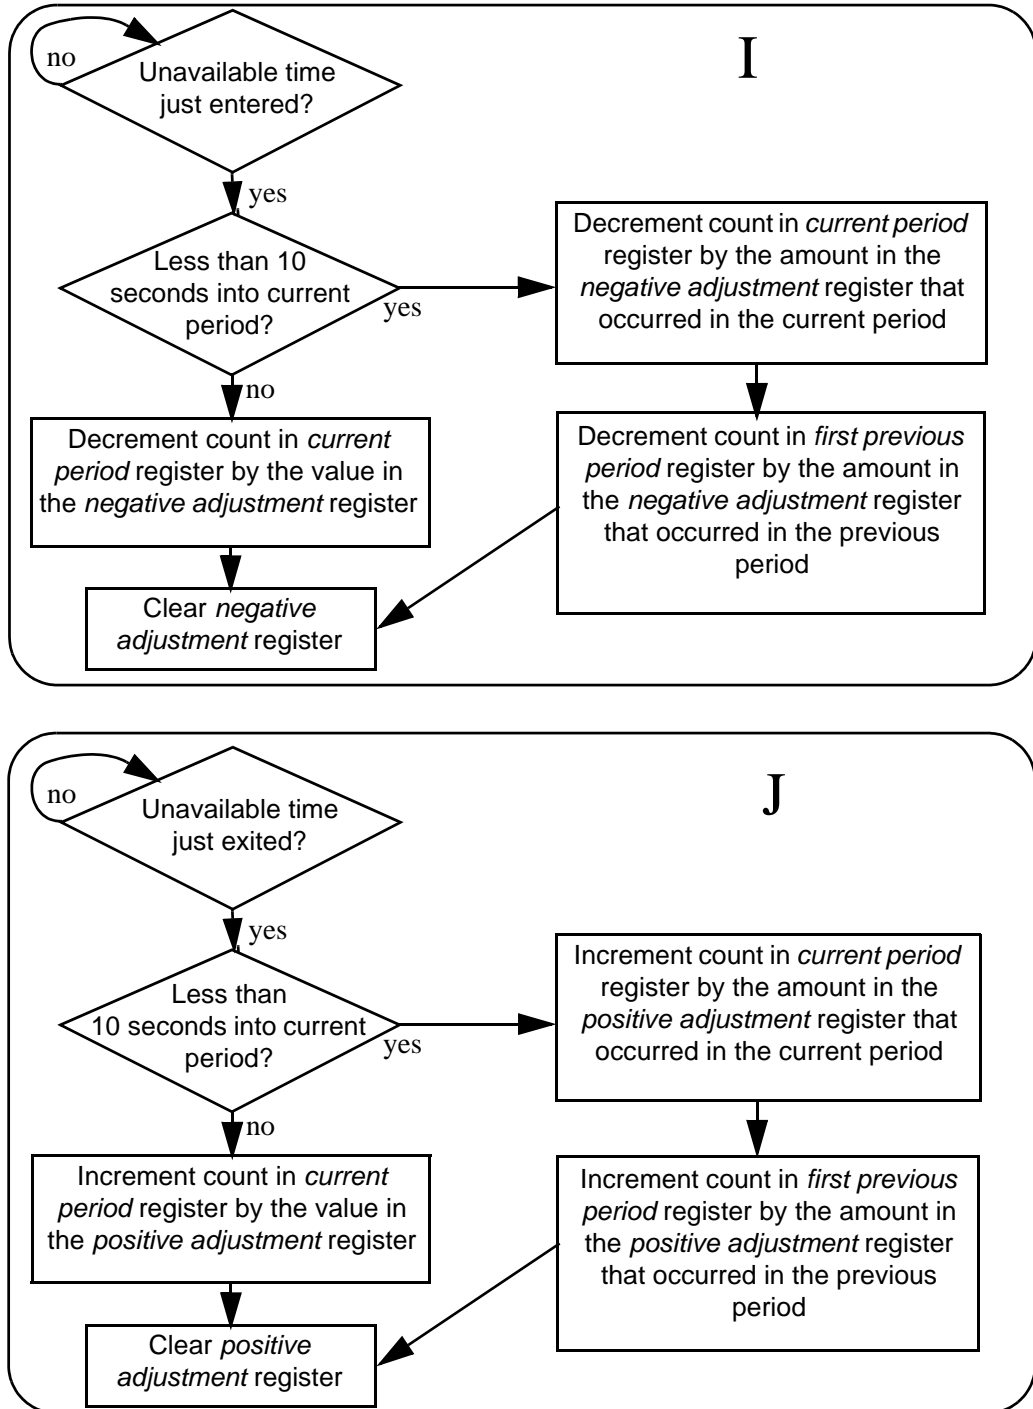


Figure 6-20 SONET PM Accumulation and Thresholding Model (Continued)



6.2.2.1 General Accumulation and Thresholding Criteria

SONET NEs are required to accumulate a variety of PM parameters. Requirements indicating the number and types of registers needed for each parameter are presented below. The wording of the criteria are meant to illustrate the requirements, not to carry any design implications.

R6-324 [634v2] For each PM parameter accumulated for a Physical, Section, Line, STS path or VT path layer entity, a SONET NE shall provide one current 15-minute, one current day, one previous 15-minute, one previous day and 31 recent 15-minute accumulation and storage registers.

R6-325 [639v2] The size of the PM parameter accumulation registers provided by a SONET NE shall be greater than or equal to the minimum accumulation register sizes shown in [Table 6-8](#).

R6-326 [636] A SONET NE shall allow the user to initialize all current 15-minute or current day registers to zero at any time, on an individual entity (e.g., STS path) basis, per direction (e.g., near-end).

R6-327 [637v2] At the end of each 15-minute period, the current 15-minute register, the previous 15-minute register and the 31 recent 15-minute registers shall behave as a push-down stack. The current 15-minute registers shall then automatically be initialized to zero.

In a push-down stack, the contents of the least recent 15-minute register is pushed off the stack, each of the 30 remaining recent 15-minute registers is pushed down the stack, the previous 15-minute register is moved to the most recent 15-minute register, and the current 15-minute register is moved to the previous 15-minute register. The (new) current 15-minute register is then initialized to zero.

R6-328 [989] At the end of each day, the contents of each current day register shall be copied to the corresponding previous day register. The current day registers shall then automatically be initialized to zero.

The invalid-data flag is set to indicate that the data stored in a register or set of registers may be corrupt, because either the period of accumulation is greater or less than the nominal period,¹⁷ or data is missing during an accumulation period (e.g., one or more of the current registers were initialized during the period, the NE was

17. According to GR-820-CORE, an invalid-data flag is supposed to be set if the time of day setting is changed and that change lengthens or shortens the current period by more than 10 seconds. In general, such a flag could be set either immediately after the time is changed [if the difference between the old time and the new time is greater than 10 seconds and, for 15-minute registers, is not an integer multiple of 15 minutes (plus or minus 10 seconds or less)], or delayed until the end of the accumulation period. If the setting of the flag is delayed, then two or more time changes made during the same period could effectively cancel each other out (i.e., result in a total period of accumulation that is within ± 10 seconds of nominal), and in that case the invalid-data flag would not be set.

unable to perform its PM data accumulation functions for some or all of the period, or the accumulation of the far-end parameters was inhibited due to the presence of a near-end defect for some or all of the period).

R6-329 [635v3] A SONET NE shall provide, as a minimum, an invalid-data flag associated with each monitored entity's near-end parameters, and another invalid-data flag associated with its far-end parameters (if applicable).

Note that it is also acceptable for an NE to provide separate invalid-data flags for each register.

R6-330 [990] Invalid-data flags shall be moved with the data to which they apply.

In addition to the accumulation and storage registers required for most PM parameters, a SONET NE must also provide threshold registers according to the following requirements.

R6-331 [638v3] The following apply regarding threshold registers:

- 15-minute and 1-day threshold registers shall be provided for the SONET PM parameters that require thresholding (see [Sections 6.2.2.2 through 6.2.2.7](#))
 - The size of the threshold registers shall be greater than or equal to the minimum threshold register sizes shown in [Table 6-8](#)
 - The value in each threshold register (with the possible exception of certain Physical layer threshold registers, see [Section 6.2.2.2](#)) shall be provisionable
 - Default values for all threshold registers shall be provided and documented by the NE supplier
 - Where equivalent near-end and far-end PM parameters are defined, the default threshold value shall be the same for the near-end and the far-end parameters
 - For PM parameters where default threshold values are shown in [Table 6-8](#), the default values provided by the NE shall be equal to those values.
-

Note that where equivalent near-end and far-end PM parameters are defined, it is sufficient to provide a single threshold register for both. An NE may optionally provide separate threshold registers for the near-end and far-end parameters, along with the capability to provision different threshold values for those parameters.

R6-332 [1081] Unless the applicable equipment (e.g., the optical transmitter) has been provisioned to a service state in which autonomous reporting is inhibited (see [Section 6.2.1.8.2](#)), an ORA shall be sent to an OS when the new "snapshot" value in a Physical layer current 15-minute or current day register is not in the acceptable range for that parameter [i.e., when it is less than or equal to the lower threshold (if defined), or is greater than or equal to the upper threshold].

R6-333 [991v2] Unless the applicable equipment (e.g., the LTE) has been provisioned to a service state in which autonomous reporting is inhibited (see [Section 6.2.1.8.2](#)), a TCA shall be sent to an OS when the value in a non-Physical layer current 15-minute or current day register reaches or exceeds the corresponding threshold and it is not possible for that value to be adjusted to less than the threshold based on entry into or exit from unavailable time.

CR6-334 [1212] The SONET NE may be required to support a failure-based TCA-suppression feature as described in **R6-335 [1213]**.

R6-335 [1213] If the failure-based TCA-suppression feature is supported, the user shall be able to activate it on a per-NE basis and the default shall be that it is deactivated. When activated this feature shall cause suppression:

- Of all CV, ES, SES, Severely Errored Framing Second (SEFS) and UAS 15-minute TCAs associated with any entity/direction combination for which a failure, derived (directly or indirectly) from a defect that is included in the definitions of the corresponding ES and SES parameters, has been declared
- From the time at which the failure is declared through the end of the 15-minute period during which it is cleared.

In addition, this feature shall have no impact on the accumulation of the various PM parameters.

As an example of **R6-335 [1213]**, if failure-based TCA-suppression has been activated at an NE and that NE declares an OC-N LOF failure at 06:03:00, any Section SEFS, CV, ES or SES (SEFS-S, CV-S, ES-S or SES-S), or Line CV, ES, SES or UAS (CV-L, ES-L, SES-L or UAS-L) TCAs associated with the failed signal would need to be suppressed for the remaining 12 minutes of the 06:00:00 15-minute time period.¹⁸ On the other hand, if the traffic being carried on the failed OC-N signal was restored as a result of a linear APS protection switch, then STS path layer TCAs would not be suppressed (at least as a result of the LOF failure). Finally, if the NE cleared the LOF failure at 06:25:00, TCAs related to the same parameters as listed above would need to be suppressed for the entire 15-minute period beginning at 06:15:00.

Based on the definitions of the UAS parameters at the Line and path layers and the requirements on the accumulation of various parameters during unavailable time, an NE must be capable of adjusting the data in certain previous period registers during the first 10 seconds of a new period (due to entry into or exit from unavailable time). In addition, such an adjustment could cause the value in the previous period register to reach or exceed the corresponding threshold, in which case the NE would need to send a TCA for the previous period. Similarly, an NE would need to send a TCA

18. The Section layer parameters would need to be suppressed because an LOF failure is derived indirectly from an SEF defect (via an LOF defect), and SEF defects are listed in the definitions of the ES-S and SES-S parameters. In addition, a persistent LOF defect would cause (at least in the functional model used in the development of the criteria in this document) an AIS-L failure to be declared by the LTE processing the signal received from the STE. AIS-L defects are listed in the definitions of the ES-L and SES-L parameters, and therefore TCAs for the near-end Line layer parameters associated with the failed signal would also need to be suppressed.

for the previous period if the value in a register has reached or exceeded its threshold, a TCA has not been sent because of the potential for entry into or exit from unavailable time, the period expires (so the data in the current register is moved to the previous period register), and then unavailable time is not entered or exited.

While the data adjustment and TCA generation capabilities described above are already required to be provided based on other criteria, it is necessary to strictly interpret and correlate those criteria in order to determine that they are required. Therefore, the following requirement explicitly addresses those capabilities. Note that this is applicable only for parameters whose accumulation is affected in some way by entry into or exit from unavailable time.

R6-336 [992v2] The following apply regarding the adjustment of PM data in previous period registers (due to entry into or exit from unavailable time) and the generation of TCAs:

- An NE shall either provide the capability to adjust the data stored in its most recent previous period registers based on entry into or exit from unavailable time during the first 10 seconds of a new current period, or it shall delay updating the values in its current period registers (and the generation of TCAs and the moving of data between registers) for up to 10 seconds
- If the capability to adjust the data stored in previous period registers is supported, a TCA shall be sent to an OS when the value in a previous 15-minute or previous day register is adjusted so that it is greater than or equal to the corresponding threshold
- If the capability to adjust the data stored in previous period registers is supported, a TCA shall be sent to an OS when the value in a previous 15-minute or previous day register can no longer be adjusted so that it will be less than the corresponding threshold (i.e., when the potential entry into or exit from unavailable time that was inhibiting the generation of the TCA at the end of a period does not occur)
- If the NE delays updating its registers (and the generation of TCAs and the moving of data between registers) for up to 10 seconds, it shall use that delay time to determine if anything has occurred that should affect the data about to be reflected in the current period registers.

In general, an NE that uses a short delay as described above would still be considered to be conforming to the objective in GR-820-CORE to accumulate PM data in “real time”.

In addition to the autonomous transmission of TCA (and ORA) messages by an NE, a user may also query the NE for the values in any of its PM registers, including its current period registers and threshold registers.

R6-337 [640v3] The SONET NE shall provide the capability for the user to retrieve the contents of any PM parameter register at any time. It shall also provide the capability to retrieve the contents of any threshold register (i.e., the provisioned threshold value).

R6-338 [641] The SONET NE shall allow the user to schedule, and shall then perform periodic (and automatic) reporting of PM data for a monitored entity. The NE shall continue to send the appropriate PM data according to the schedule until instructed to stop by the user. This instruction to stop could be part of the scheduling information that started the periodic reporting, or it could be a separate request. The NE shall support both methods of stopping periodic performance reporting.

R6-339 [642] The SONET NE shall support the ability for the user to retrieve periodic PM report schedule information.

Table 6-8 PM Register Sizes and Default Thresholds

Parameter	Rate (e.g., OC-N, STS-M)	Minimum Register Size		Minimum Threshold Register Size		Threshold Default Value	
		15 minute	1 day	15 minute	1 day	15 minute	1 day
Physical							
LBC _{normal}	Optical	-	-	255	255	-	-
OPT _{normal}	Optical	-	-	255	255	-	-
OPR _{normal}	Optical	-	-	255	255	-	-
Section							
CV-S	All	16,383	1,048,575	16,383	1,048,575	-	-
ES-S	All	900	65,535	900	65,535	-	-
SES-S	All	900	65,535	900	65,535	-	-
SEFS-S	All	900	65,535	900	65,535	-	-
Line							
CV-L CV-LFE	All	16,383	1,048,575	16,383	1,048,575	-	-
ES-L ES-LFE	All	900	65,535	900	65,535	-	-
SES-L SES-LFE	All	900	65,535	900	65,535	-	-
UAS-L UAS-LFE	All	900	65,535	900	65,535	-	-
FC-L FC-LFE	All	72	4,095	NR	NR	NR	NR
PSC	All	63	255	NR	NR	NR	NR
PSD	All	900	65,535	NR	NR	NR	NR

Table 6-8 PM Register Sizes and Default Thresholds (Continued)

Parameter	Rate (e.g., OC-N, STS-M)	Minimum Register Size		Minimum Threshold Register Size		Threshold Default Value	
		15 minute	1 day	15 minute	1 day	15 minute	1 day
STS Path							
CV-P	M = 1	16,383	1,048,575	16,383	1,048,575	15	125
CV-PFE	3c	16,383	1,048,575	16,383	1,048,575	25	250
	12c	16,383	1,048,575	16,383	1,048,575	75	750
	48c	16,383	1,048,575	16,383	1,048,575	_a	_a
	192c	16,383	1,048,575	16,383	1,048,575	_a	_a
	768c	16,383	1,048,575	16,383	1,048,575	_a	_a
ES-P	M = 1	900	65,535	900	65,535	12	100
ES-PFE	3c	900	65,535	900	65,535	20	200
	12c	900	65,535	900	65,535	60	600
	48c	900	65,535	900	65,535	_a	_a
	192c	900	65,535	900	65,535	_a	_a
	768c	900	65,535	900	65,535	_a	_a
SES-P	All	900	65,535	900	65,535	3	7
SES-PFE	All	900	65,535	900	65,535	3	7
UAS-P	All	900	65,535	900	65,535	10	10
UAS-PFE	All	900	65,535	900	65,535	10	10
FC-P	All	72	4,095	NR	NR	NR	NR
FC-PFE	All	72	4,095	NR	NR	NR	NR
PPJC-PDet	All	1,048,575	16,777,215	1,048,575	16,777,215	-	-
NPJC-PDet	All	1,048,575	16,777,215	1,048,575	16,777,215	-	-
PPJC-PGen	All	1,048,575	16,777,215	1,048,575	16,777,215	-	-
NPJC-PGen	All	1,048,575	16,777,215	1,048,575	16,777,215	-	-
PJCDiff-P	All	1,048,575	16,777,215	1,048,575	16,777,215	-	-
PJCS-PDet	All	900	65,535	900	65,535	-	-
PJCS-PGen	All	900	65,535	900	65,535	-	-

Table 6-8 PM Register Sizes and Default Thresholds (Continued)

Parameter	Rate (e.g., OC-N, STS-M)	Minimum Register Size		Minimum Threshold Register Size		Threshold Default Value	
		15 minute	1 day	15 minute	1 day	15 minute	1 day
VT Path							
CV-V CV-VFE	All	16,383	1,048,575	16,383	1,048,575	-	-
ES-V ES-VFE	All	900	65,535	900	65,535	-	-
SES-V SES-VFE	All	900	65,535	900	65,535	-	-
UAS-V UAS-VFE	All	900	65,535	900	65,535	-	-
FC-V FC-VFE	All	72	4,095	NR	NR	NR	NR
PPJC-VDet NPJC-VDet	All	32,767	2,097,151	32,767	2,097,151	-	-
PPJC-VGen NPJC-VGen	All	32,767	2,097,151	32,767	2,097,151	-	-
PJCDiff-V	All	32,767	2,097,151	32,767	2,097,151	-	-
PJCS-VDet PJCS-VGen	All	900	65,535	900	65,535	-	-

Notes:

:- To be determined

NR: Not Required

a It has been proposed that the following values be added for these parameters:

- STS-48c CV-P and CV-PFE – 10 (15 minute) and 100 (1 day)
- STS-192c CV-P and CV-PFE – 20 (15 minute) and 200 (1 day)
- STS-768c CV-P and CV-PFE – 40 (15 minute) and 400 (1 day)
- STS-48c ES-P and ES-PFE – 8 (15 minute) and 80 (1 day)
- STS-192c ES-P and ES-PFE – 16 (15 minute) and 160 (1 day)
- STS-768c ES-P and ES-PFE – 32 (15 minute) and 320 (1 day).

Note that in determining these values, it was assumed that these types of STS path signals will primarily be transported using equipment for which the receiver sensitivity parameter is based on a BER of 1×10^{-12} rather than 1×10^{-10} (i.e., primarily via OC-192 or OC-768 signals).

6.2.2.2 Physical Layer PM

6.2.2.2.1 Physical Layer Parameters

The intent of physical layer PM parameters currently defined in this section is to enable proactive monitoring of the physical devices and facilities that act as the transmitter, optical path and receiver of the SONET signal, so that an early indication of a problem is possible, before a failure actually occurs.¹⁹ These parameters are different from other PM parameters in that the recorded values are snapshot values taken at a particular time during the current period, rather than counts that may change throughout the period. If a snapshot value is either less than or equal to a lower threshold value (if defined), or greater than or equal to an upper threshold value, then an ORA is sent to an OS. In addition, each of the physical layer parameters defined here is a normalized value, referenced to a nominal value and expressed as a percentage. In some cases the nominal value is defined by the supplier, while in other cases it may be set by the user.

Note that based on the definitions provided below, it is implied that at least for calculation purposes, the various nominal and measured parameter values need to be expressed in linear units (e.g., mW, μ A) rather than logarithmic units (e.g., dBm). Also note that the LBC_{normal} parameter may not be an appropriate parameter for monitoring some technologies (e.g., for monitoring transmitters utilizing uncooled lasers), and that the nominal parameter values do not necessarily need to be set such that the expected values of the corresponding normalized parameters are 100%. [Regarding this latter point, it is more important for the nominal values to be chosen such that the upper and/or lower thresholds can be set to values within the supported range (e.g., between 0 and 255%) that will not be crossed as a result of normal variations of the monitored parameters while still providing the user with appropriate indications of potential problems.]

1. LBC_{normal} – This parameter is a measure of the Laser Bias Current (LBC). The normalized value of the LBC, expressed as an integer percentage, is the monitored parameter:

$$LBC_{normal} = \frac{LBC}{LBC_0} \times 100$$

where LBC_0 is the nominal value of the LBC provided by the NE supplier.

2. OPT_{normal} – This parameter is a measure of the average optical output power of the transmitter, or the Optical Power Transmitted (OPT). The normalized value of OPT, expressed as an integer percentage, is the monitored parameter:

19. In addition to the SONET Physical layer parameters defined here, ANSI T1.231.04 also defines an optional Loss Of Signal Seconds (LOSS) parameter. That parameter provides a count of the seconds during which a Physical layer defect (i.e., an LOS defect) is present, and thus is essentially the Physical layer equivalent of the SEFS-S parameter defined in [Section 6.2.2.3.1](#) for use at the Section layer. If an LOSS parameter is supported by a SONET NE, it is expected that it will be implemented such that it meets the specifications in T1.231.04.

$$\text{OPT}_{\text{normal}} = \frac{\text{OPT}}{\text{OPT}_0} \times 100$$

where OPT_0 is the nominal value of OPT provided by the NE supplier.

3. $\text{OPR}_{\text{normal}}$ – This parameter is a measure of the average optical power of the received signal, or the Optical Power Received (OPR). The normalized value of OPR, expressed as an integer percentage, is the monitored parameter:

$$\text{OPR}_{\text{normal}} = \frac{\text{OPR}}{\text{OPR}_0} \times 100$$

where OPR_0 is the nominal value of OPR.

Unlike LBC_0 and OPT_0 , which are equipment-specific values that are not expected to need to be set by the user, OPR_0 is an installation or application-specific value. As such, its value needs to be provisionable (see **R6-343 [1214]**). In addition, given the constraints imposed by the use of linear units, integer values and the (acceptable) minimum register and threshold register sizes of 255%, care needs to be taken in determining the value to be used in any particular installation. Among the most likely values to which OPR_0 might be set are those that are somewhat above the receiver's specified minimum received power level (i.e., somewhat greater than P_{Rmin}). Such values are likely to be appropriate in any application where the user is primarily concerned about a low OPR and the initial OPR value is relatively close to P_{Rmin} , as it would allow the lower threshold to be set to a value below 100% (e.g., to 25% if OPR_0 was set to four times P_{Rmin}) and the upper threshold to be set to a relatively large value that would be unlikely to be reached (e.g., 255%). On the other hand, for installations with relatively little attenuation in the optical path, a nominal value somewhat below the receiver's specified maximum received power level (i.e., somewhat less than P_{Rmax}) could be appropriate. This would allow the user to set the upper threshold to a value above 100% (e.g., to 200% if OPR_0 was set to 0.5 times P_{Rmax}) and the lower threshold to a small value that would be unlikely to be reached (e.g., 1%). Finally, for installations where the optical path attenuation is expected to be well within the allowed range, a nominal value relatively close to the power level received at equipment turn-up (assuming healthy equipment at turn-up) may be appropriate.

6.2.2.2.2 Physical Layer PM Criteria

The following Physical layer PM criteria apply to all SONET NEs with optical interfaces.

-
- CR6-340 [1126]** A SONET NE may be required to support the $\text{LBC}_{\text{normal}}$ parameter to provide a measurement of the health of each optical transmitter.
- O6-341 [643v5]** A SONET NE should support the $\text{OPT}_{\text{normal}}$ parameter to provide a measurement of the health of each optical transmitter.

- O6-342 [994v3]** A SONET NE should support the OPR_{normal} parameter to provide a measurement of the physical layer characteristics of the incoming signal at each optical receiver.
- R6-343 [1214]** If the OPR_{normal} parameter is supported, the value of OPR_0 used in the calculation of that parameter shall be provisionable on a per-optical-receiver basis. In addition, the default value of OPR_0 shall be clearly documented.
- R6-344 [1082]** For each Physical layer PM parameter that it supports, the SONET NE shall measure and record the parameter value once per period. This snapshot value shall be recorded at approximately the same time (i.e., within ± 10 seconds) after the start of each new period.

To provide values for the user to retrieve in the current period, it is important for the NE to record its snapshot values as soon as practical after the start of each new period. For relatively small NEs, it may be practical to record the values “immediately” after the start of the new period, while for large NEs some delay may be necessary while other end-of-period/start-of-period processing is performed. In any case, the following objective applies.

-
- O6-345 [1083]** The SONET NE should record its Physical layer PM parameter snapshot values within one minute after the start of each new period.

In addition to snapshot Physical layer PM parameter values that a user might want to store and use for purposes such as long-term trend analyses, in some situations it might be desirable for the NE to monitor for and store the largest and smallest values of the various parameters that occur during a period, or for the user to be able to determine the current value of a parameter. While neither of these capabilities is covered by the criteria in this section, the former could be supported using the tidemarking techniques defined in ANSI T1.231, and continuous monitoring and retrieval capabilities could be provided as separate Physical layer diagnostics (see [Section 6.2.3.2.1](#)). Alternatively (for the continuous monitoring case), the NE could record a new snapshot value after a current period parameter register is initialized by the user. Although an invalid-data flag would be set to indicate that the current period register had been initialized, this new snapshot value would actually be a valid value that would reflect the current value of the parameter at (approximately) the time that the register was initialized.

-
- R6-346 [1215]** If an NE supports tidemarking of a SONET Physical layer PM parameter, that function shall be provided according to the definitions and specifications in ANSI T1.231, and shall be in addition to (as opposed to in place of) the functions defined in this document.
- O6-347 [1084]** A SONET NE should record a new snapshot value within one minute after the current period register for a Physical layer PM parameter is initialized by the user.
-

As discussed in [Section 6.2.2.1](#), an NE must generate an ORA when a snapshot value for a Physical layer PM parameter is not in its acceptable range. For the OPT_{normal} and OPR_{normal} parameters, the acceptable range has been defined to be based on both upper and lower thresholds. On the other hand, only an upper threshold is considered necessary for defining the acceptable range for the LBC_{normal} parameter.

R6-348 [645v3] A SONET NE shall provide lower threshold registers for the OPT_{normal} and OPR_{normal} parameters (if supported) and shall also provide an upper threshold register for each Physical layer PM parameter that it supports.

Similar to the case discussed in [Section 6.2.2.2.1](#) regarding nominal values, the acceptable ranges for the LBC_{normal} and OPT_{normal} parameters are generally expected to be equipment specific. Therefore, the thresholds for those parameters are currently not required (but are allowed) to be provisionable. On the other hand, the acceptable range for the OPR_{normal} parameter is likely to be installation or application specific. As such, that parameter's upper and lower thresholds need to be provisionable (i.e., the third bullet in **R6-331 [638v3]** is applicable for OPR_{normal}).

6.2.2.3 Section Layer PM

6.2.2.3.1 Section Layer Parameters

The definitions of the SONET Section layer PM parameters specified in this document are shown below.²⁰ (Note that no far-end PM parameters are defined for the Section layer.)

1. Section Severely Errored Framing Seconds (SEFS-Ss) – The SEFS-S parameter is a count of the seconds during which (at any point during the second) an SEF defect was present. See [Section 5.5](#) for the definition of an SEF defect. (In addition, note that an SEF defect is expected to be present during most seconds in which an LOS or LOF defect is present. However, there may be situations when that is not the case, and the SEFS-S parameter is only incremented based on the presence of the SEF defect.)
2. Section Coding Violations (CV-Ss) – The CV-S parameter is a count of BIP errors detected at the Section layer (i.e., using the B1 byte in the incoming SONET signal). Up to eight Section BIP errors can be detected per STS-N frame, with each error incrementing the CV-S current second register.
3. Section Errored Seconds (ES-Ss) – The ES-S parameter is a count of the number of seconds during which (at any point during the second) at least one Section BIP error was detected or an SEF or LOS defect was present.

²⁰In addition to the SONET Section layer parameters defined here, ANSI T1.231.04 also defines optional Section Errored Second Type A (ESA-S) and Errored Second Type B (ESB-S) parameters. If those parameters are supported by a SONET NE, it is expected that they will be implemented such that they meet the specifications in T1.231.04.

4. Section Severely Errored Seconds (SES-Ss) – The SES-S parameter is a count of the seconds during which K or more Section BIP errors were detected or an SEF or LOS defect was present. The number of BIP errors that cause a second to be considered an SES-S has been changed several times in both TR/GR-253 and the applicable ANSI SONET standards, and thus may need to be settable.²¹ Table 6-9 contains the current values for K for the various SONET rates. These values are based on those that appeared in the 1997 version of ANSI T1.231, and now appear in ANSI T1.231.04.

Table 6-9 Section BIP Errors to Trigger a Section SES

Rate	K
OC-1	52
OC-3	155
OC-12	616
OC-48	2,392
OC-192	8,554
OC-768	22,778

6.2.2.3.2 Section Layer PM Criteria

The following requirements for Section layer performance monitoring apply to all SONET NEs.

-
- R6-349 [656v2]** A SONET NE shall be capable of accumulating the SEFS-S parameter.
- R6-350 [657]** A SONET NE shall perform thresholding for the SEFS-S parameter.
- CR6-351 [658]** A SONET NE may be required to support applications where Line and Section spans are not coincident.
- R6-352 [659v2]** A SONET NE that supports applications where Line and Section spans are not coincident shall be capable of accumulating CV-Ss, ES-Ss and SES-Ss.
- R6-353 [660v2]** A SONET NE shall perform thresholding for the CV-S, ES-S and SES-S parameters if those parameters are supported.
-

Note that the B1 byte is not required to be generated for drop-side signals (i.e., an NE transmitting a drop-side signal may set the B1 byte to either the Section BIP-8 code or all-zeros), and that an NE that is capable of receiving a drop-side signal at a

21. The ability to set the value of K (at the Section, Line, STS path and VT path layers) would be needed if future standards activities change the definitions of the SES parameters, and is not meant to imply that these numbers are variables.

particular interface must be capable of ignoring the value in the incoming B1 byte (see [Section 3.3.2.1](#)). In addition, the value of the incoming B1 byte directly affects the accumulation of CV-S, ES-S and SES-S parameters.²² Therefore, if an NE supports only one or both of the following two Section PM accumulation options:

- Accumulate all four of the defined Section PM parameters at a particular interface
- Accumulate none of the defined Section PM parameters at a particular interface

(i.e., if it cannot be provisioned so that it just accumulates SEFS-Ss, or so that it ignores the B1 byte and bases its ES-S and SES-S parameters only on the presence of SEF and LOS defects), then it will not be capable either of receiving some drop-side signals, or of accumulating any valid Section PM data for those signals. Although the following objective is primarily intended to address this issue, the capability described in the first bullet item could also be useful in other applications with coincident Section and Line spans. Also note that if the capability described in the second bullet item is supported and used, the CV-S registers will not contain valid data (i.e., they will always contain a value of zero), and in most situations the ES-S, SES-S and SEFS-S registers will contain redundant data.

O6-354 [1033v2] A SONET NE that supports the accumulation of the SEFS-S, CV-S, ES-S and SES-S parameters should provide one of the following capabilities:

- A user-provisionable, per-Section option to accumulate the CV-S, ES-S and SES-S parameters that is independent of any user-provisionable option to accumulate the SEFS-S parameter
 - A user-provisionable, per-Section option to ignore the B1 byte for the purposes of accumulating the CV-S, ES-S and SES-S parameters.
-

6.2.2.4 Line Layer PM

The definitions of the SONET Line layer PM parameters specified in this document are provided in the following two subsection.²³

6.2.2.4.1 Near-End Line Layer Parameters

The near-end (or incoming) Line layer PM parameters are as follows.

22. If one NE sets the B1 byte to all-zeros and another NE attempts to accumulate CV-Ss, ES-Ss and SES-Ss using that value, none of those parameter registers will contain valid data. In such a situation, it is likely that every second will appear to contain enough BIP-8 errors to be considered an SES-S, causing CV-S accumulation to be inhibited, the continuous accumulation of ES-Ss and SES-Ss, and the generation of unnecessary TCAs.

23. In addition to the SONET Line layer PM parameters defined here, ANSI T1.231.04 also defines optional near-end and far-end Line ESA, ESB, and AIS Second (AISS-L and AISS-LFE) parameters. If those parameters are supported by a SONET NE, it is expected that they will be implemented such that they meet the specifications in T1.231.04.

1. Near-end Line Coding Violations (CV-Ls) – The CV-L parameter is a count of BIP errors detected at the Line layer (i.e., using the B2 bytes in the incoming SONET signal). Up to 8×N BIP errors can be detected per STS-N frame, with each error incrementing the CV-L current second register.
2. Near-end Line Errored Seconds (ES-Ls) – The ES-L parameter is a count of the seconds during which (at any point during the second) at least one Line BIP error was detected or an AIS-L defect (or a lower-layer, traffic-related, near-end defect, see [Section 6.2.1.8.2](#)) was present.
3. Near-end Line Severely Errored Seconds (SES-Ls) – The SES-L parameter is a count of the seconds during which K or more Line BIP errors were detected or an AIS-L defect (or a lower-layer, traffic-related, near-end defect, see [Section 6.2.1.8.2](#)) was present. The number of BIP errors that cause a second to be considered an SES-L has been changed several times in both TR/GR-253 and the applicable ANSI SONET standards, and may need to be settable. [Table 6-10](#) contains the current values for K for the various SONET rates. These values are based on those that appeared in the 1997 version of ANSI T1.231, and now appear in ANSI T1.231.04.

Table 6-10 Line BIP Errors to Trigger a Line SES (Near-End and Far-End)

Rate	K
OC-1	51
OC-3	154
OC-12	615
OC-48	2,459
OC-192	9,835
OC-768	39,339

4. Near-end Line Unavailable Seconds (UAS-Ls) – The UAS-L parameter is a count of the seconds during which the Line was considered unavailable. A Line becomes unavailable at the onset of 10 consecutive seconds that qualify as SES-Ls, and continues to be unavailable until the onset of 10 consecutive seconds that do not qualify as SES-Ls.
5. Near-end Line Failure Counts (FC-Ls) – The FC-L parameter is a count of the number of near-end Line failure events. A failure event begins when the AIS-L failure (or a lower-layer, traffic-related, near-end failure, see [Section 6.2.1.8.2](#)) is declared, and ends when the failure is cleared. A failure event that begins in one period and ends in another period is counted only in the period in which it begins. Note that functionally, an AIS-L failure will be declared either by receiving (and timing) an AIS-L signal from another NE, or by receiving (and timing) an internally generated AIS-L signal from STE in the same NE where the LTE resides.
6. Line Protection Switching Counts (PSC-Ls) – For a working line, the PSC-L parameter is a count of the number of times that service has been switched from the monitored line to the protection line, plus the number of times it has been

switched back to the working line. For the protection line, it is a count of the number of times that service has been switched from any working line to the protection line, plus the number of times service has been switched back to a working line. The PSC-L parameter is only applicable if Line-level protection switching is used.

7. Line Protection Switching Duration (PSD-L) – For a working line, the PSD-L parameter is a count of the seconds that service was being carried on the protection line. For the protection line, it is a count of the seconds that the line was being used to carry service. The PSD-L parameter is only applicable if revertive Line-level protection switching is used.

6.2.2.4.2 Far-End Line Layer Parameters

Far-end Line layer performance is conveyed back to the near-end LTE via the K2 byte (RDI-L) and the M0 and/or M1 bytes (REI-L). The far-end Line layer PM parameters are as follows.

1. Far-end Line Coding Violations (CV-LFEs) – The CV-LFE parameter is a count of the number of BIP errors detected by the far-end LTE and reported back to the near-end LTE using the REI-L indication in the LOH. For SONET signals at rates below OC-48 and at OC-768, up to $8 \times N$ BIP errors per STS-N frame can be indicated using the REI-L. For OC-48 and OC-192 signals, up to 255 BIP errors per STS-N frame can be indicated. The CV-LFE current second register is incremented for each BIP error indicated by the incoming REI-L.
2. Far-end Line Errored Seconds (ES-LFEs) – The ES-LFE parameter is a count of the seconds during which (at any point during the second) at least one Line BIP error was reported by the far-end LTE (using the REI-L) or an RDI-L defect was present.
3. Far-end Line Severely Errored Seconds (SES-LFEs) – The SES-LFE parameter is a count of the seconds during which K or more Line BIP errors were reported by the far-end LTE or an RDI-L defect was present. The number of reported far-end BIP errors that cause a second to be considered an SES-LFE may need to be settable. [Table 6-10](#) contains the current values for K for the various SONET rates.
4. Far-end Line Unavailable Seconds (UAS-LFE) – The UAS-LFE parameter is a count of the seconds during which the Line is considered unavailable at the far end. A Line is considered unavailable at the far end at the onset of 10 consecutive seconds that qualify as SES-LFEs, and continues to be considered unavailable until the onset of 10 consecutive seconds that do not qualify as SES-LFEs.
5. Far-end Line Failure Counts (FC-LFEs) – The FC-LFE parameter is a count of the number of far-end Line failure events. A failure event begins when the RFI-L failure is declared, and ends when the RFI-L failure is cleared. A failure event that begins in one period and ends in another period is counted only in the period in which it begins.

6.2.2.4.3 Line Layer PM Criteria

The following requirements for Line layer performance monitoring apply to all SONET NEs that terminate the Line layer.

-
- R6-355 [661]** A SONET NE providing LTE functions shall accumulate CV-Ls, ES-Ls, SES-Ls, UAS-Ls and FC-Ls for each Line.
- R6-356 [662]** A SONET NE providing LTE functions shall perform thresholding for the CV-L, ES-L, SES-L and UAS-L parameters.
- R6-357 [663]** A SONET NE that supports Line-level protection switching for a given Line shall accumulate PSC-Ls for that Line.
- R6-358 [664v2]** A SONET NE that is using revertive protection switching for a given Line shall accumulate the PSD-L parameter for that Line.
-

Thresholding is not required for the PSC-L, PSD-L and FC-L parameters.

-
- R6-359 [667]** A SONET NE providing LTE functions shall provide the capability, on a per-Line basis, to accumulate the CV-LFE, ES-LFE, SES-LFE, UAS-LFE and FC-LFE parameters, and to activate and deactivate the accumulation of these parameters (as a group). The default setting shall be “not active.”
- R6-360 [668]** A SONET NE that is accumulating far-end Line PM parameters shall perform thresholding for the CV-LFE, ES-LFE, SES-LFE and UAS-LFE parameters.
-

6.2.2.5 STS Path Layer PM

The definitions of the STS path layer PM parameters specified in this document are provided in the following two subsection.²⁴

6.2.2.5.1 Near-End STS Path Layer Parameters

The near-end STS path layer PM parameters are as follows.

1. Near-end STS Path Coding Violations (CV-Ps) – The CV-P parameter is a count of BIP errors detected at the STS path layer (i.e., using the B3 byte in the incoming STS POH). Up to 8 BIP errors can be detected per frame, with each error incrementing the CV-P current second register.

24. In addition to the STS path layer PM parameters defined here, ANSI T1.231.04 also defines optional near-end and far-end STS path ESA and ESB parameters, and (for possible use by NEs that support STS path protection switching) PSC-P and PSD-P parameters. If those parameters are supported by a SONET NE, it is expected that they will be implemented such that they meet the specifications in T1.231.04.

2. Near-end STS Path Errored Seconds (ES-Ps) – The ES-P parameter is a count of the seconds during which (at any point during the second) at least one STS path BIP error was detected, or an AIS-P defect (or a lower-layer, traffic-related, near-end defect, see [Section 6.2.1.8.2](#)), an LOP-P defect or, if the STS PTE monitoring the path supports ERDI-P for that path, an UNEQ-P or TIM-P defect was present.
3. Near-end STS Path Severely Errored Seconds (SES-Ps) – The SES-P parameter is a count of the seconds during which K or more STS path BIP errors were detected, or an AIS-P defect (or a lower-layer, traffic-related, near-end defect, see [Section 6.2.1.8.2](#)), an LOP-P defect or, if the STS PTE monitoring the path supports ERDI-P for that path, an UNEQ-P or TIM-P defect was present. The number of BIP errors that cause a second to be considered an SES-P may need to be settable. As shown in [Table 6-11](#), the current values for K (which are consistent with those that appear in ANSI T1.231.04) are independent of the particular size of the STS path (e.g., STS-1, STS-3c).

Table 6-11 STS Path BIP Errors to Trigger an STS Path SES (Near-End and Far-End)

Rate	K
STS-1	2,400
STS-Nc	2,400

4. Near-end STS Path Unavailable Seconds (UAS-Ps) – The UAS-P parameter is a count of the seconds during which the STS path was considered unavailable. An STS path becomes unavailable at the onset of 10 consecutive seconds that qualify as SES-Ps, and continues to be unavailable until the onset of 10 consecutive seconds that do not qualify as SES-Ps.
5. Near-end STS Path Failure Counts (FC-P) – The FC-P parameter is a count of the number of near-end STS path failure events. A failure event begins when an AIS-P failure (or a lower-layer, traffic-related, near-end failure, see [Section 6.2.1.8.2](#)), an LOP-P failure or, if the STS PTE monitoring the path supports ERDI-P for that path, an UNEQ-P or TIM-P failure is declared. The failure event ends when these failures are cleared. A failure event that begins in one period and ends in another period is counted only in the period in which it begins. Note that functionally, an AIS-P failure will be declared either by receiving (and timing) an AIS-P signal from another NE, or by receiving (and timing) an internally generated AIS-P signal from LTE in the same NE where the STS PTE resides.
6. Positive Pointer Justification Count - STS Path Detected (PPJC-PDet) – The PPJC-PDet parameter is a count of the positive pointer justifications (i.e., valid increment operations) detected on a particular STS path in an incoming SONET signal.
7. Negative Pointer Justification Count - STS Path Detected (NPJC-PDet) – The NPJC-PDet parameter is a count of the negative pointer justifications (i.e., valid

decrement operations) detected on a particular STS path in an incoming SONET signal.

8. Positive Pointer Justification Count - STS Path Generated (PPJC-PGen) – The PPJC-PGen parameter is a count of the positive pointer justifications (i.e., increment operations) generated for a particular STS path to reconcile the frequency of the SPE with the local clock.
9. Negative Pointer Justification Count - STS Path Generated (NPJC-PGen) – The NPJC-PGen parameter is a count of the negative pointer justifications (i.e., decrement operations) generated for a particular STS path to reconcile the frequency of the SPE with the local clock.
10. Pointer Justification Count Difference - STS Path (PJCDiff-P) – The PJCDiff-P parameter is the absolute value of the difference between the net number of detected pointer justification counts and the net number of generated pointer justification counts for a particular STS path. That is:
$$\text{PJCDiff-P} = |(\text{PPJC-PGen} - \text{NPJC-PGen}) - (\text{PPJC-PDet} - \text{NPJC-PDet})|$$
11. Pointer Justification Count Seconds - STS Path Detect (PJCS-PDet) – The PJCS-PDet parameter is a count of the one-second intervals containing one or more PPJC-PDet or NPJC-PDet for a particular STS path.
12. Pointer Justification Count Seconds - STS Path Generate (PJCS-PGen) – The PJCS-PGen parameter is a count of the one-second intervals containing one or more PPJC-PGen or NPJC-PGen for a particular STS path.

Note that parameters 6 through 12 are sometimes referred to collectively as the “STS-PJ-related” parameters. Also note that if an NE happens to support ERDI-P on a provisionable basis, then the particular defects/failures that it includes in its ES-P, SES-P and FC-P parameter definitions would need to change (on either a per-STs-path or NE-wide basis) based on which version of RDI-P it is provisioned to use. On the other hand, an NE that supports ERDI-P and can differentiate between one-bit RDI-P and ERDI-P Server defects/failures (see **R6-249 [962v3]**) is not required or expected to be able to adjust its PM parameter definitions based on which type of RDI-P signals it receives.

6.2.2.5.2 Far-End STS Path Layer Parameters

Far-end STS path layer performance is conveyed back to the near-end STS PTE via bits 1 through 4 (REI-P) and 5 through 7 (RDI-P) of the G1 byte. The far-end STS path layer PM parameters are as follows.

1. Far-end STS Path Coding Violations (CV-PFEs) – The CV-PFE parameter is a count of the number of BIP errors detected by the far-end STS PTE and reported back to the near-end STS PTE using the REI-P indication in the STS POH. Up to 8 BIP errors per frame can be indicated. The CV-PFE current second register is incremented for each BIP error indicated by the incoming REI-P.
2. Far-end STS Path Errored Seconds (ES-PFEs) – The ES-PFE parameter is a count of the seconds during which (at any point during the second) at least one STS path BIP error was reported by the far-end STS PTE (using the REI-P

indication), a one-bit RDI-P defect was present, or (if ERDI-P is supported, see [Section 6.2.1.3.2](#)) an ERDI-P Server or Connectivity defect was present.

3. Far-end STS Path Severely Errored Seconds (SES-PFEs) – The SES-PFE parameter is a count of the seconds during which K or more STS path BIP errors were reported by the far-end STS PTE, a one-bit RDI-P defect was present, or (if ERDI-P is supported) an ERDI-P Server or Connectivity defect was present. The number of reported far-end BIP errors that cause a second to be considered an SES-PFE may need to be settable. [Table 6-11](#) contains the current values for K for STS paths.
4. Far-end STS Path Unavailable Seconds (UAS-PFE) – The UAS-PFE parameter is a count of the seconds during which the STS path is considered unavailable at the far end. An STS path is considered unavailable at the far end at the onset of 10 consecutive seconds that qualify as SES-PFEs, and continues to be considered unavailable until the onset of 10 consecutive seconds that do not qualify as SES-PFEs.
5. Far-end STS Path Failure Counts (FC-PFEs) – The FC-PFE parameter is a count of the number of far-end STS path failure events. A failure event begins when a one-bit RFI-P failure, or (if ERDI-P is supported) an ERFI-P Server or Connectivity failure is declared. The failure event ends when the RFI-P failure is cleared. A failure event that begins in one period and ends in another period is counted only in the period in which it begins.

Note that with the parameter definitions shown above, unless the STS PTE at both ends of a path support the same version of RDI-P, inconsistencies can be expected to occur between the near-end PM data accumulated at one NE and the far-end data accumulated at the far-end NE. The reason for this is that connectivity defects and failures (e.g., UNEQ-P, ERDI-P Connectivity) are included in the definitions of the near-end and far-end ES-P, SES-P and FC-P parameters if ERDI-P is supported, but are not included if one-bit RDI-P is being used.

6.2.2.5.3 STS Path Layer PM Criteria

The following requirements for STS path layer performance monitoring apply to all SONET NEs that terminate the STS path layer, with the possible exception of SONET NEs supporting digital switch trunk side interfaces or UPSR configurations (see GR-1400-CORE).

-
- R6-361 [669]** A SONET NE providing STS PTE functions shall accumulate CV-Ps, ES-Ps, SES-Ps, UAS-Ps and FC-Ps for each terminated STS path.
- R6-362 [670]** A SONET NE providing STS PTE functions shall perform thresholding for the CV-P, ES-P, SES-P and UAS-P parameters.
- CR6-363 [1085]** STS PTE that processes the STS pointer from an incoming SONET signal (i.e., STS PTE that terminates a path that has not been reconciled to the local clock by upstream LTE within the same NE) may be required to support the capability to accumulate the STS-PJ-related parameters defined in [Section 6.2.2.5.1](#).
-

Note that although they are included for consistency with the case where the STS-PJ-related parameters are accumulated (as intermediate-path PM parameters) at LTE, the PPJC-PGen, NPJC-PGen and PJCS-PGen parameters are expected to always be equal to zero when they are accumulated at STS PTE.

-
- R6-364 [1086]** If the accumulation of the STS-PJ-related parameters is supported, the user shall be able to activate that accumulation on a per-STS-path basis. The default setting shall be “not active”.
- R6-365 [1087]** A SONET NE that is accumulating STS-PJ-related parameters shall perform thresholding for the PPJC-PDet, NPJC-PDet, PPJC-PGen, NPJC-PGen, PJCDiff-P, PJCS-PDet and PJCS-PGen parameters.
-

Note that unlike all of the other PM parameters that are based on integer counts, the PJCDiff-P and (at the VT layer) PJCDiff-V parameters are not essentially monotonically increasing functions during any particular accumulation period. Therefore it is possible that their values may reach or exceed the corresponding thresholds for generating TCAs, decrease to values less than the thresholds, and then increase back to the thresholds, etc. Existing requirements indicate that multiple TCAs must not be generated in such situations (see **R820-38** and **R820-39** in GR-820-CORE); however, it may still be desirable to reduce the chance that those situations will occur. This can be done by providing default threshold values that are large enough that they will only be reached if there is frequency offset in the system (see GR-253-ILR Issue ID 253-30).

-
- R6-366 [672]** A SONET NE providing STS PTE functions shall provide the capability, on a per-STS-path basis, to accumulate the CV-PFE, ES-PFE, SES-PFE, UAS-PFE and FC-PFE parameters, and to activate and deactivate the accumulation of these parameters (as a group). The default setting shall be “not active.”
- R6-367 [673]** A SONET NE that is accumulating far-end STS path PM parameters shall perform thresholding for the CV-PFE, ES-PFE, SES-PFE and UAS-PFE parameters.
-

6.2.2.6 VT Path Layer PM

The definitions of the VT path layer PM parameters specified in this document are provided in the following two subsection.²⁵

²⁵In addition to the VT path layer PM parameters defined here, ANSI T1.231.04 also defines optional near-end and far-end VT path ESA and ESB parameters, and (for possible use by NEs that support VT path protection switching) PSC-V and PSD-V parameters. If those parameters are supported by a SONET NE, it is expected that they will be implemented such that they meet the specifications in T1.231.04.

6.2.2.6.1 Near-End VT Path Layer Parameters

The near-end VT path layer PM parameters are as follows.

1. Near-end VT Path Coding Violations (CV-Vs) – The CV-V parameter is a count of BIP errors detected at the VT path layer (i.e., using bits 1 and 2 of the V5 byte in the incoming VT POH). Up to 2 BIP errors can be detected per VT superframe, with each error incrementing the CV-V current second register.
2. Near-end VT Path Errored Seconds (ES-Vs) – The ES-V parameter is a count of the seconds during which (at any point during the second) at least one VT path BIP error was detected, or an AIS-V defect (or a lower-layer, traffic-related, near-end defect, see [Section 6.2.1.8.2](#)), an LOP-V defect or, if the VT PTE monitoring the path supports ERDI-V for that path, an UNEQ-V or TIM-V defect was present.
3. Near-end VT Path Severely Errored Seconds (SES-Vs) – The SES-V parameter is a count of the seconds during which K or more VT path BIP errors were detected, or an AIS-V defect (or a lower-layer, traffic-related, near-end defect, see [Section 6.2.1.8.2](#)), an LOP-V defect or, if the VT PTE monitoring the path supports ERDI-V for that path, an UNEQ-V or TIM-V defect was present. The number of BIP errors that cause a second to be considered an SES-V may need to be settable. As shown in [Table 6-12](#), the current values for K (which are consistent with those that appear in ANSI T1.231.04) are independent of the particular size of the VT path (e.g., VT1.5, VT2).

Table 6-12 VT Path BIP Errors to Trigger a VT Path SES (Near-End and Far-End)

Rate	K
VT1.5	600
VT2	600
VT3	600
VT6	600

4. Near-end VT Path Unavailable Seconds (UAS-Vs) – The UAS-V parameter is a count of the seconds during which the VT path was considered unavailable. A VT path becomes unavailable at the onset of 10 consecutive seconds that qualify as SES-Vs, and continues to be unavailable until the onset of 10 consecutive seconds that do not qualify as SES-Vs.
5. Near-end VT Path Failure Counts (FC-V) – The FC-V parameter is a count of the number of near-end VT path failure events. A failure event begins when an AIS-V failure (or a lower-layer, traffic-related, near-end failure, see [Section 6.2.1.8.2](#)), an LOP-V failure or, if the VT PTE monitoring the path supports ERDI-V for that path, an UNEQ-V or TIM-V failure is declared. The failure event ends when these failures are cleared. A failure event that begins in one period and ends in another period is counted only in the period in which it begins. Note that functionally, an AIS-V failure will be declared either by receiving (and timing) an AIS-V signal

from another NE, or by receiving (and timing) an internally generated AIS-V signal from STS PTE in the same NE where the VT PTE resides.

6. Positive Pointer Justification Count - VT Path Detected (PPJC-VDet) – The PPJC-VDet parameter is a count of the positive pointer justifications (i.e., valid increment operations) detected on a particular VT path in an incoming SONET signal.
7. Negative Pointer Justification Count - VT Path Detected (NPJC-VDet) – The NPJC-VDet parameter is a count of the negative pointer justifications (i.e., valid decrement operations) detected on a particular VT path in an incoming SONET signal.
8. Positive Pointer Justification Count - VT Path Generated (PPJC-VGen) – The PPJC-VGen parameter is a count of the positive pointer justifications (i.e., increment operations) generated for a particular VT path to reconcile the frequency of the SPE with the local clock.
9. Negative Pointer Justification Count - VT Path Generated (NPJC-VGen) – The NPJC-VGen parameter is a count of the negative pointer justifications (i.e., decrement operations) generated for a particular VT path to reconcile the frequency of the SPE with the local clock.
10. Pointer Justification Count Difference - VT Path (PJCDiff-V) – The PJCDiff-V parameter is the absolute value of the difference between the net number of detected pointer justification counts and the net number of generated pointer justification counts for a particular VT path. That is:

$$PJCDiff-V = |(PPJC-VGen - NPJC-VGen) - (PPJC-VDet - NPJC-VDet)|$$

11. Pointer Justification Count Seconds - VT Path Detect (PJCS-VDet) – The PJCS-VDet parameter is a count of the one-second intervals containing one or more PPJC-VDet or NPJC-VDet for a particular VT path.
12. Pointer Justification Count Seconds - VT Path Generate (PJCS-VGen) – The PJCS-VGen parameter is a count of the one-second intervals containing one or more PPJC-VGen or NPJC-VGen for a particular VT path.

Note that parameters 6 through 12 are sometimes referred to collectively as the “VT-PJ-related” parameters. Also note that if an NE happens to support ERDI-V on a provisionable basis, then the particular defects/failures that it includes in its ES-V, SES-V and FC-V parameter definitions would need to change (on either a per-VT-path or NE-wide basis) based on which version of RDI-V it is provisioned to use. On the other hand, an NE that supports ERDI-V and can differentiate between one-bit RDI-V and ERDI-V Server defects/failures (see **R6-265 [968v3]**) is not required or expected to be able to adjust its PM parameter definitions based on which type of RDI-V signals it receives.

6.2.2.6.2 Far-End VT Path Layer Parameters

Far-end VT path layer performance is conveyed back to the near-end VT PTE via bit 3 of the V5 byte (REI-V), and either bits 5 through 7 of the Z7 byte or bit 8 of the V5 byte (RDI-V). The far-end VT path layer PM parameters are as follows.

1. Far-end VT Path Coding Violations (CV-VFEs) – The CV-VFE parameter is a count of the number of BIP errors detected by the far-end VT PTE and reported back to the near-end VT PTE using the REI-V indication in the VT path overhead. Note that only one BIP error can be indicated per VT superframe using the REI-V bit (out of the two BIP errors that can be detected). The CV-VFE current second register is incremented for each BIP error indicated by the incoming REI-V.
2. Far-end VT Path Errored Seconds (ES-VFEs) – The ES-VFE parameter is a count of the seconds during which (at any point during the second) at least one VT path BIP error was reported by the far-end VT PTE (using the REI-V indication), a one-bit RDI-V defect was present, or (if ERDI-V is supported, see [Section 6.2.1.3.3](#)) an ERDI-V Server or Connectivity defect was present.
3. Far-end VT Path Severely Errored Seconds (SES-VFEs) – The SES-VFE parameter is a count of the seconds during which K or more VT path BIP errors were reported by the far-end VT PTE, a one-bit RDI-V defect was present, or (if ERDI-V is supported) an ERDI-V Server or Connectivity defect was present. The number of reported far-end BIP errors that cause a second to be considered an SES-VFE may need to be settable. [Table 6-12](#) contains the current values for K for the various VT paths.
4. Far-end VT Path Unavailable Seconds (UAS-VFE) – The UAS-VFE parameter is a count of the seconds during which the VT path is considered unavailable at the far end. A VT path is considered unavailable at the far end at the onset of 10 consecutive seconds that qualify as SES-VFEs, and continues to be considered unavailable until the onset of 10 consecutive seconds that do not qualify as SES-VFEs.
5. Far-end VT Path Failure Counts (FC-VFEs) – The FC-VFE parameter is a count of the number of far-end VT path failure events. A failure event begins when a one-bit RFI-V failure, or (if ERDI-V is supported) an ERFI-V Server or Connectivity failure is declared. The failure event ends when the RFI-V failure is cleared. A failure event that begins in one period and ends in another period is counted only in the period in which it begins.

Note that with the parameter definitions shown above, unless the VT PTE at both ends of a path support the same version of RDI-V, inconsistencies can be expected to occur between the near-end PM data accumulated at one NE and the far-end data accumulated at the far-end NE. The reason for this is that connectivity defects and failures (e.g., UNEQ-V, ERDI-V Connectivity) are included in the definitions of the near-end and far-end ES-V, SES-V and FC-V parameters if ERDI-V is supported, but are not included if one-bit RDI-V is being used.

6.2.2.6.3 VT Path Layer PM Criteria

The following requirements for VT path layer performance monitoring apply to all SONET NEs that terminate the VT path layer, with the possible exception of SONET NEs supporting UPSR configurations (see GR-1400-CORE).

-
- R6-368 [674]** A SONET NE providing VT PTE functions shall accumulate CV-Vs, ES-Vs, SES-Vs, UAS-Vs and FC-Vs for each terminated VT path.

- R6-369 [675]** A SONET NE providing VT PTE functions shall provide thresholding for the CV-V, ES-V, SES-V and UAS-V parameters.
- CR6-370 [1088]** VT PTE that processes the VT pointer from an incoming SONET signal (i.e., VT PTE that terminates a path that has not been reconciled to the local clock by upstream STS PTE within the same NE) may be required to support the capability to accumulate the VT-PJ-related parameters defined in [Section 6.2.2.6.1](#).

Note that although they are included for consistency with the case where the VT-PJ-related parameters are accumulated (as intermediate-path PM parameters) at STS PTE, the PPJC-VGen, NPJC-VGen and PJCS-VGen parameters are expected to always be equal to zero when they are accumulated at VT PTE.

- R6-371 [1089]** If the accumulation of the VT-PJ-related parameters is supported, the user shall be able to activate that accumulation on a per-VT-path basis. The default setting shall be “not active”.
- R6-372 [1090]** A SONET NE that is accumulating VT-PJ-related parameters shall perform thresholding for the PPJC-VDet, NPJC-VDet, PPJC-VGen, NPJC-VGen, PJCDiff-V, PJCS-VDet and PJCS-VGen parameters.
- R6-373 [676]** A SONET NE providing VT PTE functions shall provide the capability to accumulate, on a per-VT-path basis, the CV-VFE, ES-VFE, SES-VFE, UAS-VFE and FC-VFE parameters, and to activate and deactivate the accumulation of these parameters (as a group). The default setting shall be “not active.”
- R6-374 [677]** A SONET NE that is accumulating the far-end VT path PM parameters shall perform thresholding for the CV-VFE, ES-VFE, SES-VFE and UAS-VFE parameters.
-

6.2.2.7 Monitoring at DS_n Interfaces

GR-820-CORE defines DS_n path and line PM parameters.

- R6-375 [678]** A SONET NE shall provide DS_n path PM for each DS_n path that it terminates.
- R6-376 [679]** A SONET NE shall provide DS_n line PM for each DS_n line that it terminates.
- R6-377 [680]** A SONET NE shall meet the accumulation and thresholding requirements in GR-820-CORE for DS_n line and path PM parameters.
- R6-378 [681]** A SONET NE shall meet the OS reporting and data retrieval requirements in GR-820-CORE for its DS_n PM parameters.
-

Note that the accumulation of DS_n path PM where the DS_n path is not terminated is considered to be intermediate-path PM, which is covered in [Section 6.2.2.9](#).

6.2.2.8 PM During Troubles

The criteria concerning PM during troubles depend on the entity being monitored and the particular PM parameter. In general, the accumulation of a parameter is either inhibited during unavailable time, inhibited during unavailable time and during severely errored seconds, or never inhibited. For example, the accumulation of the ES and SES parameters is inhibited during periods of unavailability of the monitored entity. This is accomplished by not incrementing the current 15-minute and current day registers during unavailable seconds, where unavailable seconds are defined in [Sections 6.2.2.4, 6.2.2.5 and 6.2.2.6](#) (for SONET Lines, STS paths and VT paths). The accumulation of CVs is also inhibited during unavailable time; however, it is also inhibited for all seconds that are considered to be SESs for that entity (i.e., seconds during which certain defects are present, or during which K or more CVs are detected).

The accumulation of several PM parameters is never inhibited. For example, the accumulation of UASs and FCs (where they are defined) is not inhibited because they keep track of unavailable time and failure counts for the monitored entity. In addition, the accumulation of several of the SONET Section PM parameters is never inhibited because a UAS parameter is not defined at that layer.

Note that with these definitions a CV positive adjustment register is still needed to keep track of the number of CVs that occur (at a sub-SES rate) during the 10 seconds it takes to get out of unavailable time. In addition, those 10 seconds could easily span two periods so the NE may need to keep track of how many of the CVs occurred in each period, and possibly send a TCA for the previous period if the adjustment puts that period at or over the threshold (see [Section 6.2.2.1](#)).

R6-379 [683v4] A SONET NE shall inhibit the accumulation of near-end SONET PM parameters according to the rules listed below (and as summarized in [Tables 6-13 through 6-19](#)).

- The accumulation of all near-end Line layer parameters except for UASs, FCs, PSCs and PSDs (if applicable) shall be inhibited during periods of unavailability of the monitored line.
- The accumulation of all near-end STS or VT layer parameters except for UASs, FCs and the PJ-related parameters shall be inhibited during periods of unavailability of the monitored path.
- The accumulation of near-end CVs for a particular entity shall be inhibited for all seconds that are counted as SESs for that entity.

Note that consistent with the individual parameter specifications contained in ANSI T1.231.04 (but not with the generic specifications in ANSI T1.231), the PJ-related parameters are now (as of Issue 4) listed as exceptions in the second bullet of **R6-379 [683v4]**. Therefore, the accumulation of those parameters for a particular STS or VT path is not inhibited during periods of unavailability for that path. Note however, that in some cases the condition that causes a path to be considered unavailable (e.g., an extended AIS) may disrupt the corresponding pointer bytes, and in such cases no changes in the PJ-related parameter counts can occur until the pointer bytes are again available to be monitored.

Far-end PM parameters are derived from information carried in the incoming signal, and therefore can only be properly accumulated when that information is available. Therefore, the accumulation of all far-end parameters is also inhibited for seconds during which certain near-end defects are present.

R6-380 [995v2] A SONET NE shall inhibit the accumulation of far-end SONET PM parameters according to the rules listed below (and as summarized in Tables 6-13 through 6-19).

- The accumulation of all far-end parameters except for far-end UASs and far-end FCs shall be inhibited during periods of unavailability of the Line, STS path or VT path at the far end.
- The accumulation of far-end CVs for a particular entity shall be inhibited for all seconds that are counted as (far-end) SESs for that entity.
- The accumulation of all far-end Line parameters shall be inhibited during seconds in which a near-end AIS-L defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2) is present. An invalid-data flag shall be set for the inhibited parameters.
- The accumulation of all far-end STS path parameters shall be inhibited during seconds in which a near-end AIS-P defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2), a near-end LOP-P defect, or a near-end UNEQ-P defect is present. An invalid-data flag shall be set for the inhibited parameters.
- The accumulation of all far-end VT path parameters shall be inhibited during seconds in which a near-end AIS-V defect (or a lower-layer, traffic-related, near-end defect, see Section 6.2.1.8.2), a near-end LOP-V defect, or a near-end UNEQ-V defect is present. An invalid-data flag shall be set for the inhibited parameters.

R6-381 [682] A SONET NE that provides DS_n performance monitoring shall meet the requirements in GR-820-CORE for DS_n PM during troubles.

The following abbreviations are used in Tables 6-13 through 6-19 to indicate if the accumulation of a parameter continues when the listed defect is present and during unavailable time.

- y - Indicates the parameter shall continue to be accumulated during seconds in which the defect is present.
- N - Indicates the accumulation of the parameter shall be inhibited for all seconds during which the defect is present, as well as during periods of unavailability.
- n - Indicates the parameter shall be accumulated during seconds during which the defect is present, until unavailable time is declared for the given entity. When unavailable time is declared, the accumulation of the parameter shall be inhibited, back to the point when unavailable time began.
- 0 - Indicates the accumulation of the (far-end) parameter shall be inhibited for each second during which the (near-end) defect is present on the incoming signal. In

such cases, the parameter shall be marked as invalid for the entire accumulation interval.

Note that only defects that are detected at, or below the layer where a parameter is accumulated can affect the accumulation of that parameter. For example, the accumulation of Section PM parameters is not affected by an AIS-L defect, which is detected at the Line layer. Therefore, AIS-L is not listed in Table 6-13 on Section layer PM accumulation, but is listed in Tables 6-14 through 6-19 on Line, STS path and VT path layer PM accumulation. Finally, RDI defects do not affect the accumulation of near-end PM parameters, and therefore are not listed in Tables 6-13 through 6-16.

Table 6-13 Section Layer PM Accumulation During Defects

(Near-End) Section Parameter	Defect ^a
	LOS or SEF
SEFS-S	y ^b
CV-S	N
ES-S	y
SES-S	y

Notes:

- a. Defects detected at the Line and path layers do not affect the accumulation of Section PM parameters.
- b. See page 6-125 for definitions of y and N.

Table 6-14 Near-End Line Layer PM Accumulation During Defects

Near-End Line Parameter	Defect ^a	
	LOS, LOF, or AIS-L ^b	LOP-P, or all-ones STS pointer relay ^c
CV-L	N ^d	y
ES-L	n	y
SES-L	n	y
UAS-L	y	y
FC-L	y	y
PSC-L	y	y
PSD-L	y	y

Notes:

- a. RDI-L defects and defects detected at the path layers do not affect the accumulation of near-end Line PM parameters.
- b. In the functional model used in this document, LOS and LOF defects are detected at STE, which then generates AIS-L downstream. This AIS-L is detected at the LTE and affects the accumulation of Line PM parameters.
- c. Although LTE that processes STS pointers detects LOP-P defects and performs all-ones STS pointer relay, neither of those processes affect the accumulation of near-end Line PM parameters.
- d. See page 6–125 for definitions of y, N and n.

Table 6-15 Near-End STS Path Layer PM Accumulation During Defects

Near-End STS Path Parameter	Defect ^a		
	LOS, LOF, AIS-L, AIS-P ^b , LOP-P, UNEQ-PC ^c or TIM-PC	UNEQ-PC ^c , TIM-PC ^c , or PLM-P	LOP-V, or all-ones VT pointer relay ^d
CV-P	N ^e	y	y
ES-P	n	y	y
SES-P	n	y	y
UAS-P	y	y	y
FC-P	y	y	y
STS-PJ-related parameters	y ^f	y	y

Notes:

- a. RDI-L and RDI-P defects, and defects detected at the VT path layer do not affect the accumulation of near-end STS path PM parameters.
- b. In the functional model used in this document, defects detected at STE or LTE (e.g., LOS) result in AIS-P being generated downstream. This AIS-P is detected at the STS PTE and affects the accumulation of STS path PM parameters.
- c. If the STS PTE monitoring the path supports ERDI-P for that path, then UNEQ-P and TIM-P (if activated) are included in the second column of this table (along with LOS, LOF, etc.). On the other hand, if one-bit RDI-P is supported, then UNEQ-P and TIM-P are included in the third column (with PLM-P).
- d. Although STS PTE that processes VT pointers detects LOP-V defects and performs all-ones VT pointer relay, neither of those processes affect the accumulation of near-end STS path PM parameters.
- e. See page 6–125 for definitions of y, N and n.
- f. See the discussion on page 6–124 regarding STS-PJ-related parameter accumulation during troubles.

Table 6-16 Near-End VT Path Layer PM Accumulation During Defects

Near-End VT Path Parameter	Defect ^a	
	LOS, LOF, AIS-L, LOP-P, AIS-P, UNEQ-P, TIM-P, PLM-P, AIS-V ^b , LOP-V, UNEQ-V ^c , or TIM-V ^c	UNEQ-V ^c , TIM-V ^c or PLM-V
CV-V	N ^d	y
ES-V	n	y
SES-V	n	y
UAS-V	y	y
FC-V	y	y
VT-PJ-related parameters	y ^e	y

Notes:

- a. RDI defects do not affect the accumulation of near-end VT path PM parameters.
- b. In the functional model used in this document, defects detected at STE, LTE or STS PTE (e.g., LOS) result in AIS-V being generated downstream. This AIS-V is detected at the VT PTE and affects the accumulation of VT path PM parameters.
- c. If the VT PTE monitoring the path supports ERDI-V for that path, then UNEQ-V and TIM-V are included in the second column of this table (along with LOS, LOF, etc.). On the other hand, if one-bit RDI-V is supported, then UNEQ-V and TIM-V are included in the third column (with PLM-V).
- d. See page 6–125 for definitions of y, N and n.
- e. See the discussion on page 6–124 regarding VT-PJ-related parameter accumulation during troubles.

Table 6-17 Far-End Line Layer PM Accumulation During Defects

Far-End Line Parameter	Defect ^a		
	LOS, LOF, or AIS-L ^b	RDI-L	LOP-P, or all-ones STS pointer relay ^c
CV-LFE	0 ^d	N	y
ES-LFE	0	n	y
SES-LFE	0	n	y
UAS-LFE	0	y	y
FC-LFE	0	y	y

Notes:

- a. Defects detected at the path layers do not affect the accumulation of far-end Line PM parameters.
- b. In the functional model used in this document, LOS and LOF defects are detected at STE, which then generates AIS-L downstream. This AIS-L is detected at the LTE and affects the accumulation of Line PM parameters.
- c. Note that although all-ones STS pointer relay is performed by LTE that processes STS pointers, it is not defined as a defect and does not affect the accumulation of far-end Line PM parameters.
- d. See page 6–125 for definitions of y, N, n and 0.

Table 6-18 Far-End STS Path Layer PM Accumulation During Defects

Far-End STS Path Parameter	Defect ^a				
	LOS, LOF, AIS-L, AIS-P ^b , LOP-P, or UNEQ-P	RDI-L, TIM-P, or PLM-P	One-bit RDI-P, ERDI-P Server, or ERDI-P Connectivity ^c	ERDI-P Payload	LOP-V, or all-ones VT pointer relay ^d
CV-PFE	0 ^e	y	N	y	y
ES-PFE	0	y	n	y	y
SES-PFE	0	y	n	y	y
UAS-PFE	0	y	y	y	y
FC-PFE	0	y	y	y	y

Notes:

- a. Defects detected at the VT path layer do not affect the accumulation of far-end STS path PM parameters.
- b. In the functional model used in this document, defects detected at STE or LTE (e.g., LOS) result in AIS-P being generated downstream. This AIS-P is detected at the STS PTE and affects the accumulation of STS path PM parameters.
- c. See [Section 6.2.1.3.2](#) for information on the various RDI-P defects. Also note that the effect that an UNEQ-P or TIM-P defect detected by the far-end STS PTE will have on the accumulation of far-end STS path layer PM parameters (by the near-end STS PTE) depends on whether the near-end NE supports one-bit or enhanced RDI-P, and also on which RDI-P is supported by the far-end NE.
- d. Note that although all-ones VT pointer relay is performed by STS PTE that processes VT pointers, it is not defined as a defect and does not affect the accumulation of far-end STS path PM parameters.
- e. See page 6–125 for definitions of y, N, n and 0.

Table 6-19 Far-End VT Path Layer PM Accumulation During Defects

Far-End VT Path Parameter	Defect			
	LOS, LOF, AIS-L, LOP-P, AIS-P, UNEQ-P, TIM-P, PLM-P, AIS-V ^a , LOP-V, or UNEQ-V	RDI-L, RDI-P, TIM-V, or PLM-V	One-bit RDI-V, ERDI-V Server, or ERDI-V Connectivity ^b	ERDI-V Payload
CV-VFE	0 ^c	y	N	y
ES-VFE	0	y	n	y
SES-VFE	0	y	n	y
UAS-VFE	0	y	y	y
FC-VFE	0	y	y	y

Notes:

- a. In the functional model used in this document, defects detected at STE, LTE or STS PTE (e.g., LOS) result in AIS-V being generated downstream. This AIS-V is detected at the VT PTE and affects the accumulation of VT path PM parameters.
- b. See [Section 6.2.1.3.3](#) for information on the various RDI-V defects. Also note that the effect that an UNEQ-V or TIM-V defect detected by the far-end VT PTE will have on the accumulation of far-end VT path layer PM parameters (by the near-end VT PTE) depends on whether the near-end NE supports one-bit or enhanced RDI-V, and also on which RDI-V is supported by the far-end NE.
- c. See page 6–125 for definitions of y, N, n and 0.

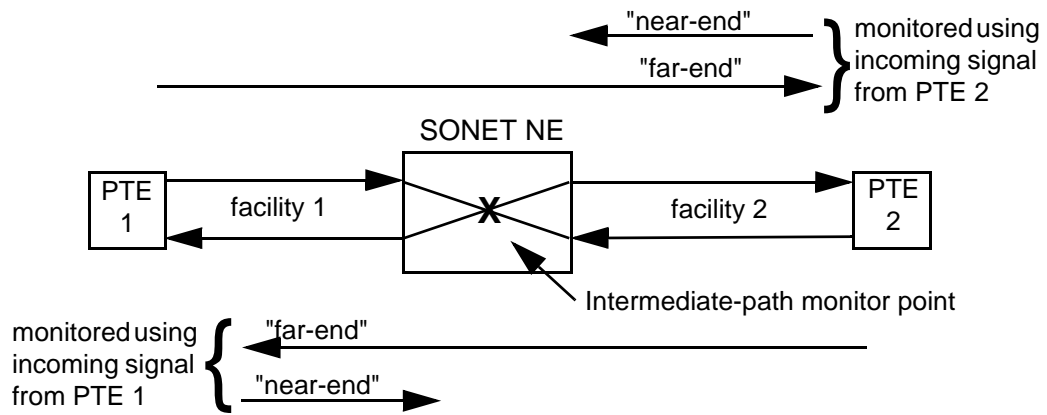
6.2.2.9 Intermediate-Path PM

Intermediate-path PM is defined to be the transparent monitoring of a constituent channel of an incoming transmission signal by an NE that, in most cases,²⁶ does not terminate that channel. For SONET NEs, intermediate-path PM is defined for the STS and VT path layers, and for DS_n paths. A SONET NE performing non-PJ-related intermediate-path PM examines the overhead in the monitored path and derives all of the near-end and far-end path PM parameters in each direction of transmission, while allowing the path signal to pass bidirectionally through the NE completely unaltered. The NE either frames on the incoming targeted path or uses the appropriate pointers to locate the appropriate overhead. This type of intermediate-path PM is useful (for example) when a monitored path is never terminated within a particular network provider’s equipment. This is particularly true for high-capacity line services that promise certain levels of performance in terms of PM parameters such as ESs, SESs and UASs. In many cases, non-PJ-related intermediate-path PM is expected to be performed where SONET and asynchronous networks meet.

26. One exception would be an NE that supports the accumulation of PJ-related intermediate-path PM parameters at LTE for an STS path that is subsequently terminated at STS PTE within the NE, or at STS PTE for a VT path that is subsequently terminated at VT PTE within the NE.

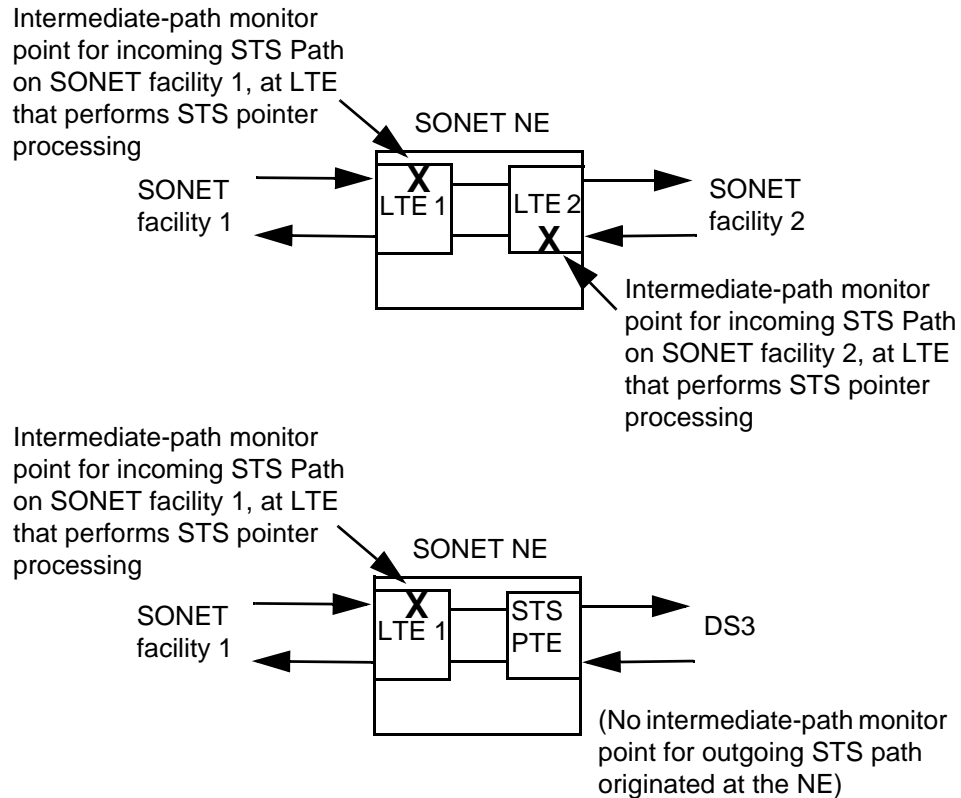
Figure 6-21 illustrates how non-PJ-related intermediate-path PM parameters are derived. In the figure, a bidirectional path is cross-connected through the “SONET NE”, and has path terminations at PTE 1 and PTE 2. As shown in the figure, four sets of PM parameters are accumulated when intermediate-path PM is activated. For the path signal received on facility 1, there are the near-end parameters monitored on the path from PTE 1 to the SONET NE and the far-end parameters monitored via the path REI and RDI signals carried on the path from PTE 1 to PTE 2. For the path signal received on facility 2, there are the near-end parameters monitored on the path from PTE 2 to the SONET NE and the far-end parameters monitored via the path REI and RDI signals carried on the path from PTE 2 to PTE 1.

Figure 6-21 Intermediate-Path PM for Non-PJ-Related Parameters



Unlike the non-PJ-related parameters discussed above, the accumulation of the PJ-related parameters is affected by both the contents of the incoming signal and the actions of the LTE or STS PTE that processes the monitored path’s pointers (i.e., that generates increments or decrements to frequency justify the SPE to a local clock). Therefore those parameters are typically accumulated only for the incoming path signal at the LTE or STS PTE that performs the pointer processing. This is illustrated in Figure 6-22 for two types of STS paths (i.e., nonterminated and terminated). Note that in Figure 6-22 there is only one incoming SONET signal shown at each of the NEs’ SONET “interfaces”, but that additional incoming signals may be present if the NE supports Line APS (e.g., linear APS or the 4-fiber BLSR architecture). In those cases it is generally expected that an NE will provide the capability to accumulate PJ-related intermediate-path PM for any of the STS paths that make up each of the working (and possibly extra-traffic) channels, independent of which lines those channels are being selected from.

Figure 6-22 Examples of Intermediate-Path PM for STS PJ-Related Parameters



Two levels of intermediate-path PM are defined. Level-1 intermediate-path PM occurs on an incoming target path when the NE does not need to perform any additional demultiplexing of the incoming signal beyond that needed to pass the signal through the NE. Level-2 intermediate-path PM occurs on an incoming target path when a “container” path must be transparently terminated/demultiplexed on the incoming signal to monitor the desired target path.

Consistent with the description of level-1 and level-2 intermediate-path PM, two types of paths may exist at a particular NE. A level-1 path is defined to be a path that is directly visible at the NE, either at its external interfaces, or in its cross-connect fabric. A level-2 path is defined to be a path carried within a level-1 path. At a given NE, there might be several different types of level-1 and level-2 paths, depending on the types of interfaces that it has and the types of cross-connections that it supports.

The combinations of path types and intermediate-path PM accumulation capabilities that might apply at a given SONET NE are as follows.

- Level-1 Intermediate DS1 path PM can occur when a DS1 is asynchronously mapped into an (originated) VT1.5 path, and when a DS1 is asynchronously demultiplexed from a (terminated) VT1.5 path.
- Level-2 Intermediate DS1 path PM can occur when a VT1.5 path carried on an incoming SONET signal and containing an asynchronously mapped DS1 is cross-connected to an outgoing SONET signal.

- Level-1 Intermediate DS3 path PM can occur when a DS3 is asynchronously mapped into an (originated) STS-1 path, and when a DS3 is asynchronously demultiplexed from a (terminated) STS-1 path.
- Level-2 Intermediate DS3 path PM can occur when an STS-1 path carried on an incoming SONET signal and containing an asynchronously mapped DS3 is cross-connected to an outgoing SONET signal.
- Level-1 Intermediate VT path PM (Non-PJ-Related Parameters) can occur when a VT path carried on an incoming SONET signal is cross-connected to an outgoing SONET signal.
- Level-1 Intermediate VT path PM (PJ-Related Parameters) can occur when STS PTE processes VT pointers to frequency justify (to a local clock) a VT path carried on an incoming SONET signal and subsequently cross-connected to an outgoing SONET signal or terminated at VT PTE within the same NE.
- Level-2 Intermediate VT path PM (Non-PJ-Related Parameters) can occur when a VT-structured STS-1 path carried on an incoming SONET signal is cross-connected to an outgoing SONET signal.
- Level-1 Intermediate STS path PM (Non-PJ-Related Parameters) can occur when an STS path carried on an incoming SONET signal is cross-connected to an outgoing SONET signal.
- Level-1 Intermediate STS path PM (PJ-Related Parameters) can occur when LTE processes STS pointers to frequency justify (to a local clock) an STS path carried on an incoming SONET signal and subsequently cross-connected to an outgoing SONET signal or terminated at STS PTE within the same NE.

CR6-382 [685v2] A SONET NE may be required to support non-PJ-related level-1 intermediate-path PM for some or all types of level-1 paths in the NE.

CR6-383 [1091] A SONET NE containing STS PTE that processes VT pointers to frequency justify (to a local clock) VT paths carried on an incoming SONET signal may be required to support the capability to accumulate VT-PJ-related level-1 intermediate-path PM parameters for those paths.

O6-384 [1092] A SONET NE containing LTE that processes STS pointers to frequency justify (to a local clock) STS paths carried on an incoming SONET signal should support the capability to accumulate STS-PJ-related level-1 intermediate-path PM parameters for those paths.

CR6-385 [686] A SONET NE may be required to support level-2 intermediate-path PM for some or all level-2 path types in the NE.

R6-386 [687v2] A SONET NE that provides a given type of intermediate-path PM for a particular type of level-1 path shall provide the capability to monitor any of the level-1 paths of that type passing through the NE.

R6-387 [688v2] A SONET NE that provides intermediate-path PM for a particular type of level-2 path shall provide the capability to monitor any of the level-2 paths of that type passing through the NE.

In general, any criteria specifying that an NE must be capable of simultaneously monitoring a certain percentage or number of a particular type of path will appear in the appropriate NE-specific. However, the following requirement applies to all NEs that support intermediate-path PM.

R6-388 [1093] If a SONET NE supports level-1 or level-2 intermediate-path PM for a particular type of path, then the number or percentage of the paths of that type that can be simultaneously monitored shall be clearly documented.

R6-389 [689v2] A SONET NE that supports a given type of intermediate-path PM for a given path type shall provide the user the ability to activate the feature on a per-path basis. The default shall be “not active” for all paths.

O6-390 [1094] A SONET NE that provides the capability to accumulate both PJ-related and non-PJ-related intermediate-path PM parameters for an STS or VT path should allow the user to independently activate and deactivate the accumulation of those two sets of parameters.

R6-391 [690] A SONET NE that provides intermediate-path PM for a given path shall not alter any bits, overhead or payload, of the monitored entity in either direction of transmission.

Note that **R6-391 [690]** means an NE that provides level-2 intermediate-path PM would need to perform a non-intrusive “termination/demultiplexing” of the nonterminated container path.

R6-392 [691v4] If a non-PJ-related intermediate-path PM feature is active for a particular path, the NE shall perform both near-end and (if defined) far-end path PM for the path signals in both directions. Except as noted below, the monitoring shall be performed as if the NE were terminating the path in each direction, using the criteria in [Sections 6.2.2.5, 6.2.2.6 and 6.2.2.7](#) (for STS, VT and DS_n paths).

R6-393 [1095v2] If a PJ-related intermediate-path PM feature is active for a particular path, the NE shall monitor that path in the direction specified. The monitoring shall be performed as if the NE were terminating the path, using the criteria in [Sections 6.2.2.5 and 6.2.2.6](#) (for STS and VT paths).

In the preceding requirements, “monitoring” includes both the accumulation of PM parameters and the generation of TCAs for those parameters. Note that for an NE to perform DS₃ intermediate-path PM “as if it were terminating the path,” it would need to be capable of being provisioned to expect M23 or C-bit parity application signals. That is not a capability that is typically provided in SONET NEs that transport DS₃s, and therefore it is considered sufficient for the NE to accumulate DS₃ intermediate-path PM according to the criteria for M23 applications (for which

only near-end PM parameters are defined, and for which the bit error information is provided by the P-bits). The NE may (but is not required to) allow the user to provision it to expect a C-bit parity application DS3 signal. In that case, the NE would be required to accumulate both the near-end and far-end DS3 intermediate-path PM data according to the criteria for C-bit parity applications (assuming that it is a bidirectional path).

Also note that although TIM defects may affect or contribute to the accumulation of various PM parameters at STS or VT PTE (e.g., see the parameter definitions in [Sections 6.2.2.5.1](#) and [6.2.2.6.1](#)), due primarily to provisioning issues they are generally not expected to be detected at non-PTE that performs intermediate STS or VT path PM, or level-2 intermediate DS_n path PM. Therefore, the detection of TIM defects is currently an allowed exception to the “as if the NE were terminating the path” statement in **R6-392 [691v4]**, and the possible impacts (both positive and negative) need to be carefully considered before the detection of TIM defects for a path is activated. For example, if the detection of TIM-P defects is activated for a particular STS path, those defects could cause differences between the set of near-end PM data accumulated at an intermediate NE that is performing STS intermediate-path PM but does not support TIM-P monitoring and the set of near-end PM data accumulated at the downstream NE where the path is terminated (or the sets of far-end PM data accumulated at any NEs monitoring the upstream path).

-
- O6-394 [693]** An NE that is capable of performing intermediate-path PM for a particular type of level-2 path should allow the user to provision whether it also performs level-1 intermediate-path PM on the container paths.
- R6-395 [694v4]** An NE that is performing non-PJ-related intermediate-path PM for a particular path shall (non-intrusively) detect and terminate AIS, LOP and RDI defects (and for DS_n paths, DS_n OOF) as if it were terminating that path.
- R6-396 [1034v3]** An NE that is performing non-PJ-related intermediate-path PM for a particular SONET path, and that supports the detection of ERDI defects for that path, shall (non-intrusively) detect and terminate UNEQ defects as if it were terminating the path.
- R6-397 [1127v2]** An NE that is performing non-PJ-related intermediate-path PM for a particular SONET path, and that supports the detection of ERDI defects and the capability to detect TIM defects on that path shall allow the user to provision whether it monitors for the TIM defects. In addition, if the NE has been provisioned to monitor for TIM defects on the path, it shall (non-intrusively) detect and terminate those defects as if it were terminating that path.
-

Note that as the various PM parameters are currently defined, the detection of a PLM-P defect does not affect the accumulation of any of the STS path PM parameters, and the detection of a PLM-V defect does not affect the accumulation of any of the VT path PM parameters. Therefore, it is not necessary for an NE that is performing intermediate-path PM on an STS or VT path to detect and terminate PLM

defects on that path. However, PLM defects detected on a container path do affect the accumulation of the target path PM parameters. Therefore, those defects are included in the following requirement.

R6-398 [996] A SONET NE that is performing level-2 intermediate-path PM for a particular path shall (non-intrusively) detect and terminate AIS, LOP, UNEQ, PLM and RDI defects for the container path as if it were terminating that path.

R6-399 [1128v2] An NE that is performing level-2 intermediate-path PM for a particular path and supports the capability to detect TIM defects on that path's container path shall allow the user to provision whether it monitors for those defects. In addition, if the NE has been provisioned to monitor for TIM defects on the container path, it shall (non-intrusively) detect and terminate those defects as if it were terminating that path.

As discussed in [Section 6.2.1.2.4](#), when an AIS, LOP, UNEQ, PLM or (possibly) TIM defect is detected on a terminated STS-1 or VT path carrying an asynchronously mapped DS_n, the SONET NE is required to generate DS_n AIS downstream. In addition, the detection of DS_n AIS defects affect the accumulation of DS_n path PM parameters (see GR-820-CORE). Therefore, when level-2 intermediate-path PM is being performed on a DS_n path, the detection of an AIS, LOP, UNEQ, PLM or (possibly) TIM defect on the container path must affect the accumulation of the PM parameters in the same way that the detection of DS_n AIS would.

Several of the preceding criteria (e.g., **R6-395 [694v4]** and **R6-398 [996]**) use the phrase “an NE that is performing”, which means they are applicable when an intermediate-path PM feature is active for the particular path. Note that it is also acceptable for an NE that is capable of performing intermediate-path PM on a particular path, but that is currently not provisioned to do so, to (non-intrusively) detect and terminate defects as if the intermediate-path PM feature were active.

Finally, see [Section 6.2.1.8.7](#) for criteria concerning the declaration and clearing of failures based on the defects detected and terminated according to the preceding criteria.

6.2.3 Testing Process

Testing deals with procedures that result in isolation of a failure to a replaceable or repairable entity. Maintenance tools that achieve this isolation, besides those built into the SONET signal format, include test access, diagnostics and loopbacks.

The long-term objective for testing is to evolve toward self-diagnosing NEs. Steps toward this goal are effected in the alarm and PM plans. Another step is reflected in the requirements for diagnostics, discussed in [Section 6.2.3.2](#). Testing activities include:

- Analyzing alarms, PM data and maintenance signals
- Executing diagnostics
- Executing controls, such as switching to protection
- Activating loopbacks

- Test access for signal measurements.

In each of these activities, operations personnel gain access through either a local craftsperson interface or the remote operations interface.

6.2.3.1 Test Access

Test access allows access to a signal for the purposes of non-intrusive monitoring and intrusive testing. For SONET NEs, this access falls into three possible categories: access to the fiber for monitoring and testing the optical signal and the fiber; access to the SONET signal for monitoring and testing the format, mapping, and equipment specifications; and digital test access to lower-speed digital signals to test the service.

6.2.3.1.1 Fiber Access

GR-1309-CORE, *TSC/RTU and OTAU Generic Requirements for Remote Optical Fiber Testing*, describes functional requirements for fiber testing.

-
- R6-400 [706]** Access to the fiber for fiber testing (e.g., for identifying the location of a fiber break) shall be provided.
-

This fiber access currently implies disconnecting the optical source or receiver.

-
- O6-401 [707]** To aid in the remote sectionalization of a problem to the fiber or to one of the terminals, the ability to switch the fiber to an alternate source or an alternate receiver should be provided.
-

6.2.3.1.2 SONET Signal Test Access

This section provides the following objective for local test access.²⁷

-
- O6-402 [709]** A SONET NE should provide local test access to individual STS-1s and STS-3s within an OC-N or STS-N electrical signal.
-

If an NE provides local test access, the following requirements apply.

-
- R6-403 [710]** The SONET electrical level test access shall be in accordance with the specifications defined in [Section 4.4](#) for STS-1 and STS-3 electrical signals.

- R6-404 [711]** The SONET electrical level test access shall provide a non-intrusive and hitless monitor mode.

²⁷The support of remotely controlled access and testing for OC-N signals, and requirements for remotely controlled access and testing of SONET electrical interfaces, are FFS.

- R6-405 [712]** The SONET electrical level test access shall provide the ability to perform intrusive tests.
- R6-406 [713]** When intrusive testing occurs, the NE shall provide the appropriate path AIS in the non-test direction.
- R6-407 [714]** In NEs supporting APS, the test access shall not impair the working channel on the protection line.
-

6.2.3.1.3 *Digital Test Access*

SONET NEs that are VT (or DS1) programmable (i.e., that allow the flexible assignment of low-speed signals to timeslots within an OC-N signal) may provide DS1 remote test access.

- CR6-408 [715]** A SONET NE that is VT programmable may be required to provide DS1 remote test access.
-

Generic criteria on testing are described in GR-818-CORE, *Network Maintenance: Access and Testing - Generic Test Architecture*, and GR-819-CORE, *Network Maintenance: Access and Testing - Special Services (SS) and SS-Like Networks*. GR-818-CORE gives the Telcordia view of a generic test architecture for circuit testing, and provides the framework for proposed generic requirements for Test System Controller/Remote Test Units (TSC/RTUs), Metallic Test Access Units (MTAUs) and Digital Test Access Units (DTAUs). GR-819-CORE discusses access and testing of Special Services (SS) and SS-like networks. In particular, it contains proposed generic criteria for DS1 test access.

CR6-408 [715] may be changed to a requirement in GRs or TRs covering specific types of SONET NEs. If a SONET NE provides DS1 remote test access, the following requirements apply.

- R6-409 [1129]** The DS1 test access shall meet the test architecture criteria in GR-818-CORE.
- R6-410 [1130]** A facility loopback capability shall be provided for all test access ports as defined in GR-819-CORE.
- R6-411 [1131]** The DS1 test access capability shall support the monitor access and the split access capabilities as defined in GR-819-CORE.
- R6-412 [717]** When TL1 interfaces are used, SONET NEs shall use the TL1 messages of GR-834-CORE, *OTGR Section 12.4: Network Maintenance: Access and Testing Messages*, for any test access functions provided.

- R6-413 [718]** When CMISE interfaces are used, SONET NEs shall adhere to the object model and use the CMISE service mappings of GR-1031-CORE, *OTGR Section 15.6: Operations Interfaces Using OSI Tools: Test Access Management*, for any test access functions provided.
-

6.2.3.2 Diagnostics

The term diagnostic, as used in this section, has several connotations. For example, there are those diagnostics that the equipment manufacturer provides for internal troubleshooting. Some diagnostics in this category may run continuously, some may run on a designed schedule (e.g., once an hour or triggered by an event), and some may run on a schedule the user is permitted to specify. Typically, these routine diagnostics are design dependent, and during normal operation the user is unaware of them. The ability to manually initiate routine diagnostics between the scheduled times is referred to as “on demand”. A second sense in which the term diagnostics is used is as a user’s testing tool. These tools are functions that are activated on demand (initiated at a WS or OS interface) to retrieve information from the overhead or operate special circuitry to support trouble analysis (e.g., the corrupted BIP), or to run an NE diagnostic and see the result.

- R6-414 [719]** A SONET NE shall provide diagnostic capabilities. As a minimum, diagnostics shall be provided to detect the equipment failures listed in [Section 6.2.1.1.4](#). The diagnostic capabilities shall run routinely and on demand. When these diagnostics are run on demand, the NE shall provide the user with the results.
- O6-415 [720]** An NE should support additional diagnostics that provide the ability to isolate an equipment trouble to the replaceable circuit pack or module. These diagnostics should not interfere with working services and may be run routinely or on demand.
- R6-416 [721]** Diagnostics that interfere with service shall not run routinely unless permitted by the user.²⁸
-

In addition to these general diagnostics criteria, the following subsections provide criteria related to specific functions or layers supported by the SONET NE. In addition, Section 9.3 of GR-1244-CORE contains several criteria that apply if an NE utilizes automatic clock diagnostic routines that can adversely impact its synchronization performance (e.g., that cause phase transients on the NE’s output signals). Those criteria are applicable to all NEs with outputs that are synchronized to internal clocks, and thus apply to SONET NEs that contain SMCs or stratum 2, 3E or 3 clocks.

²⁸.It is expected that the supplier would provide a complete list of any diagnostics that interfere with service.

6.2.3.2.1 Physical Layer

As discussed in [Section 6.2.2.2.2](#), in some situations it may be useful for a user to be able to determine the current value of an optical transmitter's LBC or OPT, or an optical receiver's OPR. If an NE supports the Physical layer PM parameters defined in [Section 6.2.2.2.1](#) and meets **O6-347 [1084]**, then this capability can be provided as a part of its Physical layer PM feature. On the other hand, if the NE does not support one or more of the Physical layer PM parameters or does not meet **O6-347 [1084]**, then the following objectives apply.

-
- O6-417 [722v3]** An NE that does not support the OPR_{normal} PM parameter for an optical receiver should provide an on-demand diagnostic that reports the received optical power, preferably in units of dBm.
- O6-418 [723v3]** An NE that does not support the OPT_{normal} PM parameter for an optical transmitter should provide an on-demand diagnostic that reports the transmitted optical power, preferably in units of dBm.
- CR6-419 [724v3]** An NE that does not support the LBC_{normal} PM parameter for an optical transmitter may be required to provide an on-demand diagnostic that reports the laser bias current, preferably in microamperes (μ A).
-

Note that for the diagnostics described above to be of significant value, it would be important for the user to be able to easily determine whether the reported values are within their expected or specified ranges, as well as the approximate accuracy limitations of the measurements. Such information could be provided, for example, in the product documentation or as part of the NE's response to the command that initiated the diagnostic.

-
- O6-420 [1096]** An NE that supports the OPR_{normal} PM parameter but does not meet **O6-347 [1084]** should provide an on-demand diagnostic that reports the current value (not the snapshot value) of OPR_{normal} .
- O6-421 [1097]** An NE that supports the OPT_{normal} PM parameter but does not meet **O6-347 [1084]** should provide an on-demand diagnostic that reports the current value (not the snapshot value) of OPT_{normal} .
- O6-422 [1098]** An NE that supports the LBC_{normal} PM parameter but does not meet **O6-347 [1084]** should provide an on-demand diagnostic that reports the current value (not the snapshot value) of LBC_{normal} .
-

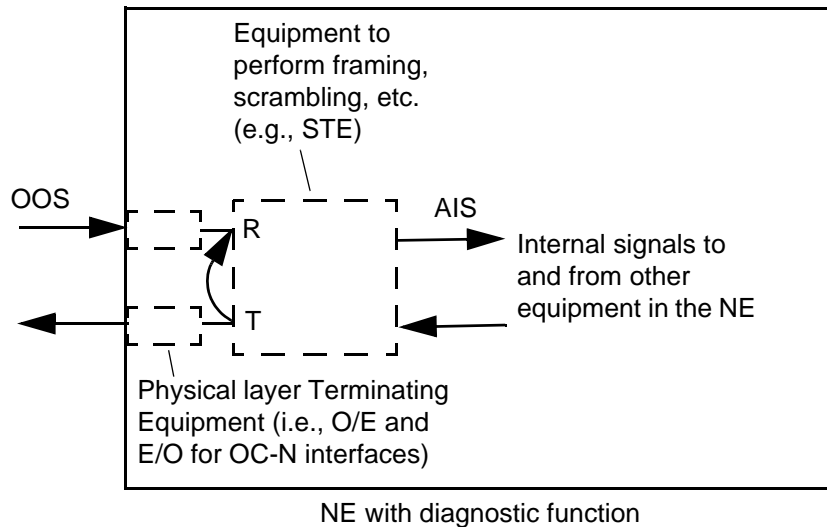
6.2.3.2.2 Section Layer

The criteria that appear in this section describe a diagnostic that involves looping back a SONET signal immediately before it is transmitted, to verify the integrity of the electronics associated with the receiver or the transmitted signal (see [Figure 6-23](#)). In Issue 1 of this document, these criteria appeared in [Section 6.2.3.3.1](#)

and the feature was identified as a SONET Terminal Loopback. However, the term “loopback” is generally associated with features that involve the use of external test equipment to monitor the looped back signal. The feature described here is strictly internal to the SONET NE at which it is performed, and therefore is considered to be a diagnostic.

-
- R6-423 [740v2]** The SONET NE shall provide the functional equivalent of the diagnostic illustrated in [Figure 6-23](#).
-

Figure 6-23 Section Layer Diagnostic



In this diagnostic, the outgoing signal is looped back to the associated receiver at a point as close to the physical interface as possible (e.g., immediately before the electrical-to-optical conversion in the case of an optical interface). Note that this will interrupt an incoming SONET signal, and therefore it must not be initiated by a command carried in the Section or Line DCC of that signal (because it may interrupt communications between the NE performing the loopback and the user). The use of a DCC in this manner could cause the diagnostic to be initiated without providing a means for it to be terminated. Also note that if the NE supports linear APS at the interface where the diagnostic is being performed, then the traffic could be selected from the line that is not affected by the diagnostic and no service interruption need occur.

-
- R6-424 [741]** An NE shall deny a request for this diagnostic if the incoming facility associated with the receiver is In_Service or Out_of_Service-Autonomous, as defined in GR-1093-CORE.
- R6-425 [743]** While the diagnostic is being performed, the NE shall recognize that it is in a test state, as defined in GR-1093-CORE, and it shall not take action to revert to a previous state.

- R6-426 [745]** The NE shall exit the test state, as defined in GR-1093-CORE, when the diagnostic is deactivated.
- R6-427 [744]** An NE shall notify the OS when the diagnostic has been activated and when it has been deactivated.
- R6-428 [746]** The NE shall determine if the receiver is able to frame on the delivered signal.
- R6-429 [747]** The NE shall detect errors in the Section BIP-8 (B1 byte), or if the B1 byte is not processed, in the Line BIP-8s (B2 bytes).
- R6-430 [748v2]** The NE shall report results indicating that the equipment could not frame, that CVs were detected, or that no framing problems or CVs were detected.
-

6.2.3.2.3 *Signal Identification*

This section contains criteria for diagnostics that allow a signal to be traced back to its source (e.g., for troubleshooting purposes). Diagnostics are defined that use the Section, STS path and VT path traces (J0, J1 and J2 bytes), and the STS and VT signal labels (C2 byte and V5 bits 5 through 7).

6.2.3.2.3.A *Section, STS Path and VT Path Trace*

The Section, STS path or VT path trace may be used for connectivity verification even in applications where the detection of TIM defects (see [Section 6.2.1.1.9](#)) is not used. The following criteria, along with the criteria in [Sections 3.3.2.1](#) and [6.2.1.1.9](#) related to the provisioning and transmission of Section and STS/VT path trace strings, are intended to support that use.

- O6-431 [1216]** STE should provide an on-demand diagnostic to detect and report the contents of the received Section trace.
- CR6-432 [1217]** STE may be required to provide a diagnostic that, when activated, continuously monitors the incoming Section trace.
- R6-433 [728]** STS PTE shall provide an on-demand diagnostic to detect and report the contents of the received STS path trace.
- O6-434 [730]** STS PTE should provide a diagnostic that, when activated, continuously monitors the incoming STS path trace.
- CR6-435 [729]** LTE with STS cross-connection capabilities may be required to provide an on-demand diagnostic to detect and report the contents of the STS path trace in the (nonterminated) STS path designated by the user.

CR6-436 [1218] LTE with STS cross-connection capabilities may be required to provide a diagnostic that, when activated, continuously monitors the incoming STS path trace.

O6-437 [1219] VT PTE should provide an on-demand diagnostic to detect and report the contents of the received VT path trace.

CR6-438 [1220] VT PTE may be required to provide a diagnostic that, when activated, continuously monitors the incoming VT path trace.

CR6-439 [1221] STS PTE with VT cross-connection capabilities may be required to provide an on-demand diagnostic to detect and report the contents of the VT path trace in the (nonterminated) VT path designated by the user.

CR6-440 [1222] STS PTE with VT cross-connection capabilities may be required to provide a diagnostic that, when activated, continuously monitors the incoming VT path trace.

In general, it is possible for a SONET NE that supports on-demand trace diagnostics to simply report the detected trace, or to also report whether the detected trace matches a user-provisioned “expected” path trace. Similarly, a SONET NE that supports a continuous trace monitoring diagnostic could compare the current incoming trace to either a previously received trace or to a user-provisioned expected trace.

CR6-441 [1223] A SONET NE may be required to support a feature to allow the user to provision, for diagnostics purposes, the expected trace for each SONET line-side signal that it terminates.

O6-442 [1224] If a SONET NE supports a feature that allows the user to provision an expected Section trace for diagnostics purposes, then the following apply:

- The feature should allow the user to enter the contents of the expected Section trace such that it meets the specifications in ANSI T1.269 related to 16-byte trace messages
- The feature should allow the user to enter the contents of the expected Section trace such that it consists of any single value from ‘00’ through ‘FF’ (hex).

CR6-443 [999v2] A SONET NE may be required to support a feature to allow the user to provision, for diagnostics purposes, the expected path trace for each STS path that it terminates.

R6-444 [1000v3] If an NE supports a feature that allows the user to provision an expected STS path trace for diagnostics purposes, then the following apply:

- The feature shall allow the user to enter a string of up to 62 ASCII printable characters
- The feature shall place no restriction on the contents of the string (other than that the characters shall be ASCII printable characters).

O6-445 [1225] If an NE supports a feature that allows the user to provision an expected STS path trace for diagnostics purposes, that feature should also allow the user to enter an expected path trace that meets the specifications in ANSI T1.269 related to 64-byte trace messages.

CR6-446 [1226] A SONET NE may be required to support a feature to allow the user to provision, for diagnostics purposes, the expected path trace for each VT path that it terminates.

R6-447 [1227] If an NE supports a feature to allow the user to provision an expected VT path trace for diagnostics purposes, that feature shall allow the user to enter the contents of the expected path trace such that it meets the specifications in ANSI T1.269 related to 16-byte trace messages.

Note that support of a feature that allows the user to provision the “expected” trace is required in [Section 6.2.1.1.9.A](#) for TIM-P defect detection purposes, and conditionally required in [Section 6.2.1.1.9.B](#) for TIM-V defect detection purposes. Also note that the use of an expected path trace feature is currently not recommended for STS or VT paths that are not terminated. The use of such a feature could result in excessive reprovisioning at intermediate NEs when the two end-points (e.g., the STS PTE) are changed.

R6-448 [731v4] A SONET NE that is continuously monitoring an incoming trace and does not support an expected trace feature for diagnostics purposes shall compare the incoming trace with a previously received path trace.

CR6-449 [1002v3] A SONET NE that is continuously monitoring an incoming trace and supports an expected trace feature for diagnostics purposes, but that is not provisioned with an expected trace, may be required to compare the incoming trace with a previously received trace.

R6-450 [1004v2] An NE that compares the incoming trace with a previously received trace shall be capable of performing that comparison on a byte-by-byte basis (e.g., independent of whether the contents are ASCII-based or contain <CR> and <LF> characters).

R6-451 [735v4] A SONET NE that is monitoring an incoming trace for changes or mismatches for diagnostics purposes shall suspend the monitoring function when the J0, J1 or J2 byte cannot be accessed (e.g., when an LOP-P or AIS-P defect has been detected by the STS PTE, or an LOP-V or AIS-V defect has been detected by the VT PTE). If the NE is monitoring for changes in the incoming trace, the contents of the trace before the disruption shall be used as the starting value following a restart.

R6-452 [736v3] A SONET NE that has suspended monitoring of an incoming trace because the J0, J1 or J2 byte could not be accessed shall resume monitoring when that byte can again be accessed.

- R6-453 [732v3]** A SONET NE that is monitoring for changes of the incoming trace shall detect when a sustained change in the trace content occurs. Upon detecting a sustained change, the NE shall send a message to an OS. The level of the message shall be Not Alarmed, and it shall include either the hexadecimal values (1-byte Section trace case) or ASCII-based portions (all multi-byte trace cases) of both the previously received trace and the new trace.
- R6-454 [1005v3]** A SONET NE that is monitoring for a mismatch between the incoming trace and an expected trace for diagnostics purposes shall detect when a sustained mismatch occurs. Upon detecting a sustained mismatch, the NE shall set an indication for that signal or path and send a message to an OS. The default level of the message shall be Not Alarmed, and it shall include either the hexadecimal values (1-byte Section trace case) or ASCII-based portions (all multi-byte trace cases) of both the expected trace and the new trace.
- R6-455 [1006v3]** A SONET NE that is monitoring the incoming trace for diagnostics purposes and that has detected a sustained mismatch shall detect when the incoming trace matches the expected trace. Upon detecting a match, the NE shall clear the indication for that signal or path and send a clear message to the OS (if the mismatch was reported to an OS).

A sustained change or mismatch of the Section, STS path or VT path trace is one where the new trace is consistently being received, as opposed to a short-term disruption or mismatch due to (for example) a burst of errors. Note that a change in the “starting position” of an incoming multi-byte trace is not considered a sustained change or mismatch. Such a change could be caused by (for example) an upstream protection switch where transmission delay differences cause a J1 byte to appear to be dropped or repeated in successive frames (so that one path trace string would appear to contain 63 or 65 bytes). Also note that **R6-453 [732v3]**, **R6-454 [1005v3]** and **R6-455 [1006v3]** do not specify a time period within which an NE must report that it has detected a change or mismatch. In general, the goal is for such events to be reported without excessive delays, but for the detection mechanism to be robust in the presence of performance degradations. Timing similar to that required for the defects and failures defined in [Section 6.2.1](#) (e.g., report a change or mismatch if it persists for approximately 2.5 seconds) would be appropriate.

- CR6-456 [1007v3]** An NE that monitors the incoming trace for mismatches for diagnostics purposes may be required to be user provisionable to report a detected mismatch as alarmed or Not Alarmed.

Note that if an NE is monitoring a particular trace for changes rather than mismatches, no condition exists that would cause it to send a clear message after it has reported a change (e.g., a detected change back to the old trace would cause another change to be reported, not a clear message). Therefore, a Section, STS path or VT path trace change cannot be alarmed. In addition, since a continuous trace monitoring diagnostic is a feature that must be activated by the user, it is assumed that the user will want any detected changes or mismatches to be autonomously reported. Therefore, it is not necessary for an NE to allow the user to provision a change or mismatch to be Not Reported.

6.2.3.2.3.B STS and VT Path Signal Label

STS and VT signal label diagnostics can be used to identify the construction of STS and VT SPEs for troubleshooting purposes.

-
- R6-457 [737]** A SONET NE shall provide an on-demand diagnostic to report the value currently being received in the STS and VT signal labels at STS PTE and VT PTE, respectively.
- R6-458 [1008]** A SONET NE that supports the detection of PDI-P defects (e.g., at an STS-level path selector in a UPSR NE) shall provide an on-demand diagnostic to report the value currently being received in the STS signal label contained in each STS path that it is monitoring.
- CR6-459 [1009]** An NE that supports the generation of PDI-P signals may be required to support an on-demand diagnostic that reports the code that is currently being inserted on the specified STS signal label.
-

6.2.3.2.4 Error Monitoring

-
- R6-460 [738]** A SONET NE shall provide the ability to transmit, on demand, a fully corrupted BIP value (all parity check bits inverted) for the Section, Line, STS path or VT path (as appropriate to the NE). The NE shall provide the user the capability to specify the approximate duration of the corrupted BIP diagnostic as some number of units of time or some number of frames (or VT superframes for a VT path BIP).
-

The above demand diagnostic can be used (for example) to verify that the performance monitoring functions of [Section 6.2.2](#) in another NE are working properly.

6.2.3.2.5 Internal Clock

As discussed in GR-1244-CORE, some NEs may perform clock-related diagnostic routines that have externally observable impacts on their performance. If such diagnostics are utilized by a SONET NE, the criteria in GR-1244-CORE apply.

-
- R6-461 [1228]** If a SONET NE utilizes clock-related diagnostic routines that have externally observable impacts on its (the NE's) performance, it shall conform to the criteria in Section 9.3 of GR-1244-CORE.
-

6.2.3.3 Loopbacks

To support pre-service operations practices and test-related activities in some applications, SONET NEs may need to provide loopbacks for SONET and DS_n signals. Three types of SONET and DS_n signal loopbacks are discussed in this document. These are terminal loopbacks, cross-connect loopbacks, and facility loopbacks. In general, loopbacks interrupt the flow of traffic, change the normal transmission, and often require coordinated activity as two or more NEs are affected. Because of this potential impact on the network, the use of a loopbacks in the SONET network as a routine practice is discouraged. In addition, a DCC must not be used to initiate a SONET loopback if that loopback will interrupt communications between the NE performing the loopback and the user. The use of a DCC in this manner could cause a loopback to be initiated without providing a means for it to be terminated.

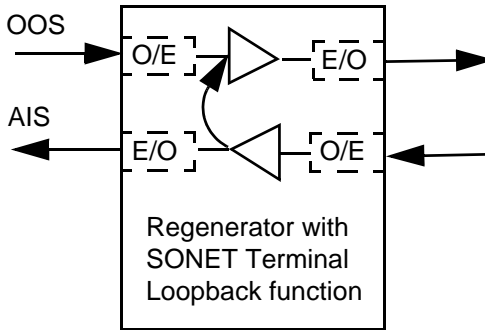
R6-462 [755v2] A SONET NE shall notify the OS when any loopback has been activated and when it has been deactivated.

6.2.3.3.1 SONET Terminal Loopback

In general, loopbacks involve the use of external test equipment to monitor the looped back signal. Conversely, the feature that was described in this section in Issue 1 of this document was strictly internal to the SONET NE. Therefore, the criteria related to that feature were moved to [Section 6.2.3.2.2](#) and there currently are no applicable “SONET terminal loopback” criteria. Although no such criteria are expected to be added to this document in the future, several issues related to terminal loopbacks are discussed below.

In a terminal loopback, the signal that is about to be transmitted is connected to the associated, incoming signal receiver. For example, a SONET regenerator could provide a terminal loopback capability such that an incoming signal is looped back at the “far side” of the regenerator (see [Figure 6-24](#)). Such a capability could allow a user to test a SONET line system step-by-step (i.e., test the facilities between the test equipment and regenerator “A” using the SONET facility loopback described in [Section 6.2.3.3.3](#) at “A”, test those facilities and “A” using the terminal loopback capability at “A”, test the additional facilities between “A” and regenerator “B” using the SONET facility loopback at “B”, etc.). Note that if an NE terminates the SONET Line layer, it would generally not be recommended that it perform terminal loopbacks at its high-speed SONET interfaces, or at tributary SONET interfaces supporting signals that are being used to transport multiple (independent) payloads. Such loopbacks could be used to determine if the NE is multiplexing and demultiplexing a particular payload signal (e.g., a DS_n) correctly; however, they would also disrupt all of the other payload signals carried on the looped back signal. Conversely, it may be appropriate for an NE that terminates the SONET Line layer to provide terminal loopback capabilities for a tributary SONET interface supporting a signal that is being used to transport a single payload.

Figure 6-24 SONET Terminal Loopback Example

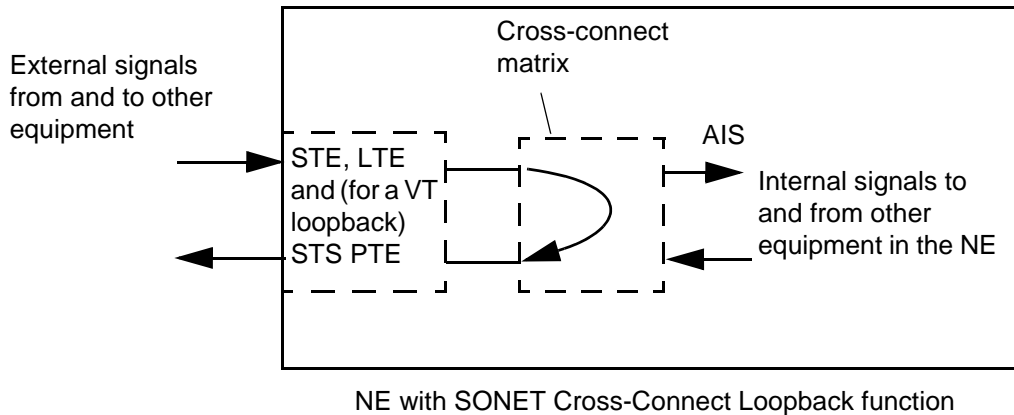


6.2.3.3.2 SONET Cross-Connect Loopback

In general, a cross-connect loopback connects a portion of a received signal to the transmitter in the return direction. In addition, the portion of the signal that is looped back is a unit that can be cross-connected between the various interfaces supported by the NE (e.g., a VT1.5 or STS-1 path signal), and is performed within the NE's cross-connect matrix (see [Figure 6-25](#)).

CR6-463 [1229] A SONET NE may be required to provide a SONET cross-connect loopback, as illustrated in [Figure 6-25](#).

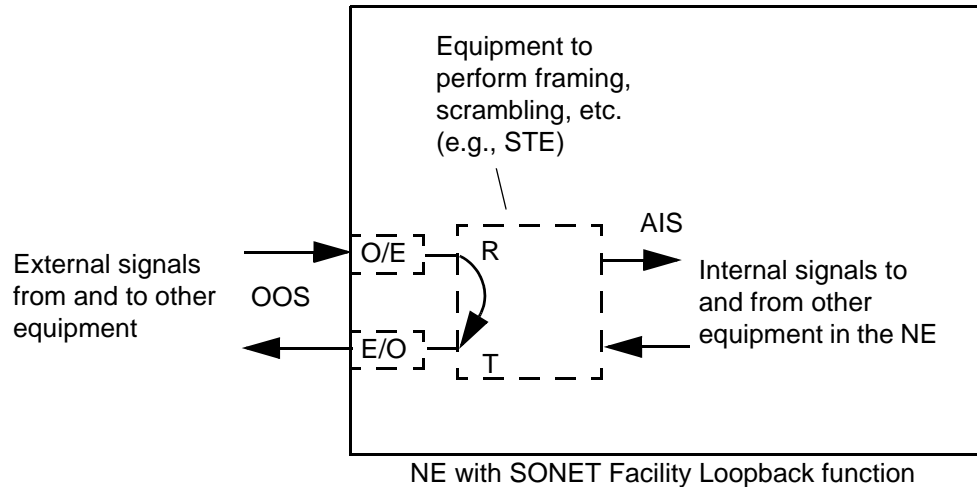
Figure 6-25 SONET Cross-Connect Loopback



6.2.3.3.3 SONET Facility Loopback

In general, a facility loopback connects the incoming received signal to the transmitter in the return direction (see [Figure 6-26](#)).

Figure 6-26 SONET Facility Loopback



O6-464 [749] A SONET NE should provide a SONET facility loopback, as illustrated in [Figure 6-26](#).

In a SONET facility loopback, the signal is looped back by connecting the incoming received signal immediately following the optical-to-electrical conversion (before descrambling) to the associated return transmitter. In the case of an electrical interface (i.e., if there is no optical-to-electrical conversion), the loop occurs before descrambling.

R6-465 [750] A SONET NE shall deny a SONET facility loopback request if either the impacted direction of transmission (i.e., the direction of the looped signal) or the incoming signal to be looped is In_Service or Out_of_Service-Autonomous, as defined in GR-1093-CORE.

R6-466 [751v2] When a SONET facility loopback is activated, the SONET NE shall place the associated out-of-service facilities into a test state, as defined in GR-1093-CORE.

R6-467 [756] The SONET NE shall exit the test state, as defined in GR-1093-CORE, when the loopback is deactivated.

6.2.3.3.4 DS_n Loopback

06-468 [757v2] A SONET NE providing DS_n line terminations should provide DS_n terminal loopback capabilities, as shown in Figure 6-27.

06-469 [1010] A SONET NE providing DS_n line terminations should provide DS_n facility loopback capabilities, as shown in Figure 6-28.

Note that in the DS_n terminal loopback, the DS_n is looped back toward the SONET system just before being transmitted on a DS_n line. This DS_n can be monitored (for example) at the point where the DS_n entered the SONET system as a check on the performance of that system. In the DS_n facility loopback, on the other hand, the DS_n is looped back immediately after entering the SONET system. Therefore, it can be used as a check of the performance of the DS_n facility. Both of these DS_n loopbacks are initiated and removed by commands sent to the NE (e.g., from a craft interface or an OS, either directly or via an NE/NE communications channel such as the Section DCC).

In addition to the types of DS_n loopbacks covered by the preceding criteria, an NE may support DS_n loopbacks that are activated by messages carried in the DS_n signals (e.g., by ESF data link messages in an ESF DS₁, or by FEAC signals in a C-bit parity DS₃). See GR-499-CORE for additional information concerning those types of DS_n loopbacks. In addition, GRs covering specific types of SONET NEs may contain additional criteria on DS_n loopback capabilities.

Figure 6-27 DS_n Terminal Loopback

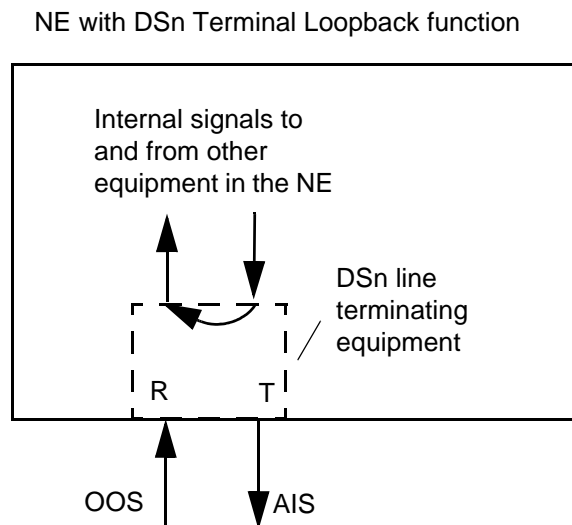
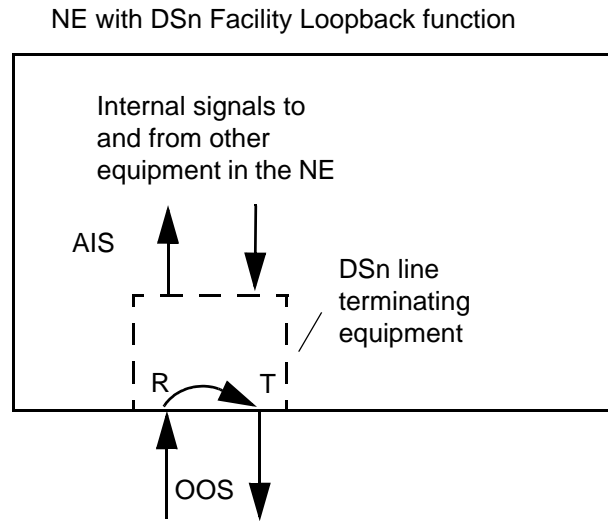


Figure 6-28 DS_n Facility Loopback



6.2.4 Control Features

This section highlights control features for SONET NEs that are necessary to maintain the NE. Other control features explicitly mentioned earlier in [Section 6](#) are not restated.

R6-470 [759] The following control functions shall be provided.

1. Reinitialize system – System reinitialization, or “hard boot,” reloads the operating system of the NE and may affect the state of memory and other resources.
2. Restart system – System restart, or “soft boot,” reloads only an application onto the system and not the operating system. System restart should not affect the state of nonvolatile memory and other resources. Failure states, protection switching configuration, performance parameters and other information necessary for fault isolation should not be affected.
3. Reestablish a failed entity – Reestablishing a failed entity involves temporary techniques to restore a service when an entity has failed. For example, when a facility fails, restoration of a service may involve reconfiguring routing tables or rerouting facilities.
4. Remove an entity from service to run tests – For such tests, traffic should be switched off the entity and subsequent indications should be suppressed.
5. Inhibit and allow alarmed and nonalarmed indications – This capability permits the user to suppress and restart messages from the NE.
6. Check status of equipment configuration – Equipment configuration status shall be retrievable. Equipment configuration status includes the indication of the active hardware entities in replicated equipment and the active synchronization source.

7. Protection switch capabilities – See the required functions identified in [Section 5.3.6](#).
8. Manual switch from active to standby – This capability is used for any replicated hardware or software.
9. Manual switch between synchronization sources – See the required functions identified in [Section 5.4.6](#).

R6-471 [760] A SONET NE shall notify the OS when any control function is executed.

7 Other Generic Criteria

7.1 Physical and Environmental Criteria

This section primarily provides references for physical and environmental criteria for equipment and outside plant cable. In addition, for the reader's convenience some of the information and criteria related to the operating temperature and humidity ranges provided in other documents is restated below.

7.1.1 Operational Environment for Equipment

An operational environment is a set of conditions such as temperature, humidity and airborne contaminant concentration in which the equipment's performance parameters are supposed to be maintained within their specified ranges. In general, the operational environment criteria differ depending on the type of structure in which the equipment is specified to be deployed. For central office equipment and remote terminal equipment intended to be installed only in environmentally controlled structures such as CEVs, GR-63-CORE, *NEBS™ Requirements: Physical Protection*, is the primary document of interest. Along with other information and criteria, that document includes temperature and humidity ranges for both normal operations and short-term periods (e.g., up to 96 hours) where the temperature and humidity conditions are more severe than normal.

R7-1 [1132] All SONET NEs (including those specified for use only in controlled environments such as COs and CEVs) shall operate properly in ambient conditions that are within the normal and short-term (as appropriate) temperature and humidity ranges given in GR-63-CORE.

For equipment that is specified to be able to be installed in structures with limited or no environmental controls, relevant operational environment information and criteria appear in several documents, including GR-63-CORE, GR-487-CORE, *Generic Requirements for Electronic Equipment Cabinets*, and GR-57-CORE, *Functional Criteria for Digital Loop Carrier (DLC) Systems*. Information from all of those documents is combined to arrive at the following requirement for SONET NEs.

R7-2 [1133] SONET NEs specified for use in uncontrolled environments shall operate properly in any ambient conditions that are within a minimum temperature range of -40 to $+65^{\circ}\text{C}$ and relative humidity levels up to at least 90% (with a maximum absolute humidity condition of 0.024 kg of water/kg of dry air).

7.1.2 Electromagnetic Compatibility (EMC)

GR-1089-CORE, *Electromagnetic Compatibility and Electrical Safety - Generic Criteria for Network Telecommunications Equipment*, contains criteria relevant to SONET NEs in the areas of EMC and electrical safety.

7.1.3 Outside Plant Cable

GR-20-CORE contains criteria for physical and environmental requirements for outside plant cable.

7.2 Equipment Design

GR-78-CORE, *Generic Requirements for the Physical Design and Manufacture of Telecommunications Products and Equipment*, contains physical design requirements.

-
- 07-3 [761]** The equipment should be designed to minimize the investment in the frame and bay-work by using the modular design concept to minimize the costs associated with installation and the ongoing operation of the system.
-

7.3 Documentation and Training

The service provider may assume responsibility for engineering, constructing and installing the transmission system.

-
- R7-4 [762v2]** To support the service provider in engineering, constructing and installing the system in a safe and cost-effective manner, the supplier shall provide appropriate documentation and training. In general, the criteria for supplier documentation for NEs in GR-454-CORE, *Generic Requirements for Supplier Provided Documentation*, shall be met.
-

In addition to the documentation criteria referenced above, Section 9.2 of GR-1244-CORE contains several relevant requirements related to the documentation of clock performance characteristics. Those requirements (i.e., **R1244-106** and **R1244-107**) are applicable for SONET products that provide SMC or higher quality internal clocks. Criteria for supplier documentation for outside plant cable are contained in GR-20-CORE.

-
- R7-5 [763]** The supplier shall provide training in accordance with GR-839-CORE, *Generic Requirements for Supplier-Provided Training*.
-

GR-839-CORE provides requirements for:

- Existing training
- Training to be developed by suppliers
- Course content and training documentation
- Course delivery
- Training maintenance and updates
- Product support.

7.4 Safety

This section is concerned with station equipment safety and fiber optic cable safety.

7.4.1 Station Equipment Safety

In general, most of the safety-related criteria applicable to SONET NEs can be found in Sections 12 and 14 of GR-499-CORE, which covers areas such as safety labels (format and locations), and user access to high voltages or temperatures. In addition to those criteria, the requirements in this section are also applicable to SONET NEs.

-
- R7-6 [764]** All safety labels shall be visible to craftspersons when equipment covers are in place and when they are removed.
-

To minimize the exposure of personnel to hazardous voltages, the following requirements are applicable:

-
- R7-7 [765]** Voltages at or above 140 V_{dc} or 50 V_{rms ac} shall be enclosed or guarded to prevent contact. Safety labels shall be conspicuous when the guards are in place and when they are removed.
- R7-8 [766]** The design shall allow craftspersons safe access to parts if metal tools are to be used (e.g., insulating sleeves to guide screwdrivers to recessed potentiometers when nearby parts have hazardous voltages present).
- R7-9 [767]** Arrangements shall be provided to discharge large capacitors (e.g., “bleeder” resistors).
- R7-10 [768]** All external metal parts shall be grounded.
-

For additional information on grounding, refer to GR-63-CORE.

-
- R7-11 [771]** The fiber optic system and required optical test equipment shall be registered and certified with the Department of Health, Education and Welfare Bureau of Radiological Health as specified in 21 CFR 1040.10, *Performance Standard for Light-Emitting Products – Laser Products*. Documentation demonstrating system certification shall be available to assist in the determination of fiber optic safety precautions required to install, operate and maintain the system.
- R7-12 [772]** The equipment involved shall conform to the applicable performance, labeling and informational requirements as specified in 21 CFR 1040.10.
-

7.4.2 Fiber Optic Cable Safety

-
- R7-13 [773]** The fiber optic cable and required optical splicing and test equipment shall be registered and certified with the Department of Health, Education and Welfare Bureau of Radiological Health as specified in 21 CFR 1040.10. Documentation demonstrating system certification shall be available to assist in the determination of fiber optic safety precautions required to install, operate and maintain a fiber optic system.
- R7-14 [774]** The equipment involved shall conform to the applicable performance, labeling and informational requirements as specified in 21 CFR 1040.10.
-

7.5 Quality and Reliability

This section is concerned with the quality and reliability of station equipment and fiber optic cable. Reliability requirements are intended to help ensure that a product will be able to meet all technical specifications consistently throughout the lifetime of the product. In general, these criteria cover such areas as system availability, system qualification, manufacturing tests and inspections, software reliability, and customer support and field reliability.

7.5.1 Network Equipment Reliability

GR-418-CORE is the primary Telcordia document for system reliability requirements. System reliability is often calculated by conducting a system reliability analysis (i.e., a detailed review of the system architecture from a reliability perspective). This review is conducted using system documentation along with technical discussions with system designers and reliability engineers. After this information gathering process, reliability block diagrams, failure mode tables and Markov models are developed to represent the various hardware failure modes in the system. From these reliability models and the circuit pack failure or Mean Time Between Failure (MTBF) rates provided by the equipment supplier, system availability or downtime estimates are calculated for total system downtime and downtime per interface type (e.g., DS3, STS-1 electrical, OC-3, OC-12).

A functional block diagram is often used in system reliability modeling and analysis. This approach is used to represent the physical or electrical relationship between the hardware components (circuits), which are shown as functional blocks. If the failure of any one of these functional blocks interrupts service, it is considered a single point failure. For a function that is protected, service can still be affected when the active and standby circuits both fail, or the active circuit failure does not result in protection switching. Failure group size is then defined as the number of channels (e.g., equivalent DS0s) that become unavailable by these various combinations of failures. Redundancy and segmentation of capacity are often used to reduce failure group size.

In general, the impact of the equipment architecture on the failure group size needs to be considered in the development and design of SONET systems. This is particularly true for systems that support OC-192 or OC-768 interfaces, where an outage could affect the equivalent of as many as 129,024 or 516,096 DS0s.

07-15 [1134v2] Any failure group size that exceeds the equivalent OC-48 capacity should not account for more than 10% of the downtime for an OC-192 or OC-768 system.

7.5.2 Fiber Optic Cable Quality and Reliability

General quality and reliability requirements for fiber optic cable are given in GR-20-CORE. They deal with documentation, manufacturing program analysis, quality surveillance, process and product verification, and initial product qualification and periodic requalification.

7.5.3 Component Reliability Assurance

Component reliability assurance criteria address the necessary attributes and minimum practices of an equipment supplier's component qualification and lot control procedures. Listed below are other Telcordia GRs and TRs that may contain information and criteria applicable to SONET systems:

- GR-326-CORE, which includes two sections addressing the reliability of single-mode optical connectors
- GR-357-CORE, *Generic Requirements for Assuring the Reliability of Components Used in Telecommunication Equipment*, which contains criteria for general components
- GR-468-CORE, *Generic Reliability Assurance Requirements for Optoelectronic Devices Used in Telecommunications Equipment*, which addresses optoelectronic components in central office and loop applications
- TR-NWT-000930, *Generic Requirements for Hybrid Microcircuits Used in Telecommunications Equipment*, which covers hybrid microcircuits
- GR-1221-CORE, *Generic Reliability Assurance Requirements for Passive Optical Components*, which applies to components such as optical splitters and optical couplers
- GR-1312-CORE, which contains requirements for OFAs.

Note that greater use of certain relatively new component technologies is likely in systems supporting OC-192 or OC-768 interfaces. These technologies include high-speed ICs (e.g., GaAs MESFET, Si bipolar, AlGaAs/GaAs HBT), dispersion compensators, external modulators, high-speed lasers (e.g., multi-quantum-well lasers) and/or OFAs. The reliability of many of these devices is currently unknown or supported by a minimal amount of data. In most cases, reliability will likely improve as the technologies continue to mature. However, it is important to determine if there are any inherent reliability problems with these new components, determine their absolute reliability in order to accurately predict system reliability

and service (un)availability, and obtain a better understanding of their failure modes and mechanisms so that appropriate qualification and lot-to-lot controls can be specified in detail.

R7-16 [1135v2] The OC-192 or OC-768 equipment supplier shall identify, for review by the service provider or its representatives, the critical components and technologies used by the product that are not commonly employed in lower speed systems. Device structures, manufacturing processes and controls, qualification, and lot-to-lot controls shall also be provided for review upon request by the service provider.

In some cases, the Telcordia documents listed above either do not currently support some of these new technologies or need to be reexamined for their adequacy at 10 or 40 Gb/s [e.g., high-speed integrated circuits (GR-357-CORE)]. In addition, new test procedures may be necessary for devices such as high-speed lasers (e.g., screening for laser chirp), and specific issues such as sensitivity to Electrostatic Discharge (ESD), temperature and humidity (i.e., hermeticity requirements) need to be addressed.

7.6 Human Factors

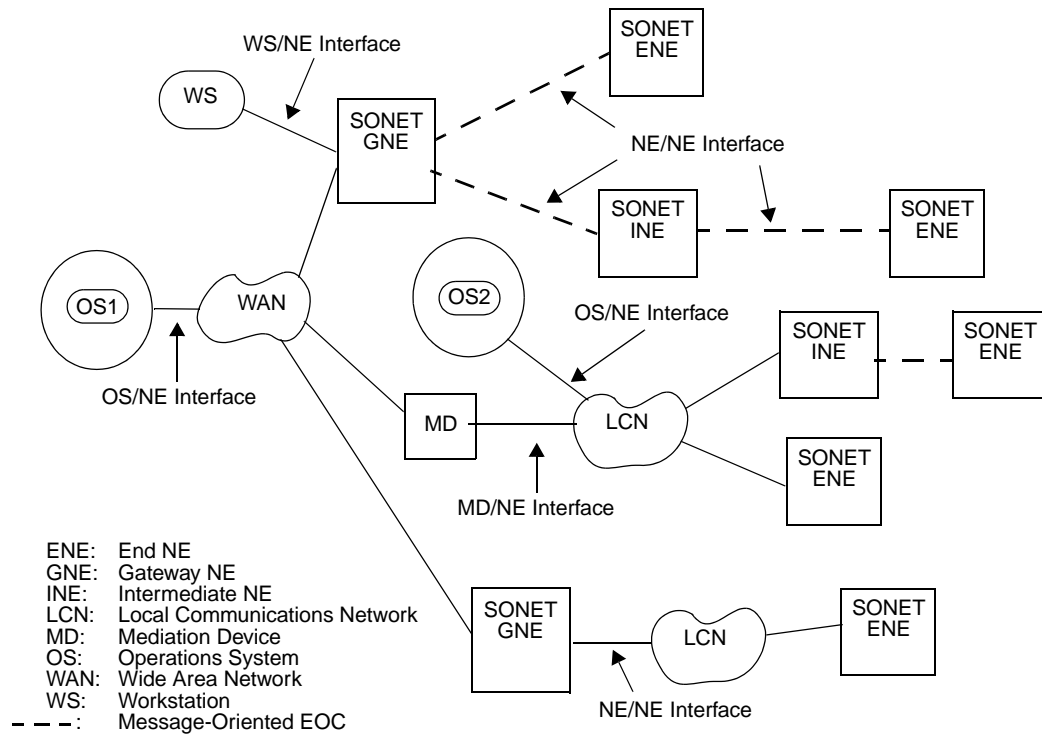
Human factors criteria are contained in Section 12 of GR-499-CORE. In addition, TR-TSY-000824, TR-TSY-000825, *OTGR Section 10.A: User System Interface – User System Language* and GR-826-CORE, *OTGR Section 10.2: User Interface Generic Requirements for Supporting Network Element Operations*, contain criteria concerning the craft interface.

8 SONET Operations Communications

In general, SONET NE operations communications criteria are consistent with the TMN concept in ITU-T M.3010 *Principles for a telecommunications management network*. A TMN is a support network that provides operations communications paths for SONET OS/NE, MD/NE, NE/NE, and WS/NE communications. This section briefly describes the role of SONET NEs in a TMN, and focuses on the implementation of the communications network and mediation functions by using SONET NEs and MDs, and SONET DCCs, WANs and LCNs. In addition, it should be noted that this section is primarily focused on implementations that utilize either full seven-layer protocol stacks (i.e., OSI protocol stacks) or one particular “short stack” that was supported by many OSs and NEs when SONET was initially defined (i.e., the “non-OSI TL1/X.25” protocol stack). Alternate implementations that may be supported either now or in the future [e.g., TL1 directly over Transport Control Protocol/Internet Protocol (TCP/IP)] are currently not considered in detail, and the possibility of revising this section or developing a new GR to fully cover such implementations is FFS.

Operations communications criteria for SONET NEs depend on the role of the NE in the operations communications architecture. More specifically, operations communications criteria depend on whether the NE serves as a Gateway NE (GNE), Intermediate NE (INE), or End NE (ENE), as illustrated in Figure 8-1.

Figure 8-1 SONET Operations Communications: Example NE and Interface Types

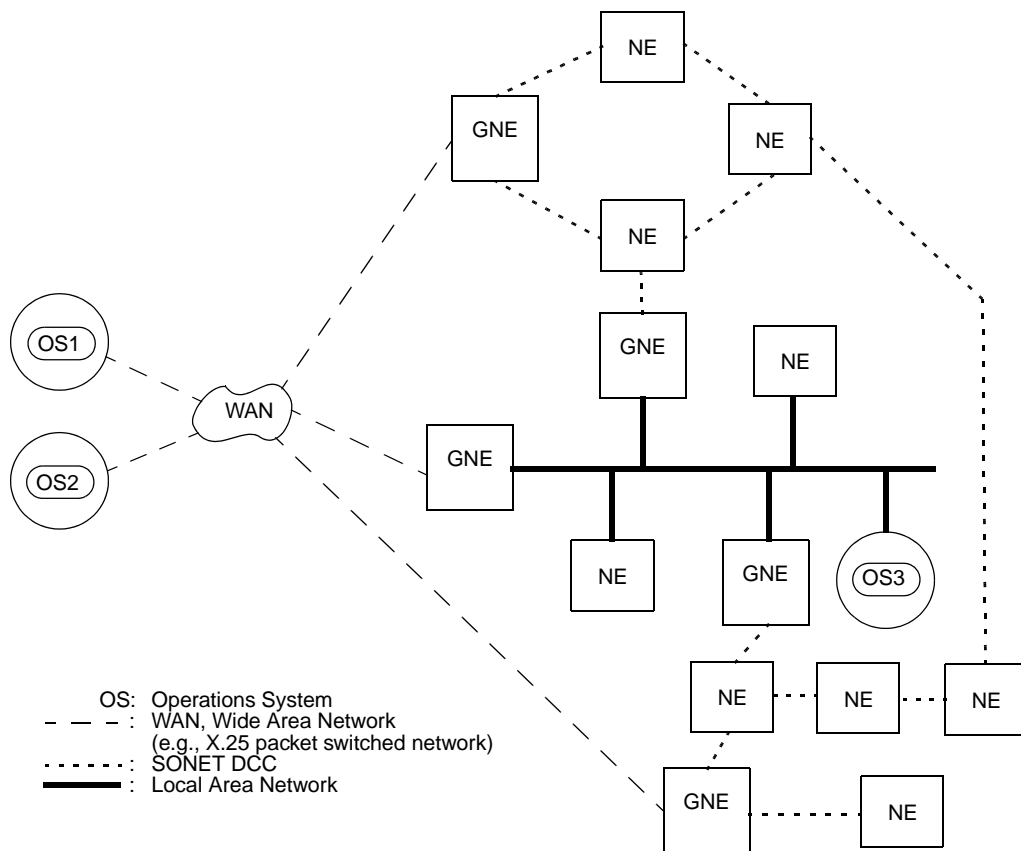


8.1 SONET Operations Communications Architecture

8.1.1 Architecture Overview

SONET operations communications architectures will vary depending on the configuration (e.g., communications within a site or between sites) and application (e.g., OS/NE or NE/NE, tree or ring). This section will look at some operations communications architectures that may be used by network providers, starting with [Figure 8-2](#), which shows a generalized view of a SONET operations communications architecture that includes a WAN, a LAN,¹ and DCC tree and ring connections.

Figure 8-2 Example SONET Operations Communications Architecture



1. Note that in many cases drop-side SONET interfaces are used to connect SONET NEs within a site. In those cases the SONET DCCs are not available, and therefore LANs are typically used to support intra-site operations communications. On the other hand, there may be cases where issues such as security, reliability or cost may make DCCs (rather than LANs) the preferred choice for intra-site operations communications. (For example, this might be the case in a CEV where only a few SONET NEs reside.) In such cases it will be necessary to utilize line-side SONET interfaces for intra-site transport connections.

While some network providers' deployments may be as complicated as (or more complicated than) the example shown in Figure 8-2, others are likely to be more limited. Two examples of this are illustrated in Figures 8-3 and 8-4. In particular, Figure 8-3 shows two OSs (such as a memory administration OS and a surveillance OS) communicating with a SONET network via a WAN. In this example, most of the SONET NEs are in a CO with a LAN, while one far-end SONET NE is connected via a point-to-point DCC configuration. This example also illustrates how mediation devices might be used to provide gateway functions such as interworking a WAN with a LAN. (Section 8.1.3 discusses gateway functions, and Section 8.4 discusses interworking.) Figure 8-4 shows an example of how the SONET DCC may be used in a survivable ring application. In this example, two of the NEs on the ring support gateway functions for added OS/NE operations communications reliability (one gateway is primary and the other provides a backup).

Figure 8-3 Example Intra-site LAN and Point-to-Point DCC

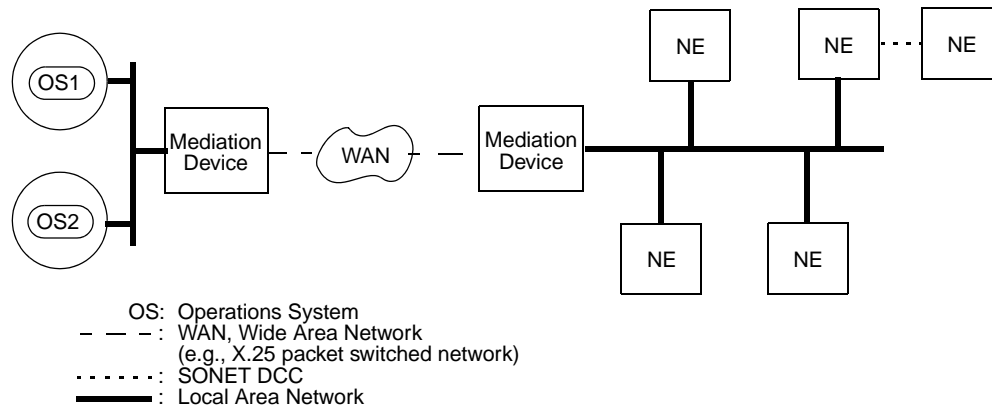
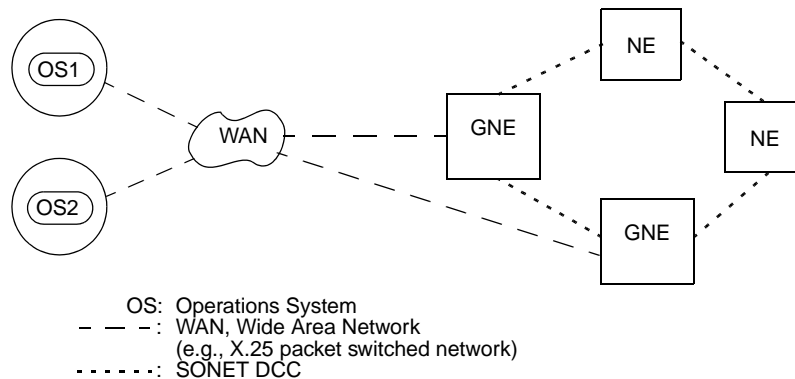
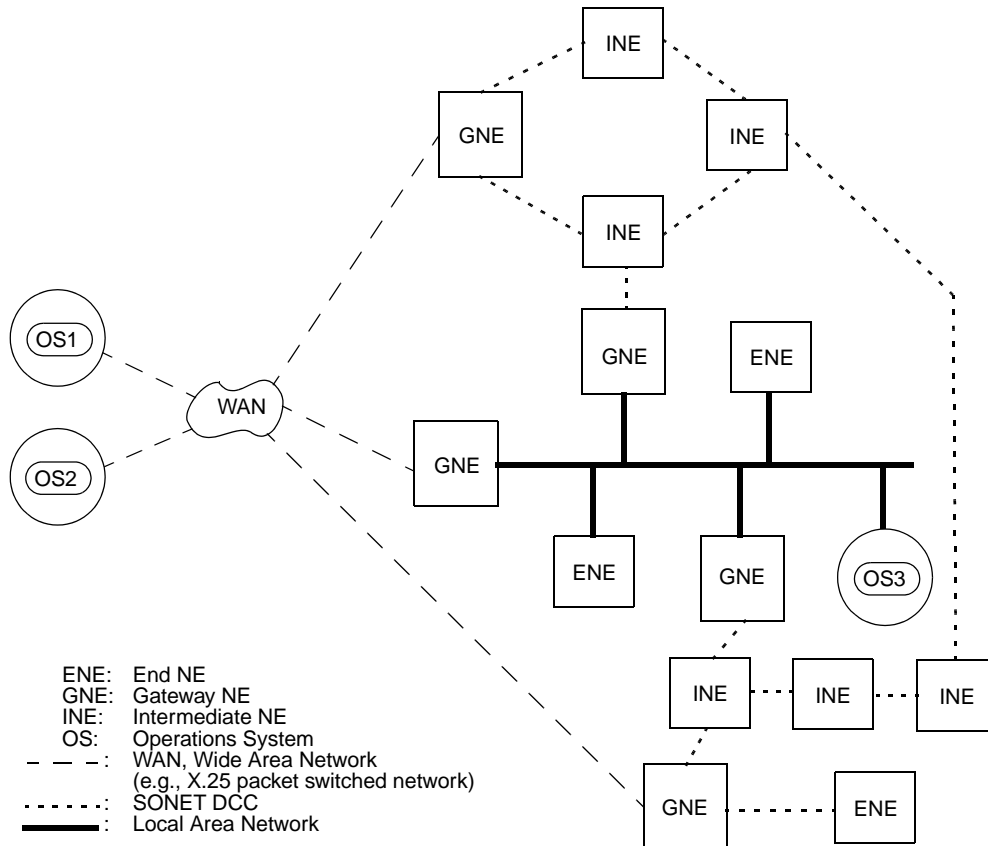


Figure 8-4 Example Operations Communications Network for a Survivable Ring



Depending on an NE's placement and application within a network, it may be a GNE, INE or ENE. Figure 8-5 shows the same operations communications architecture as Figure 8-2, but the NEs have been labeled by the operations communications role they perform. Note that the OSs are not explicitly labeled as End Systems.

Figure 8-5 Operations Communications Functions



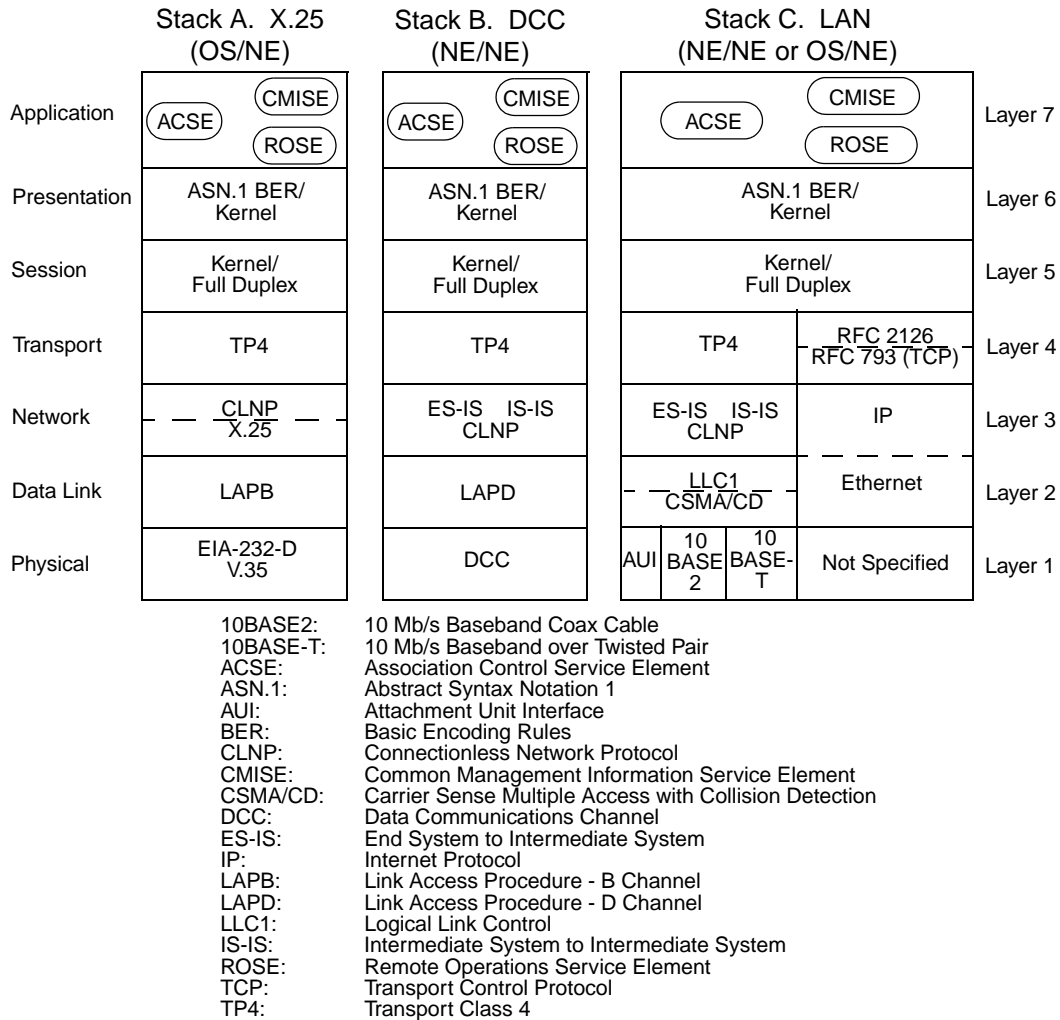
The role that a given SONET NE supports (i.e., GNE, INE or ENE) is determined by the operations communications network architecture. Thus, network providers should work closely with equipment suppliers to ensure that the operations communications functions provided by SONET NEs meet the needs of individual architectures. (This may include possible migration strategies to more complex operations communications network architectures.)

8.1.2 Operations Communications Protocol Overview

Figure 8-6 shows the CMISE versions of the seven-layer protocol stacks specified in this GR for use on WANs, DCCs, and LANs. In all three cases, the Connectionless Network layer Protocol (CLNP ISO 8473) can reside at layer 3, and therefore it is possible for interworking of these protocol stacks to be done by standard routing and relaying functions. (Section 8.5 contains specific requirements for support of routing protocols.) However as shown on the right side of stack C, it is also possible for IP to reside at layer 3 (with TCP above it and Ethernet below it). In applications

where IP is used on some links and CLNP is used on others, interworking will require a number of additional gateway functions, beyond those currently considered in this GR, to be performed. This is an issue FFS.

Figure 8-6 Interactive CMISE/OSI Protocol Stacks for SONET Operations Communications



In cases where full seven-layer/CMISE functionality is not available in OSs and SONET NEs, TL1 can be used for operations messages. In particular, TL1 can be used at the application layer (with ACSE, and instead of CMISE/ROSE) in any of the protocol stacks shown in Figure 8-6, or over a three-layer X.25 protocol stack as defined in TR-TSY-000827, *OTGR Section 11.1: Generic Operations Interface: Non-OSI Communications Architecture*, (i.e., the non-OSI TL1/X.25 stack). Section 8.3 contains detailed requirements for NE support of seven-layer protocol stacks on a DCC or LAN, and the use of CMISE or TL1 messages on those stacks. In applications where the non-OSI TL1/X.25 stack is used, the GNE must perform some application-layer processing of the TL1 messages to interwork that stack and the OSI TL1 protocol stacks used on a DCC or LAN. Details of this application-layer

processing for TL1/X.25 interworking with a TL1/DCC or LAN are provided in [Section 8.4](#). [Section 8.5](#) describes the ISO routing protocols and contains generic criteria for SONET NE support of these protocols.

8.1.3 Gateway NE Requirements

A GNE interworks two different kinds of network technologies. Three possible types of GNEs are considered in this section:²

1. A GNE that interworks an X.25 WAN (protocol stack A in [Figure 8-6](#)) and the SONET DCC (protocol stack B in [Figure 8-6](#))
2. A GNE that interworks the X.25 WAN and an intra-site CLNP-based LAN (the left side of protocol stack C in [Figure 8-6](#))
3. A GNE that interworks an intra-site CLNP-based LAN and the SONET DCC.

In all three of these cases, interworking takes place at layer 3, and is transparent to the management operations, such as alarm surveillance. On the other hand, interworking of any of the three protocol stacks referenced above with an IP-based LAN (the right side of stack C in [Figure 8-6](#)) would require additional functions that are currently not addressed in this GR.³

SONET GNEs that support DCCs and/or CLNP-based LANs must support both level 1 and level 2 IS-IS routing protocol (ISO/IEC 10589) functions as well as the IS role of the ES-IS protocol (ISO 9542).

R8-1 [776v2] All SONET GNEs that support DCCs and/or CLNP-based LANs shall support level 1 and level 2 IS-IS routing functions as defined in ISO/IEC 10589 and Appendix C.

R8-2 [777v2] A SONET GNE shall support the IS Role of the ES-IS routing protocol, as defined in ISO 9542 and Appendix C, over DCC and CLNP-based LAN interfaces.

R8-3 [778v2] A SONET GNE that supports DCCs and/or CLNP-based LANs shall support the IS-IS Reachable Address Prefix functionality defined in ISO/IEC 10589.⁴

2. Also see [Section 8.4](#), which includes criteria for a GNE interworking a non-OSI TL1/X.25 OS/NE interface with various types of NE/NE interfaces.

3. In determining the functions that would need to be supported, it should be noted that RFC 2126 indicates that it “defines a Transport Service which appears to be identical to the Services and Interfaces offered by the ISO Transport Service Definition (ISO 8072), but which will in fact implement the ISO Transport Protocol (ISO 8073) on top of TCP/IP (IPv4 or IPv6), rather than the ISO Network Service (ISO 8348).”

4. Reachable address prefixes are used to facilitate routing of messages from NEs to OSs over dynamically assigned circuits.

A detailed description of level 1 and level 2 routing is provided in [Section 8.5](#). A feature of IS-IS which may be used when both level 1 and level 2 routing is supported is partition repair. If a routing area becomes partitioned because one or more links have failed, level 2 routing (with partition repair) can be used to repair the partition.⁵

-
- O8-4 [779v2]** A SONET GNE that supports DCCs and/or CLNP-based LANs should support the level 1 partition repair function as defined in ISO/IEC 10589 and in [Appendix C](#).
-

In addition to the functions discussed above, some GNEs may provide a message concentration function for the X.25 WAN. Instead of having one X.25 Virtual Circuit (VC) to each SONET NE, the GNE can provide an X.25 VC between it and the OS that can be used for messages to and from the OS and other subtending NEs on the network. The DCC or LAN is used to transport the messages within the SONET operations communications network.

8.1.4 Intermediate NE Requirements

An INE has operations communications connections with one or more subtending NEs and performs forwarding for through traffic. An INE that supports DCCs and/or CLNP-based LANs must support IS-IS level 1 routing and the IS role of the ES-IS protocol.

-
- R8-5 [780v2]** A SONET INE shall support level 1 IS-IS routing functions, as defined in ISO/IEC 10589 and [Appendix C](#), over DCC and CLNP-based LAN interfaces.
- R8-6 [781v2]** A SONET INE shall support the IS role of the ES-IS protocol, as defined in ISO 9542 and [Appendix C](#), over DCC and CLNP-based LAN interfaces.
-

8.1.5 End NE Requirements

ENEs only originate and receive traffic (i.e., they do not support packet forwarding), and they must have direct access to either a DCC, a LAN, or an X.25 WAN (to an OS). An example of such a system may be a SONET NE with only one DCC connection (e.g., a terminal multiplexer at the end of a feeder route) or a SONET NE on a LAN that is not a GNE.

-
- R8-7 [782v2]** A SONET ENE shall support the ES role of the ES-IS protocol, as defined in ISO 9542 and [Appendix C](#), over DCC and CLNP-based LAN interfaces.
-

5. Appropriate links must be in place for partition repair to take place.

8.1.6 Mediation Device

An MD is a network entity that performs mediation functions. Mediation functions include communications gateway functions and possibly TMN-related information processing functions for a specified set of SONET NEs (see [Figure 8-1](#)). The applicable communications gateway functions are essentially the same as those a GNE performs. The decision to use a stand-alone MD rather than a GNE depends on the complexity of the communications gateway functions to be performed. For example, if extensive message translation functions are needed, a stand-alone MD may be warranted. The decision to use an MD (either stand-alone or as an independent module of an NE) to provide gateway communications functions can also be influenced by a supplier's plans to support information processing functions within the MD. Information processing functions are sometimes included in MDs to provide common NE processing functions or subnetwork management functions. Although generic requirements for such capabilities are not discussed in this GR, the possible use of MDs (primarily to support complex communications gateway functions) is included in the discussions below. (See the Glossary for further definition of mediation functions.) Additional information regarding Element Management Layer (EML) applications that may be provided by suppliers in an MD can be found in GR-3000-CORE, *Generic Requirements for SONET Element Management Systems (EMSs)*.

8.2 Communications Types

The primary types of operational entities that a SONET NE must be able to communicate with are OSs, MDs, other NEs and Craftsperson interfaces. The corresponding types of communications are described in the following sections.

8.2.1 OS/NE Communications

Communications with OSs are required at all SONET NEs in order to support network operations and management functions. An OS/NE communications path can be direct or indirect. A direct OS/NE communications path is one with no gateway or intermediate systems between the OS and the NE. The direct path may be provided by a dedicated physical connection, a LAN, or a WAN. An indirect OS/NE path contains at least one NE/NE or MD/NE interface.

SONET OS/NE communications require the use of a message-oriented channel. In initial deployments and in certain network configurations, network providers may want SONET NEs to be directly connected to an OS (i.e., a direct communications path). Such NEs would not necessarily be supporting GNE functions even though they support direct OS/NE communications.

R8-8 [783] All SONET NEs shall provide an OS/NE communications path.

R8-9 [784v3] A SONET NE shall be capable of being equipped with a direct OS/NE interface. For this interface, the protocols at layers 1 through 4 shall conform to either (or both) of the following protocol profiles:

- SONET Operations Communications Interface lower layers profile as specified in Appendix C, *SONET Operations Communications Protocol Profile - Lower Layers*
- The RFC 2126/RFC 793 profile as specified in ITU-T Rec. Q.812, *Upper layer protocol profiles for the Q and X interfaces*, and the IP profiles as specified in ITU-T Rec. Q.811, *Lower Layer protocol profiles for the Q and X interfaces*.

8.2.2 MD/NE Communications

MDs may aid in various operations communications functions as well as perform operations (information processing) functions for a collection of SONET NEs. In this case an MD's primary purpose is to facilitate OS/NE communications and, perhaps, to support remote craftsperson access. To perform these functions, an MD must be able to communicate with all of the supported NEs, OSs, and local craftsperson terminals. MD/NE communications would include a message-oriented MD/NE interface (see Figure 8-1) and, possibly, one or more message-oriented NE/NE interfaces. The MD would also need to support an OS/MD interface and a WS/MD interface. The requirements and protocol specifications for these interfaces are identical to the OS/NE and WS/NE interfaces, respectively.

If an MD is used, it must communicate with each supported NE for OS/NE communications and remote craftsperson communications (because of communications gateway functions of the MD) or for MD/NE communications (because of information processing functions of the MD). An NE may interface directly with an MD via an MD/NE interface, or may do so indirectly via an MD/NE interface and one or more NE/NE interfaces (see Figure 8-1).

If mediation functions are provided within an NE, that NE will assume the MD role. If mediation functions are provided in the form of a stand-alone MD, then a LAN is typically used for the MD/NE interface.

-
- R8-10 [785]** When a stand-alone MD is used in a SONET subnetwork, the language and protocol stack for the MD/NE interface shall be identical to that for the NE/NE interface via a LAN.
-

8.2.3 NE/NE Communications

SONET NEs need to communicate with each other to report alarm, failure, status, and error indications (such as AIS), and to perform protection switching. NE/NE communications are provided in two forms: bit-oriented EOCs and message-oriented operations channels (such as a Section or Line DCC, or a LAN). The use of a bit-oriented EOC versus a message-oriented channel depends on the type of information to be communicated and the time-critical nature of the communications (e.g., AIS needs to be transmitted and received within microseconds and therefore utilizes a bit-oriented EOC).

If two SONET NEs need to communicate with each other using a message-oriented channel, the same channel as utilized in indirect OS/NE communications will be used. NE/NE message-oriented operations communications are currently primarily used to support indirect OS/NE communications. However, it is possible that NE/NE message-oriented communications to support peer-to-peer interactions among SONET NEs will become more prevalent as more complicated SONET architectures are implemented and as operation of the network requires more sharing of information among the SONET NEs in these architectures.

Note that the choice of utilizing a Section or Line DCC, a LAN, or both for intra-site NE/NE communications is up to the network provider based on their operations communications architecture plans.

R8-11 [786] All SONET NEs shall support NE/NE operations communications paths.

R8-12 [787v3] To support indirect OS/NE communications paths and message-oriented NE/NE operations communications, a SONET NE shall be capable of being equipped with Section DCC interfaces to other NEs in accordance with the protocol profiles specified in Appendix C, *SONET Operations Communications Protocol Profile - Lower Layers*.

CR8-13 [1230] To support indirect OS/NE communications paths and message-oriented NE/NE operations communications, a SONET NE may be required to be capable of being equipped with Line and/or (if OC-768 interfaces are supported) extended Line DCC interfaces to other NEs in accordance with the protocol profiles specified in Appendix C.

R8-14 [1231] To support indirect OS/NE communications paths and message-oriented NE/NE operations communications, a SONET NE shall be capable of being equipped with a LAN interface. For this interface, the protocols at layers 1 through 4 shall conform to either (or both) of the following protocol profiles:

- SONET Operations Communications Interface lower layers profile as specified in Appendix C, *SONET Operations Communications Protocol Profile - Lower Layers*
- The RFC 2126/RFC 793 profile as specified in ITU-T Rec. Q.812, *Upper layer protocol profiles for the Q and X interfaces*, and the IP profiles as specified in ITU-T Rec. Q.811, *Lower Layer protocol profiles for the Q and X interfaces*.

8.2.4 Craftsperson/NE Communications

The interfaces defined for craftsperson/NE communications are shown in [Figure 8-12](#). Access to the local NE involves both the craftsperson/WS interface and the WS/NE interface. These interfaces are further defined in [Section 8.6](#).

R8-15 [788v2] All SONET NEs shall support local craftsperson access by means of a WS.

8.3 SONET Operations Communications Interface

All of the communications types described in [Section 8.2](#) may utilize a common set of 7-layer OSI protocols collectively referred to as the SONET Operations Communications Interface. Since the OS/NE WAN subnetwork technology may differ from the NE/NE DCC or LAN subnetwork technology, the physical, data link and network layers are addressed separately for the OS/NE and NE/NE interfaces.⁶ In particular, layers 1 to 3 of the OS/NE X.25 interface rely on GR-828-CORE, *OTGR Section 11.2: Generic Operations Interface – OSI Communications Architecture*, for detailed requirements. However, the SONET Operations Communications Interface transport, session, presentation and application layers are defined to be identical for both SONET NEs and OSs that support SONET NEs.

Most of the criteria for the SONET Operations Communication Interface are based on the criteria in GR-828-CORE. Both the Interactive Protocol Stack and the File-oriented Protocol Stack defined by GR-828-CORE are used for the SONET OS/NE X.25 interface. Additional detail is being supplied here (and in [Appendices C](#) and [D](#)) to further specify the interface and help assure interoperability both between various SONET NEs and between SONET NEs and OSs. For example, this document specifies as mandatory for the SONET Operations Communications Interface some protocol capabilities that are optional or out of scope in GR-828-CORE.

[Appendices C](#) and [D](#) provide a protocol profile for the SONET Operations Communications Interface. [Appendix C](#) addresses the data link layer (LLC for LANs and LAPD for the DCC), the network layer (CLNP, ES-IS and IS-IS routing protocols) and the transport layer. [Appendix D](#) addresses the session and presentation layers and ACSE in the application layer. In the future, these appendices may be expanded to include other protocols.

The use of X.500-based Directory Services (see [08-21 \[1036\]](#)) is based on ANSI T1.245, *Directory Service for Telecommunications Management Network (TMN) and Synchronous Optical Network (SONET)*.

R8-16 [790v3] For Interactive Class communications, SONET NEs shall support CMISE (in accordance with the criteria in [Sections 8.3.7.1](#) and [8.3.7.2](#) and the protocol profile specified in [Appendix D](#)), or TL1 (in accordance with the criteria in [Sections 8.3.7.1](#) and [8.3.7.5](#)) at the application layer.

CR8-17 [793v4] When TL1 is used across the SONET OS/NE interface, either (or both) of the following non-OSI communications interfaces may be required to be supported:

- TL1 over X.25
- TL1 (directly) over TCP/IP.

R8-18 [1137] If TL1 over X.25 is supported, it shall be supported in accordance to the criteria in TR-TSY-000827.

6. For the NE/NE interface, the protocols for the physical and data link layers are also addressed separately for the DCC case and the LAN case.

- R8-19 [1138]** If TL1 over TCP/IP is supported, it shall be supported in accordance to the specifications in NSIF-033-1999, *Requirements for the TCP/IP Protocol Suite on the SONET Access DCN*.
- R8-20 [791v2]** When file-oriented applications are supported, SONET NEs shall support either FTAM or FTP as the application-layer protocol in accordance with the protocol profiles specified in GR-1250-CORE.
- O8-21 [1036]** When a SONET NE supports CMISE on the OS/NE or NE/NE interface, it should also support the X.500-based Directory Services for TMN and SONET, as defined in ANSI T1.245, for the name/address translation service.

The following sections contain additional criteria that apply to seven-layer protocol stacks supported by SONET NEs.

8.3.1 Physical Layer

8.3.1.1 OS/NE

- R8-22 [794v2]** At OS/NE OSI X.25 interfaces, SONET NEs shall support the physical layer requirements of the TP4/Connectionless-mode Network Service (CLNS) Protocol Case as described in GR-828-CORE.

At OS/NE LAN interfaces, the applicable physical layer criteria are those specified in [Section 8.3.1.2](#) (for NE/NE LAN interfaces).

8.3.1.2 NE/NE – LAN

The following criteria apply to cases where CLNP is provided at the network layer of the LAN protocol stack (e.g., the left side of stack C in [Figure 8-6](#)). In cases where IP is provided at the network layer (e.g., the right side of stack C in [Figure 8-6](#)), no particular physical layer specifications are provided.

- R8-23 [795v3]** The physical layer shall support the following 10 Mb/s baseband Media Dependent interface (including electrical interface and connector specifications):
- 10BASE-T per ISO/IEC 8802-3:1996.
- CR8-24 [1011v2]** The physical layer may also be required to support the following 10 Mb/s baseband Media Dependent interfaces:
- 10BASE2, as specified in ISO/IEC 8802-3
 - The media independent Attachment Unit Interface (AUI) as specified in ISO/IEC 8802-3.

8.3.1.3 NE/NE – DCC

As discussed in [Section 3.3.2](#), up to three sets of overhead bytes are defined as DCCs. Two of these are defined at all of the SONET bit rates, and are referred to as the Section DCC (consisting of the D1, D2 and D3 bytes located in the SOH of the first STS-1 in the STS-N signal) and the Line DCC (consisting of the D4 through D12 bytes located in the LOH of the first STS-1). In addition, for OC-768 signals the extended Line DCC is located in the D13 through D156 bytes, which are in a number of the D4, D7 and D10 byte positions of certain STS-1s within the STS-768 (see [Section 3.3.2.3](#)). Together, the D1 through D3 bytes form one 192-kb/s data channel, the D4 through D12 bytes form a 576-kb/s data channel, and the D13 through D156 bytes form a 9216-kb/s channel. Use of the Line and extended Line DCCs is also described in ANSI T1.105.04, *Synchronous Optical Network (SONET): Data Communication Channel Protocols and Architectures*.

R8-25 [796v2] A SONET NE shall be capable of using the Section DCC as the physical layer of the message-oriented EOC.

CR8-26 [1232] A SONET NE may be required to be capable of using the Line DCC (i.e., the D4 through D12 bytes) as the physical layer of a message-oriented EOC.

CR8-27 [1233] A SONET NE that supports OC-768 interfaces may be required to be capable of using the extended Line DCC (i.e., the D13 through D156 bytes) as the physical layer of a message-oriented EOC at those interfaces.

R8-28 [1234] The order of transmission shall be bit 1 of D1 (most significant) through bit 8 of D3 (least significant) for the Section DCC, bit 1 of D4 through bit 8 of D12 for the Line DCC, or bit 1 of D13 through bit 8 of D156 for the extended Line DCC. Data link protocols shall transmit bits across this channel by placing them into the next most significant bits.

The EOC that uses the Section DCC for its physical layer is referred to as the Section EOC, while an EOC that uses the Line DCC or extended Line DCC for its physical layer is referred to as a Line EOC. The following EOC protection criteria apply if linear APS or 4-fiber BLSR protection is provided.

R8-29 [797v2] Section or Line EOCs shall be protected in the same way as working traffic is protected. The protection switch for the EOC shall follow the protection architecture and mode of operation (e.g., if the traffic is protected bidirectionally, the EOC is also protected bidirectionally; if the traffic is protected unidirectionally, then the EOC is also protected unidirectionally).

R8-30 [798] A SONET RGTR that accesses the Section EOC shall read the K1 and K2 bytes in both directions to determine when an EOC is being carried with working traffic (i.e., to determine when an EOC is usable).

This protection scheme results in the loss of a Section or Line EOC if a loss of working traffic occurs. Also, diversely routed regenerators are not supported by this scheme. For protected lines, the protection scheme can generally protect the

Section EOC very quickly. If protection is not available, the network layer routing information distribution protocols (i.e., ES-IS and IS-IS) may still be used to maintain operations communications connectivity (see [Section 8.5](#)). The EOC hardware may also have to be protected against failure by an EOC hardware redundancy feature. GRs and TRs covering specific types of SONET NEs contain any EOC hardware protection criteria.

8.3.2 Data Link Layer

8.3.2.1 OS/NE

R8-31 [799v2] At an OS/NE X.25 interface, SONET NEs shall support the data link layer requirements of the TP4/CLNS Protocol Case as described in GR-828-CORE.

At OS/NE-LAN interfaces, the applicable data link layer criteria are those specified in [Section 8.3.2.2](#) (for NE/NE-LAN interfaces).

8.3.2.2 NE/NE – LAN

The following criteria apply to cases where CLNP is provided at the network layer of the LAN protocol stack (e.g., the left side of stack C in [Figure 8-6](#)). In cases where IP is provided at the network layer (e.g., the right side of stack C in [Figure 8-6](#)), the applicable specifications are those in or referenced by ITU-T Rec. Q.811 (which shows Ethernet as the protocol supporting IP). No additional detailed criteria are provided here.

R8-32 [800v2] Media Access Control functionality for the LAN shall be as specified in ISO/IEC 8802-3 CSMA/CD specifications.

R8-33 [801v2] Logical Link Control functionality for the LAN shall be as specified in ISO/IEC 8802-2 LLC Class 1 Type 1 service and as described in Appendix C of this document.

R8-34 [802v2] The value 0111 1111, in which the leftmost bit is the least significant bit, shall be used for the Destination Service Access Point (DSAP) address and the Source Service Access Point (SSAP) address within each Logical Link Control (LLC) PDU. This value would be represented as 'FE' (hex).

8.3.2.3 NE/NE – DCC

R8-35 [803v2] The data link layer protocol for the SONET DCC shall be based on Link Access Protocol on the D-channel (LAPD) as specified in ITU-T Rec. Q.921, *ISDN user-network interface – Data link layer specification*, and as described in Appendix C.

Note that the data link channel can be in one of two states:

1. Active channel state, where LAPD is sending a frame, an abort sequence, or interframe time fill (contiguous flags between frames), or
2. Idle state, where continuous 1's are sent for at least 15 times.

R8-36 [804] Both the Unacknowledged Information Transfer Service (UITS) and the Acknowledged Information Transfer Service (AITS) shall be supported. AITS shall be the default mode of operation.

The Globally Unique Network Layer Quality of Service (QoS) parameter is used to select between UITS and AITS (see the network layer protocol discussion).

Procedures defined in ITU-T Q.921 for using Service Access Point Identifier (SAPI) and Terminal Endpoint Identifier (TEI) subfields of the LAPD address field for LAPD D-channel applications do not apply to the SONET DCC applications. LAPD TEI management procedures for D-channel applications also do not apply to SONET DCC applications.

R8-37 [805] The SAPI value shall be preassigned, and shall be settable locally or remotely by an OS.

R8-38 [806] The data link layer procedures, with the exception of the TEI management procedure, shall follow the rules ITU-T Q.921 specifies.

R8-39 [807v2] A SAPI value of 62 shall be used for SONET DCC applications.⁷

R8-40 [808] The assignment of user-side/network-side roles (and, hence, the C/R bit value) shall be settable and made before initialization of the data link.

R8-41 [809v2] A TEI value of 0 (zero) shall be used for SONET DCC applications.

8.3.3 Network Layer

The following requirements apply to OS/NE and NE/NE interfaces where CLNP is provided at the network layer (e.g., stack A, stack B and the left side of stack C in [Figure 8-6](#)). In cases where IP is provided at the network layer (e.g., the right side of stack C in [Figure 8-6](#)), the applicable specifications are those in or referenced by ITU-T Rec. Q.811 (which shows Ethernet and TCP as the protocols on either side of IP). No additional detailed criteria are provided here.

Requirements on the N-SEL portion of the NSAP are provided to allow NEs to differentiate between TP4 PDUs and TL1 TID Address Resolution Protocol (TARP) PDUs (see [Section 8.7](#) for a description of TARP).

7. The need to reserve additional SAPI values for specific purposes (e.g., DCC maintenance) is FFS.

-
- R8-42 [810]** When TP4 is being run over CLNP, the N-SEL portion of the NSAP shall be set to an initial value of '1D' (hex).
- R8-43 [811]** When TARP is being run over CLNP, the N-SEL portion of the NSAP shall be set to an initial value of 'AF' (hex).
- R8-44 [812]** The capability to change the value of the N-SEL when the OSI stack is reinitialized shall be supported.
-

8.3.3.1 OS/NE

-
- R8-45 [813v4]** At an OS/NE OSI X.25 interface, SONET NEs shall support the network layer requirements of the CL-WAN profile (CLNS2) as described in ITU-T Q.811, except that use of the ES-IS protocol shall not be supported.
-

ES-IS will not be used over the X.25 WAN by either the OSs or the NEs. This differs from the requirements found in ITU-T Q.811.

-
- R8-46 [1099v2]** At an OS/NE OSI X.25 interface, SONET NEs shall also support the X.25 Subnetwork Service and Protocol requirements found in Section 5.3.2.4 of GR-828-CORE (except as noted below).
-

The exception mentioned in **R8-46 [1099v2]** is that **R828-159** is to be replaced by the following:

To operate the Connectionless Network Layer Protocol (CLNP) over X.25 subnetworks, the Subnetwork Dependent Convergence Function defined in ISO/IEC 8473-3 ITU-T X.622 shall be used.

Telcordia intends to make this change in GR-828-CORE when it is next reissued.

At OS/NE-LAN interfaces, the applicable network layer criteria are those specified in [Section 8.3.3.2](#) (for NE/NE-LAN and DCC interfaces).

8.3.3.2 NE/NE – LAN and DCC

-
- R8-47 [814v2]** The network layer protocol for DCCs and LANs shall be CLNP (ISO/IEC 8473-1) as specified in [Appendix C](#), using the appropriate subnetwork dependent convergence function(s) specified in ISO/IEC 8473-2 (for LAN) and ISO/IEC 8473-4 (for DCC).

- R8-48 [816]** The ISO 8473 Category 3 QoS function shall be supported.
-

The QoS parameter is used to select either UITS or AITS service in the LAPD data link layer protocol.

- R8-49 [817]** The coding of the QoS parameter for the selection of UITS/AITS in the data link shall be as follows:
- A. The absence of a QoS parameter shall select AITS
 - B. Bits 7 and 8 set to '1' (Globally Unique QoS) and bit 1 set to '1' shall select AITS
 - C. Bits 7 and 8 set to '1' (Globally Unique QoS) and bit 1 set to '0' shall select UITS.
- CR8-50 [818]** Other ISO 8473 Category 3 functions may be required except where prohibited below.
- R8-51 [819v2]** The following service/protocol parameters shall have the values specified below:
- A. Error Reporting (E/R) Flag – As stated in ANSI T1.204, *OAM&P – Lower Layer Protocols for Telecommunication Management Network (TMN) Interfaces, Q3 and X Interfaces*, the setting of E/R flag is a local matter. The default value of this flag shall be zero to avoid excessive network traffic that can result during broadcast routing.
 - B. Partial Source Routing – Partial source routing shall not be supported because NBSIR 88-3824-1, containing OSI implementation agreements, has identified a defect with the partial source routing option that can cause NPDUs to loop in the network until their lifetime expires.
 - C. Inactive Subset – Implementations shall not transmit NPDUs encoded using the ISO 8473:1988 inactive subset. Received NPDUs encoded with the inactive subset shall be discarded.
 - D. Segmentation – The non-segmenting subset shall not be used. However, implementations shall be capable of receiving and correctly processing NPDUs that do not contain the segmentation part.
 - E. Segmentation Permitted (SP) Flag – Implementations shall not generate data NPDUs without a segmentation part (i.e., the SP flag shall be set to '1' and the segmentation part shall be included).
 - F. Lifetime Control – The lifetime parameter shall be used as Paragraph 6.4 of ISO 8473:1988 specifies. This parameter shall have an initial default value of at least three times the network span (number of network entities) or three times the maximum transit delay (divided by 500 ms), whichever is greater, as ISO 8473 specifies. The initial default PDU Lifetime Control shall be 10 seconds.
- O8-52 [820]** The CLNS Congestion Notification should be used. The default value of '0' should be used when originating NPDUs.
- O8-53 [821]** The mandatory and optional approaches to congestion avoidance and recovery given in NIST Publication 500-202, Part 4, Section 5.1.2.5 should be used.
- R8-54 [822v2]** The destination and source addresses used for SONET applications shall be NSAP, as specified in ISO/IEC 8348. The Domain Specific Part (DSP) shall be the ISO Data Country Code (ISO DCC) format as described in ANSI T1.204.

The NSAP is used in the network layer to address SONET NEs and their supporting OSs in this OSI environment. The NSAP is divided into two components: the Initial Domain Part (IDP) and the DSP. The IDP is further subdivided into the Authority and Format Identifier (AFI) and the Initial Domain Identifier (IDI). The AFI identifies the IDI format and the DSP syntax. For SONET, the value of the AFI is 39 (decimal), which identifies the ISO DCC as the address format and preferred binary encoding for the DSP. The ISO DCC is a three-digit numeric code allocated according to ISO 3166. The IDI portion of the IDP has the value of 840 (decimal) for the United States. ISO/IEC 8348 defines a fixed total length for the IDP of 5 digits, or 2 1/2 octets. Under the preferred binary encoding rules, each digit of the IDP is encoded in Binary Coded Decimal (BCD), where each decimal digit is encoded and transmitted using one semi-octet. Due to the odd number of digits allocated to the IDP, it is necessary to pad the IDP with a semi-octet to ensure an integral number of octets as defined in ISO/IEC 8348. The AFI and ISO DCC IDI define the DSP to be composed of 17 binary octets. The breakdown of the entire NSAP, and the 17 octets of the DSP, is shown in [Figure 8-7](#).

Figure 8-7 SONET NSAP Format

Field Name	IDP (Including Pad)			DSP						
	AFI	IDI	IDI PAD	DFI	ORG	RES	RD	AREA	SYSTEM	SEL
Number of Octets	1	1 1/2	1/2	1	3	2	2	2	6	1

IDP: Initial Domain Part
 DSP: Domain Specific Part
 AFI: Authority and Format Identifier
 IDI: Initial Domain Identifier
 DFI: DSP Format Identifier
 ORG: Organization Identifier
 RES: Reserved
 RD: Routing Domain
 AREA: Identifier for a Routing Area within a Routing Domain
 SYSTEM: Routing Entity Identifier for Routing Entity within an NE or OS
 SEL: NSAP Selector

The encoding of the IDP, including the IDI Pad semi-octet, is shown in [Figure 8-8](#).

Figure 8-8 IDP Encoding

Field Name	IDP (Including Pad)		
	AFI	IDI	IDI PAD
Decimal Value	39	840	none
BCD encoding in NSAP	0011 1001	1000 0100 0000	1111

The DSP values are allocated by the ISO member body or sponsored organization to which the ISO DCC value has been assigned. For the United States, ANSI has been chosen as the organization responsible for the DSP format. ANSI assigns values to the network providers for the ORG fields. The DSP Format Identifier (DFI) portion of the DSP is 128 (decimal) to identify the SONET DSP format. This is encoded in binary as 1000 0000. The remaining fields of the DSP (also encoded in binary) are assigned by the network provider and the equipment supplier. The DSP fields are used by the IS-IS routing protocol, to provide information about the hierarchical structure of routing areas and domains. The NSAP Selector serves to differentiate multiple entities (for example, TP4 or TARP entities) operating over the same network entity. [Section 8.5](#) contains additional information about the routing protocols.

-
- R8-55 [823]** The System ID field shall be assigned a 6-octet IEEE address by the equipment supplier.
-

The other assignable fields are provisioned by the network provider.

8.3.4 Transport Layer – OS/NE and NE/NE

-
- R8-56 [824v2]** Class 4 of ISO/IEC 8073 shall be supported as specified in [Appendix C](#).
-

Transport Protocol Class 4 (TP4) over Connection-mode Network Service (CONS) and Transport Protocol Classes 0, 1, 2 and 3 are not used for SONET applications.

-
- R8-57 [825]** TP4 implementations shall be capable of receiving and processing all possible parameters for all possible Transport PDUs (TPDUs), dependent upon the class and optional functions implemented.
- R8-58 [826]** If the ISO session layer is being run over TP4, then the Transport Service Access Point (TSAP)-ID shall be set to an initial value of “TT” (ASCII), i.e., ‘5454’ (hex).
- R8-59 [827]** The capability to change the value of the TSAP-ID (within a range of 0 to 4 octets) when the OSI stack is reinitialized shall be supported.
-

The following transport layer requirements are included in International Standardized Profile (ISP) 10608-1.

-
- R8-60 [828]** An unknown parameter in any received Connection Request (CR) TPDU shall be ignored.
- R8-61 [829]** Known parameters with invalid lengths in a CR or Connection Confirm (CC) TPDU shall be handled as a protocol error.
- R8-62 [830]** Known parameters with valid lengths but invalid values in a CR TPDU shall be handled as follows:
-

- A. TSAP-ID: Send a Disconnect Request (DR) TPDU
- B. TPDU size: Ignore parameter, use default
- C. Version: Ignore parameter, use default
- D. Checksum: Discard CR TPDU
- E. Alternate protocol classes: Protocol error

R8-63 [831] Unrecognized or inapplicable bits of the additional options parameter shall be ignored.

8.3.5 Session Layer – OS/NE and NE/NE

R8-64 [832] The ISO session layer shall be supported as specified in Appendix D.

R8-65 [833] If the ISO presentation layer is being run over the ISO session layer, then the session selector parameter shall be set to an initial value of “SS” (ASCII), i.e., ‘5353’ (hex).

R8-66 [834] The capability to change the value of the session selector parameter when the OSI stack is reinitialized shall be supported.

O8-67 [835] A session service-user-data size of up to 65,535 octets should be supported.

The above objective may become a requirement in the future.

8.3.6 Presentation Layer – OS/NE and NE/NE

R8-68 [836] The ISO presentation layer shall be supported as specified in Appendix D.

R8-69 [837] The P-SEL shall be no greater than four octets in length.

The following presentation selector values are used to differentiate between various SONET Application Service Elements (ASEs). These values apply to the called presentation selector that must be included in connect-presentation Presentation PDUs (PPDUs, see Appendix D).

R8-70 [838] When CMISE PDUs are sent, the P-SEL shall be set to an initial value of ‘01’ (hex).

R8-71 [1037] When X.500 Directory Access Protocol (DAP) PDUs are sent from the Directory User Agent (DUA) to the Directory System Agent (DSA), the P-SEL shall be set to an initial value of ‘04’ (hex).

- R8-72 [839]** When FTAM PDUs are sent, the P-SEL shall be set to an initial value of '02' (hex).
- R8-73 [841]** When TL1 PDUs are sent, the P-SEL shall be set to an initial value of 'AF' (hex).
- R8-74 [842]** The capability to change the value of the P-SEL when the OSI stack is reinitialized shall be supported.
-

8.3.7 Application Layer – OS/NE and NE/NE

8.3.7.1 ACSE

- R8-75 [843]** The Association Control Service Element (ACSE) shall be supported as specified in Appendix D.
-

The use of the ACSE Authentication Functional Unit is under study. In the future, its use may be required for SONET applications.

8.3.7.2 ROSE/CMISE

The following requirements apply to SONET NEs that utilize CMISE at the application layer.

- R8-76 [844v2]** SONET NEs that utilize CMISE at the application layer shall support Remote Operations Service Element (ROSE)/CMISE as specified by GR-828-CORE.
- R8-77 [845v2]** The TMN Application Context defined in ITU-T M.3100, *Generic network information model*, Section 11, shall be used during association establishment.
- R8-78 [846]** The CMISE objects and service mappings for SONET that are contained in GR-1042-CORE and GR-1042-IMD, and the objects and service mappings supporting surveillance and memory administration that are contained in GR-836-CORE and GR-836-IMD shall be supported as per the criteria in those documents.
-

8.3.7.3 FTAM or FTP

When file-oriented applications are supported, SONET NEs must support FTAM or FTP as the application-layer protocol (see **R8-20 [791v2]**).

8.3.7.4 Name/Address Translation Services

In order to establish an association between communicating systems over an OSI network, an address, or NSAP, is required. Typically, an OS would know the name of the system it wishes to establish an association with; however, the address of that system is needed. There are a number of ways a translation between the name of a system and the corresponding address can be achieved. A local static mapping table can be used at the OS; in the case of systems named by TIDs (i.e., when TL1 is used) TARP can be used; or an X.500-based Directory Service as defined in ANSI T1.245 can be used. A local static mapping table is a local matter at the OS or NE. TARP is defined in [Section 8.7](#). An objective for the use of X.500 Directory Services is contained in [Section 8.3](#).

8.3.7.5 TL1

The criteria in this section apply when TL1 is used at the application layer.

R8-79 [847] The presentation context shall contain the TL1 abstract syntax and TL1 transfer syntax that these identifiers specify:

```
tl1AbstractSyntax OBJECT IDENTIFIER ::= { 1 3 17 104 11 2 bellcoreSONETSyntax(1) }
```

```
tl1TransferSyntax OBJECT IDENTIFIER ::= { 1 3 17 104 12 2 bellcoreSONETSyntax(1) }
```

R8-80 [848] The transfer syntax for TL1 messages shall be the ASCII encoding of the character string constituting each TL1 message.

R8-81 [849] The defined context set shall contain the following:

```
ABSTRACT SYNTAX { 1 3 17 104 11 2 bellcoreSONETSyntax(1) }  
  { joint-iso-ccitt association-control (2), abstract-syntax (1), apdus (0), version (1) }
```

```
TRANSFER SYNTAX { 1 3 17 104 12 2 bellcoreSONETSyntax(1) }  
  { joint-iso-ccitt asn1 (1), basic-encoding (1) }
```

R8-82 [850v3] PPDUs containing TL1 messages exchanged between communicating SONET NEs shall use the choice of “fully encoded data” for user data defined in ISO/IEC 8823-1. Also, presentation data values shall use the “octet-aligned” choice as shown below.

```
User-data ::= CHOICE {  
  [APPLICATION 0] IMPLICIT Simply-encoded-data,  
  [APPLICATION 1] IMPLICIT Fully-encoded-data }
```

```
Fully-encoded-data  
  ::= SEQUENCE OF PDV-list
```

```
PDV-list ::= SEQUENCE {  
  Transfer-syntax-name OPTIONAL,  
  Presentation-context-identifier,  
  presentation-data-values  
  CHOICE {  
    single-ASN1-type [0] ANY,
```

octet-aligned	[1] IMPLICIT OCTET STRING,
arbitrary	[2] IMPLICIT BIT STRING }

Note that when ACSE is used to establish an OSI association for TL1, the “single-ASN1-type” choice (above) is used for ACSE PDU parameters that are mapped to user data parameters in PPDUs.

- R8-83 [851]** The TL1-ASE shall consist of the appropriate subset of the TL1 messages defined in GR-833-CORE, GR-199-CORE and GR-834-CORE.
- R8-84 [852]** All application context definitions for associations over which TL1 messages will be exchanged shall include ACSE and the TL1-ASE.
- R8-85 [853]** TL1 messages shall be exchanged using data presentation PDUs.
- R8-86 [1038]** The upper limit on TL1 message size at the application layer, specified in TR-TSY-000827 as 4096 bytes for TL1 over X.25, shall be 4096 bytes for TL1 over any protocol stack or transport mechanism.
- R8-87 [854]** Peer NE/NE communications shall use an association established with the Application Context Identifier (ACI) below:

```
tl1PeerComm OBJECT IDENTIFIER ::= { 1 3 17 104 10 3 tl1PeerComm(1) }
```

See [Section 8.4.1](#) for a discussion of application contexts for indirect OS/NE communications via the TL1-based NE/NE interface and a GNE.

8.4 Interworking between OSs and SONET NEs

This section provides requirements for SONET NEs and GNEs to interwork a non-OSI X.25 OS/NE interface with DCC or CLNP-based LAN NE/NE interfaces for interactive class OS/NE communications. It also discusses interworking the DCC NE/NE interface with the CLNP-based LAN NE/NE interface for interactive communications. This latter discussion also applies to interworking the DCC-based NE/NE interface with the CLNP-based LAN OS/NE interface.

There are several interworking cases (based on protocols and messages supported) that can be individually examined:

- **TL1/X.25 [OS] – TL1/OSI [SONET]**

- *SONET LAN Interworking:* TL1 messages are sent between an OS and a SONET NE, using an intra-site LAN (see [Figure 8-9](#)).
- *SONET DCC Interworking:* TL1 messages are sent between an OS and a SONET NE, using a DCC network.
- *SONET LAN and DCC Interworking:* TL1 messages are sent between an OS and a SONET NE, using a DCC network and an intra-site LAN (see [Figure 8-10](#)).

- **TL1/X.25 [OS] – CMISE [SONET]**

- *SONET LAN and/or DCC Interworking:* TL1 messages are used by OSs, and CMISE messages are used by NEs.

- **CMISE [OS] – CMISE [SONET]**

- *SONET LAN and/or DCC Interworking:* CMISE messages are used by both OSs and SONET NEs.

- **CMISE or TL1/OSI [SONET] – CMISE or TL1/OSI [SONET]**

- *SONET DCC and LAN Interworking:* CMISE or TL1 messages are sent between a SONET NE using a DCC network and a SONET NE using an intra-site LAN. This case also applies to messages sent between a SONET NE using a DCC network and an OS using a LAN.

8.4.1 TL1/X.25 [OS] – TL1/OSI [SONET]

In this interworking scenario, the OS/NE interface is TL1 messages carried over X.25, and the NE/NE interface is TL1 messages carried over a SONET DCC or CLNP-based LAN protocol stack. The ASEs used on the NE/NE interface in this case are ACSE and the TL1 ASE (as described in [Section 8.3.7](#)).

Three OS/NE interworking issues need to be addressed:

1. The GNE must determine the NSAP of the destination NE for OS-to-NE messages
2. The GNE must direct autonomous messages from remote NEs to the appropriate OS(s)
3. Responsibility for establishing connections between OSs, GNEs and target NEs must be determined.

These issues are addressed below.

8.4.1.1 Determine Destination NSAP

As TL1 messages destined for a specific NE are sent from the OS to a GNE, the GNE must determine the destination NE's NSAP address before it routes the message toward the destination. If the GNE supports the multiplexing of TL1 messages for multiple Remote NEs (RNEs) on a single X.25 VC between the OS and the GNE (called "multiple RNEs/VC"), the GNE has to do a TID-to-NSAP mapping for each TL1 message. If the GNE supports one VC between the OS and GNE for each destination NE (called "single RNE/VC"), then the GNE has to do an LCN-to-NSAP mapping. With the multiple RNEs/VC method, fewer X.25 VCs are used; with the single RNE/VC method, the need to extract TIDs from each TL1 message is eliminated.

GNEs that support multiple RNEs/VC or a single RNE/VC have different criteria that they need to meet. The following requirement is common to both types of GNEs.

-
- R8-88 [855]** The GNE shall maintain static information to map subtending NEs' TIDs to NSAP addresses.
-

This static information (e.g., a table) can be populated and maintained by TARP (see [Section 8.7](#)).

- R8-89 [856]** If a GNE supports multiple RNEs/VC, then for an OS-to-NE message, the GNE shall:

- A. Receive the TL1 message from an X.25 VC
- B. Extract the TID information from the message and map the TID to the destination NE's NSAP
- C. Determine the appropriate association to be used or established
- D. Establish the association (if necessary)
- E. Forward the TL1 message over the appropriate association to the destination NE.

- R8-90 [857]** If a GNE supports single RNE/VC, then the following apply.

- A. The GNE shall maintain dynamic information to map X.25 LCNs to NSAP addresses of subtending NEs.
 - B. When an X.25 VC is established between an OS and a GNE for communications with a remote NE, the GNE shall:
 - Listen on that VC for a TL1 command
 - Extract the TID from the command
 - Map the TID and the X.25 VC LCN to the NSAP
 - Establish the association with the remote NE.
 - C. For an OS-to-NE message, the GNE shall:
 - Receive the TL1 message from an X.25 VC
 - Map the LCN from the X.25 VC over which the TL1 message is received to the destination NE's NSAP
 - Determine the appropriate association to be used
 - Forward the TL1 message over the appropriate association to the destination NE.
-

8.4.1.2 Directing Autonomous Messages

There are two kinds of TL1 messages that a remote NE can send to an OS. These are responses to commands, and autonomous messages. Responses to TL1 commands are sent by the remote NE back over the same association over which the request was received. The GNE forwards the response to the correct OS based on the association over which it was received.

-
- R8-91 [858]** An NE shall send responses to TL1 commands using the same application association to the GNE on which the request from the OS was received.
-

Autonomous messages present a problem in this case because the OSs do not have NSAPs and autonomous messages do not have TIDs. Thus there is no way that the PDU itself can contain any information that helps the GNE forward it to the correct OS.

Using the modified multiple ACI method,⁸ one association is established between the GNE and the RNE for each OS with which the RNE communicates. Each association between the RNE and the GNE is established using an ACI that identifies a specific OS.

Three ACIs are defined for this purpose: one for a maintenance OS, one for a memory administration OS, and one for a testing OS. GR-199-CORE contains a TL1 message that can be used to establish and maintain the table that maps X.121 (*International numbering plan for public data networks*) addresses to ACIs. If an X.25 VC is established between an OS and a GNE that uses an unknown X.121 address, then the GNE should establish the association with the remote NE using the t11PeerComm ACI. The GNE maintains information that maps each association to a particular OS, so that when a message (either request-response or autonomous message) is sent from a remote NE to the GNE, the GNE knows which OS should get the message.

-
- R8-92 [859]** When an RNE needs to send an autonomous TL1 message to a particular OS, it shall send the message to the GNE over the association established with the appropriate ACI (t11Maintenance, t11MemoryAdministration or t11Test).

- R8-93 [860]** To support RNEs that need to send autonomous messages to a particular OS, the GNE shall support the following ACIs on the NE/NE interface in order to direct autonomous messages to the correct OS.

t11Maintenance OBJECT IDENTIFIER ::= {1 3 17 104 10 3 t11Maintenance(2)}

t11MemoryAdministration OBJECT IDENTIFIER ::= {1 3 17 104 10 3
t11MemoryAdministration(3)}

t11Test OBJECT IDENTIFIER ::= {1 3 17 104 10 3 t11Test(4)}

- R8-94 [861]** The GNE shall be capable of establishing associations with subtending NEs for communication between a subtending NE and an OS using the following ACI:

t11PeerComm OBJECT IDENTIFIER ::= {1 3 17 104 10 3 t11PeerComm(1)}

Note that this ACI is the same one defined in [Section 8.3.7.5](#) for peer TL1 communication on the NE/NE interface. Use of associations established with the t11PeerComm ACI are a local matter.

8. This method is modified from the original multiple ACI method described in TR-NWT-000253, Issue 2.

-
- R8-95 [862]** The GNE shall maintain static information to map the X.25 address of each OS to its related ACI. This static information shall be provisionable.
- R8-96 [863]** When an association is established between a GNE and a subtending NE for OS/NE communications, the GNE shall dynamically relate the association with the appropriate OS's X.25 virtual circuit.
- R8-97 [864]** For NE-to-OS messages, the GNE shall route the TL1 message over an X.25 VC to the appropriate OS based on the association on which it was received.
- R8-98 [865]** When an OS needs to communicate with an NE through a GNE, the GNE shall use or establish an association with the "target" NE using the appropriate ACI (tl1Maintenance, tl1MemoryAdministration or tl1Test). If the OS's X.121 address is not present as part of the GNE's mapping information, then the GNE shall use the tl1PeerComm ACI.
-

How an NE determines which association to use is a local matter. One example of how this might be done is by having it built into the application in the NE. Another example is through the use of a TL1 message telling the NE what autonomous messages should be sent over which associations.

8.4.1.3 Establishing Connections

Establishing end-to-end OS/NE connectivity in this TL1 interworking environment is a two-step process: the X.25 VC between the OS and the GNE must be established, and the OSI association between the GNE and the target NE must be established. Session establishment is one-way, from the OS through the GNE to the target NE. The communication between the OS and the target NE is two-way.

-
- R8-99 [866]** The OS shall initiate the X.25 VC to the GNE over the TL1/X.25 interface.
- R8-100 [867]** The GNE shall establish OSI application associations with remote NEs.
- R8-101 [868]** If the X.25 VC between the OS and the GNE goes down, either deliberately or through a communications failure, then the GNE shall take down the association(s) with the RNE(s) using that X.25 VC.
-

If multiple RNEs/VC are supported, then associations to several remote NEs will be taken down. If a single RNE/VC is supported, then an association to one remote NE is taken down.

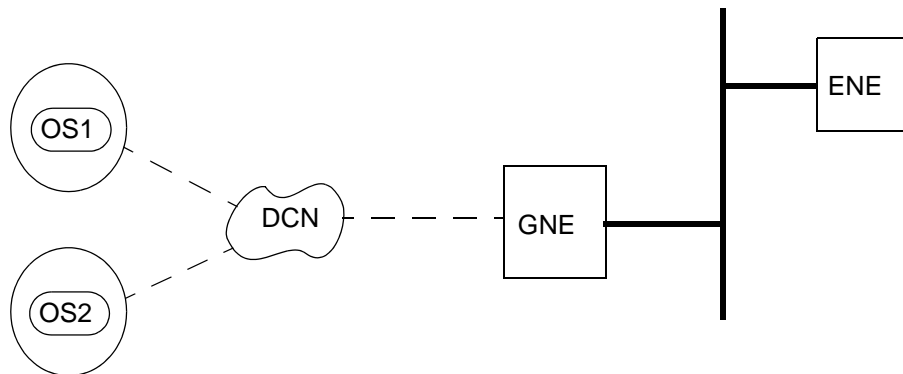
8.4.1.4 SONET LAN Interworking

Figure 8-9 shows an example architecture to illustrate this interworking case. In this example, OS1 and the ENE want to exchange messages. OS1 sends a TL1 message containing the TID of the ENE to the GNE. The GNE maps the TID of the ENE to its NSAP (using TID-NSAP information provided by TARP) and maps this NSAP to the

ENE's LAN Media Access Control (MAC) address. This NSAP-to-LAN MAC address mapping was learned by the GNE via the ES-IS routing protocol. The GNE then puts the PDU destined for the ENE onto the LAN. This is done using the Subnetwork Dependent Convergence Function (SND CF) for CLNP and ISO 8802 (LAN protocol) as described in ISO 8473. The ENE receives the PDU over the LAN and processes the TL1 message.

When the ENE wants to send a TL1 message to OS1, the ENE puts the PDU on the LAN using the MAC address and the layer 3 NSAP address of the GNE. (ES-IS is used to learn the LAN MAC address of the GNE.) If this message is a response to a TL1 request by an OS, the ENE would send the request using the application association on which the message was received. The GNE would then forward the message to the destination OS. If this was an autonomous message, the GNE would forward the message to the destination OS.

Figure 8-9 TL1/X.25 – LAN Interworking



8.4.1.5 SONET DCC Interworking

In this case, the only interworking that takes place is OS/NE interworking by the GNE. If an OS wants to send a message to a remote NE, it sends a TL1 message containing the TID of the remote NE to the GNE. The GNE maps the TID to the NE's NSAP address and forwards the message along the appropriate DCC.

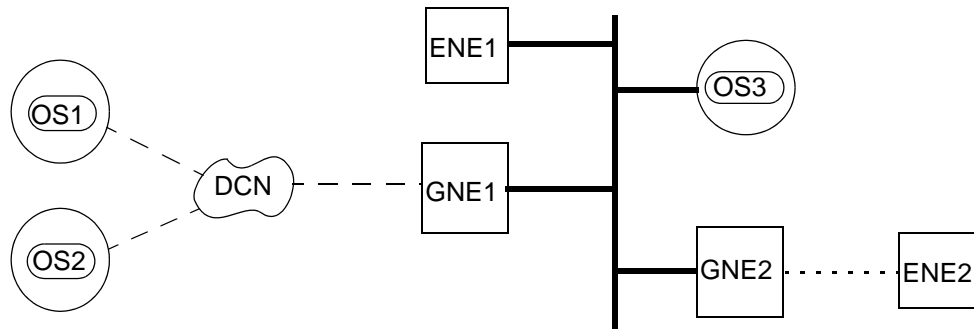
ES-IS and IS-IS are used to determine the links on which to forward messages. Responses or autonomous messages are forwarded from the remote NE to the OS using the modified multiple ACI method discussed earlier.

8.4.1.6 SONET LAN and DCC Interworking

Figure 8-10 contains an example illustrating this interworking case. In this example, OS1 and ENE2 want to exchange TL1 messages. For this to occur, two stages of interworking must take place: X.25-LAN interworking and LAN-DCC interworking. TID-to-NSAP mapping or LCN-to-NSAP mapping and multiple ACI processing is performed by GNE1, as appropriate.

In addition, the LAN MAC address mapping must be done. When OS1 sends GNE1 a TL1 message containing the TID of ENE2, GNE1 first looks up ENE2's NSAP. It then maps this NSAP to the MAC address of GNE2. This mapping is learned via the IS-IS routing protocol. GNE1 then sends the PDU over the LAN to GNE2 which routes it to ENE2.

Figure 8-10 TL1/X.25-LAN-DCC Interworking



8.4.2 TL1/X.25 [OS] – CMISE [SONET]

The OS/NE interface may be TL1 on a three-layer X.25 protocol stack. If the NE/NE interface is CMISE on an OSI stack, then the TL1/X.25 protocol stack must be interworked with the CMISE/OSI protocol stack. In addition, the GNE must perform message translation functions for both OS-to-NE and NE-to-OS messages. These translation functions are assumed to be supplier-specific.

8.4.3 CMISE [OS] – CMISE [SONET]

Since many SONET NE operations communications networks use CLNP (ISO 8473) and OSs often use the connection-mode network layer protocol X.25 (ISO/IEC 8208), a means to support instances of OS/NE communications must be specified. The most flexible approach is to use the network layer relay operational mode of the interworking functional unit described in ISO/IEC TR 10172:1991: *Information technology - Telecommunications and information exchange between systems - Network/transport protocol interworking specification*. This approach essentially extends the CLNP out from the SONET DCC, over the X.25 (ISO/IEC 8208) network, to the OS. It requires that the convergence functions described in ISO 8473 be included in the GNE and in the OS, and that the ISO 8473 protocol be included in the OS. In this method, the ISO 8473 PDUs are embedded in the data field of the X.25 (ISO/IEC 8208) packets. Packets arriving at the OS (or GNE) containing CLNP PDUs are distinguished from others by using the protocol identification conventions described in ISO/IEC TR 9577, *Information technology - Protocol identification in the network layer*. The OSI architecture (ISO/IEC 7498-1, *Information technology - Open System Interconnection Basic Reference Model: The Basic Model*), states that interworking between networks, such as those

described above for OSs and NEs, should take place at the network layer. The transport layer, and higher layers, should operate strictly on an end-to-end basis between ESSs. The requirements in [Section 8.3](#) specify the protocol stack required to achieve this interworking.

The previous discussion and proposals are valid for LAN NE/NE interfaces, DCC NE/NE interfaces and combinations of such interfaces (e.g., an OS/NE interface connected to a LAN NE/NE interface, which is then connected to a DCC NE/NE interface) when all of the interfaces support the target seven-layer protocol stacks. The ES-IS (ISO 9542) protocol maps local network-dependent addresses to layer 3 NSAPs and the IS-IS (ISO/IEC 10589) protocol determines routes for PDUs as they are transmitted between subnetworks. Specific requirements for support of the routing protocols are provided in [Section 8.5](#).

8.4.4 CMISE or TL1/OSI [SONET] – CMISE or TL1/OSI [SONET]

[Figure 8-10](#) can be used to illustrate DCC-LAN interworking for CMISE or TL1 for NE/NE or OS/NE transaction-oriented messages. In this example, ENE1 and ENE2 want to exchange messages, but the discussion also applies to an exchange of messages between OS3 and ENE2. For TL1 NE/NE messages, [Section 8.3.7.5](#) provides requirements on the use of ACIs and encoding of TL1 messages. Since both networks support CLNP at layer 3, the interworking between them is done at layer 3 and is independent of the application protocol used to exchange messages.

When ENE1 sends a message to ENE2, ENE1 puts the NSAP of ENE2 into the PDU sent on the LAN. The LAN MAC address used will be the address of either GNE1 or GNE2 (ES-IS is used to get the LAN MAC addresses of the GNEs). If GNE2 receives the PDU, it will forward it via the DCC to ENE2. If GNE1 receives the PDU, it will forward it to GNE2 (using IS-IS to determine that ENE2 is reachable via GNE2). GNE2 will then forward the PDU to ENE2. Lastly, GNE1 will send a Redirect message (an ES-IS function) to ENE1 noting that ENE2 is reachable via GNE2. This will cause ENE1 to route the next PDU to ENE2 via GNE2.

When ENE2 sends a message to ENE1, it puts the NSAP of ENE1 in the PDU and forwards it to GNE2 via the DCC. GNE2 then inserts the LAN MAC address of ENE1 into the PDU (using ES-IS to map ENE1's NSAP to its LAN MAC address) and places it on the LAN. ENE1 receives the PDU and processes it.

Note that the route-learning process described for LANs also applies to SONET DCC networks. The ES-IS and IS-IS routing protocols determine the existence, reachability, and “best routes” to other entities on a DCC-based operations network. IS-IS is also used to learn which ISs have connections to remote networks.

8.5 SONET Operations Communications Routing

8.5.1 Routing Overview

In SONET operations communications networks that utilize the Operations Communication Interfaces defined in [Section 8.3](#), ISO routing protocols (ISO 9542: ES-IS, and ISO/IEC 10589: IS-IS) are to be used for selective routing of network layer

data PDUs. These routing protocols automatically determine the “best” route to all destinations on the network. If link or node failures occur and alternate links are available, the ISO/IEC 10589 protocol can automatically reconfigure the routing information to route around the failure.⁹

ISO has defined an *administrative domain* as a collection of ESs, ISs and subnetworks operated by a single organization or administrative authority. Network providers will establish their own administrative domains, an example of which might be a region or a state. A *routing domain* is a set of ESs and ISs that operate according to the same routing procedures and is entirely contained within a single administrative domain. Administrative domains may be broken down into one or more routing domains.

Each routing domain can be hierarchically organized into routing subdomains, called *areas*. This is helpful when trying to maintain and process all of the information necessary to perform the routing function, as it can keep the size of the routing information base and the resources needed to compute routes reasonable. Each routing area maintains detailed routing information about its own internal composition and also maintains information that allows it to reach other routing areas. Since each routing area needs to maintain detailed routing information only about its own internal composition, the size of routing information bases is minimized. This, in turn, reduces the amount of data that needs to be exchanged and the overhead associated with computing routes within a routing domain.

Routing within an area is referred to as level 1 routing, while routing between areas is referred to as level 2 routing. Thus, Level 1 ISs keep track of the routing within their own areas, while Level 2 ISs keep track of routes to external areas. When an NPDU is sent to an external area, an ES (the source of the PDU) will send the NPDU to an IS in its area. This IS, a level 1 IS, then sends the NPDU to the nearest level 2 IS in its own area.¹⁰ The NPDU then travels via level 2 routing to the destination area, where it again travels via level 1 routing to the destination ES.

Figure 8-11 shows an example of a routing domain organized into three areas. An NPDU traveling from ES1 (in Area 1) to ES8 (in Area 3) will first go from ES1 to IS1. IS1 will then route the NPDU to IS4 via level 2 routing, and IS4 will route it to IS5 via level 1 routing. Finally, and IS5 will route the NPDU to ES8 (the destination system).

It is possible for an entity to fill several routing roles. In the previous example, IS1, IS2, IS3 and IS4 support level 1 and level 2 routing, while IS5 supports level 1 routing only. This example shows one way of partitioning a network into areas, and illustrates how level 1 and level 2 routing can be used. If the network in the example was configured as a single area, then only level 1 routing would be used.

Routing within a routing domain (e.g., as described above) is supported by the ISO routing information exchange protocols, ES-IS and IS-IS. The ES-IS protocol permits ESs and ISs to exchange configuration and routing information to facilitate routing and relaying. The IS-IS protocol permits ISs to exchange configuration and routing information to facilitate routing between ISs. In the example illustrated in

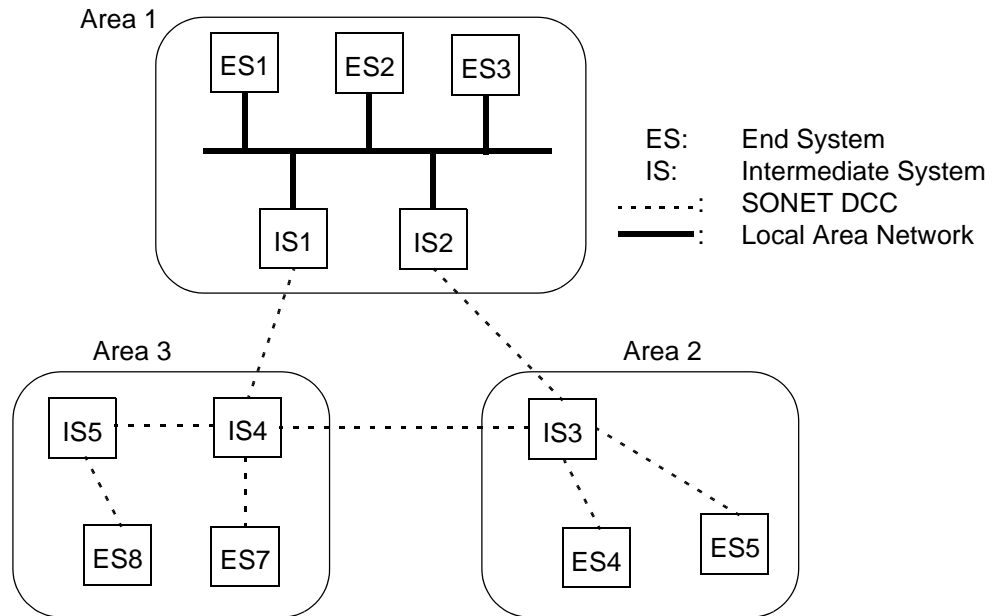
9. Similar protocols are used for routing in IP-based networks; however, the particular protocols to be used by SONET NEs that support interfaces to such networks (e.g., that support the right side of stack C in Figure 8-6) are FFS.

10. Since an IS may support both level 1 and level 2 routing, the nearest level 2 IS may be itself.

Figure 8-11, ES1, ES8, IS1 and IS5 support the ES-IS protocol. ES1 used it to discover IS1, and IS5 used it to discover ES8. All the ISs in that example support the IS-IS protocol, which supports both level 1 and level 2 routing.

Figure 8-7 shows the structure that is used to specify SONET NSAP addresses. A unique area address consists of the area field and all fields to the left of the area field in **Figure 8-7** (i.e., the area address includes the AFI, IDI, DFI, ORG, RES, RD and AREA). An IS uses the area address to determine the next hop for an NPDU by examining the destination NSAP field, which is part of the NPDU. If the area address of the destination NSAP matches its own area address, the IS will use level 1 routing. If the area address of the destination NSAP differs from its own area address, then the IS will use level 2 routing.

Figure 8-11 Example Routing Domain



8.5.2 ES-IS Requirements

The ES-IS protocol information is organized into two groups: Configuration Information and Redirection Information. Configuration Information is used by ESs to discover the existence and reachability of ISs, and it is used by ISs to discover the existence and reachability of ESs within the same area. Redirection Information is used by ISs to inform ESs of potentially better routes to a particular destination. Note that since the DCC is point-to-point, Redirection Information is not used, since an ES has only one route to an IS.

Redirection Information is used on a LAN. In **Figure 8-11**, for example, ESs in Area 1 have two ways to send PDUs out of the area. If ES1 were to send a PDU destined for ES7 (in Area 3) to IS2, IS2 would forward the PDU towards its destination and then send a Redirect PDU back to ES1 indicating that IS1 is a better choice for sending PDUs to ES7.

R8-102 [869v2] All SONET NEs shall support the ES-IS protocol, as specified in [Section C.5](#) of Appendix C, at NE/NE Operations Communications Interfaces as defined in [Section 8.3](#).

O8-103 [870] All SONET NEs supporting the Redirect capability in the ES role of the ES-IS protocol should support the ISO 9542 Address Mask Generation function.

O8-104 [871] All SONET NEs supporting the Redirect capability in the IS role of the ES-IS protocol should be able to send the ISO 9542 PDU Address Mask field.

The Address Mask appears only in ES-IS Redirect PDUs, and can be used to indicate that the redirection information applies to a larger population of NSAP addresses than the Destination Address contained in the Redirect PDU. In particular, the Address Mask establishes an equivalence class of NSAP addresses to which the same redirection information applies.

O8-105 [872] All SONET NEs supporting the ES or the IS role of the ES-IS protocol should be able to send and receive the ISO 9542 PDU Security field.

Note that the manner in which the security field may be used to augment routing security is an area FFS. In the future, the use of the security field may progress from an objective to a requirement.

8.5.3 IS-IS Requirements

R8-106 [873] All SONET NEs supporting the IS-IS protocol shall do so in accordance with the protocol specifications in ISO/IEC 10589 and Appendix C.

O8-107 [874] SONET NEs supporting the IS-IS protocol should authenticate IS-IS PDUs based on passwords as specified in ISO/IEC 10589.

For X.25-based OS/NE communications, the reachable address prefix capability of IS-IS can be used to facilitate routing of messages from an NE to an OS. This is done by assigning OSs NSAP addresses in an “external” routing area via the NSAP address prefix. When an NE sends a PDU to an OS, the level 1 routers in the SONET operations communication network recognize that the address is outside of their routing area and forward the PDU to a level 2 router, which forwards it to an OS/NE GNE. The OS/NE GNE maintains a table of NSAP reachable addresses for each OS and the OS’s corresponding X.121 address. Using this table, the GNE routes the message on to the destination OS.

R8-108 [875] A SONET GNE shall maintain static information to map NSAP addresses of OSs to their corresponding X.121 addresses. This static information shall be provisionable locally or remotely.


The reachable address prefix capability could also be of use if an organization wishes to divide its Administrative Domain into a number of separate Routing Domains. An Inter-Domain Routing Protocol (IDRP) is required to route between such routing domains; however, the suitability of any particular IDRP (e.g., ISO/IEC 10747) for use in SONET operations communications networks is an area FFS. To facilitate the construction of multi-domain topologies, inter-domain routing information can be made available to IS-IS. This information is in the form of a set of Reachable Address Prefixes, which may be manually provisioned or, in the future, automatically entered by an IDRP.

8.6 Craftsperson/NE Interfaces

The interfaces defined for craftsperson/NE communications are shown in Figure 8-12. Access to the local NE involves both the craftsperson/WS interface and the WS/NE interface (which are collectively referred to as the local craftsperson interface).

Remote Login is defined as the function or capability that enables a craftsperson to log into a remote NE that has one or more intermediate NEs between itself and the local WS (where the craftsperson is physically located). Remote Login would enable the craftsperson to connect to any remote SONET NE interconnected to the local SONET NE via the DCC or a LAN and appear as if the craftsperson were physically present at the remote NE. SIF-009-1997, *NE-NE Remote Login Implementation Requirements Specifications*, contains implementation requirements for a remote login function between a WS (Craft Interface Terminal) and a remote SONET NE. The DCC and/or LAN portions of the DCN provide the communications paths.

Figure 8-12 Craftsperson/NE Communications Network



8.6.1 Craftsperson/WS Interface

R8-109 [876] All operations functions supported by a SONET NE via OS/NE communications shall be supported at the local craftsperson interface.

R8-110 [877] Any features available exclusively at the local craftsperson interface shall be clearly identified in supporting documentation.

- R8-111 [878v2]** The workstation shall provide the craftsperson a command-line mode of operation (or interaction) as defined in TR-TSY-000824.
- R8-112 [879]** The command-line interface between the craftsperson and the workstation shall conform to TL1 requirements in GR-831-CORE, *OTGR Section 12.1: Operations Application Messages – Language for Operations Application Messages*, which is based on User System Language (USL) of TR-TSY-000825.
- R8-113 [880v2]** The user interface requirements specified in GR-826-CORE shall be supported (e.g., support of a graphical user interface).
-

8.6.2 WS/NE Interface

- R8-114 [881v2]** A WS/NE interface shall be provided in accordance with TR-TSY-000824.
- O8-115 [1039]** It is an objective that SONET NEs support the LAN-based interface defined by Section 4 of SIF-009-1997.
- R8-116 [1040v2]** SONET NEs shall support the WS/RNE interface defined by Section 5 of SIF-009-1997.
-

8.7 TARP

TARP is used on the NE/NE interface when there is a need to translate the TID of a TL1 message to the CLNP address (NSAP) of an NE.¹¹ Such a need may arise in (for example) the following scenarios:

- When the non-OSI TL1/X.25 protocol stack is used on the OS/NE interface (see [Section 8.4](#)), the GNE needs to be able to map TIDs to CLNP addresses for its subtending NEs
- If Remote Login (see [Section 8.6](#)) is initiated by entering the TID of the RNE at the local WS, the local NE needs to be able to map that TID to the CLNP address.

This TID-to-NSAP translation occurs by mapping TIDs to Network Entity Titles (NETs) and then deriving NSAPs from NETs by using the Network Selector Values specified in [Section 8.3.3](#).

TARP uses a selective PDU propagation methodology in conjunction with a distributed database (within NEs) of learned TID/NET mappings. TARP enables NEs to translate between TIDs and NETs by automatically exchanging mapping information with other NEs. No additional address provisioning is needed at the NE to support TARP.

11. TARP is not required to operate in an isolated subnetwork with no INEs present.

R8-117 [882] When a SONET NE supports TL1/OSI on the NE/NE interface (DCC or LAN), it shall also support TARP on the NE/NE interface according to the requirements of [Sections 8.7](#) and [C.8](#).

R8-118 [1041] For all TARP-related TID/NET mappings, case shall be ignored for the TIDs.

A SONET NE may have multiple NETs associated with the same TID. For such NEs, the following requirement applies:

R8-119 [883v2] When an NE supports TARP and has multiple NETs associated with the same TID, the NE shall designate one NET as the NET to be used for mapping purposes for all TARP-related TID/NET mappings.

8.7.1 Network Layer Protocol to Support TARP

The TARP PDU is carried by the standard ISO 8473 Data (CLNP DT) PDU. When sending TARP PDUs, TARP places some constraints on the values within the CLNP DT PDU header fields as specified below. When no TARP constraints are given, these fields will be used according to their specification in ISO 8473.

R8-120 [884] TARP PDUs shall be carried as CLNP DT PDUs.

R8-121 [885] When TARP PDUs are sent, the following constraints shall apply to the header of the CLNP DT PDU:

- A. The PDU Lifetime field shall be set to a value of 25000 milliseconds
 - B. The SP flag shall be set to a value of one (1) indicating that segmentation is permitted
 - C. The Error Report flag shall be set to a value of zero (0) indicating that discard of the PDU will not cause generation of an Error Report PDU.
-

8.7.2 TARP PDU Specification

This section specifies the TARP PDU fields that are carried in total by the *Data Part* of the CLNP DT PDU (see [Table 8-1](#)). The following subsections describe each of the TARP PDU fields.

R8-122 [886] The TARP PDU fields shown in [Table 8-1](#) shall be supported and sent in the order shown by [Table 8-1](#) (starting with tar-lif).

If a node receives a CLNP Service Data Unit (SDU) containing a TARP PDU where the length of the CLNP SDU is greater than N octets [where N is the length of a TARP PDU, which equals nine octets (the size of the fixed TARP PDU header) + the TID Target Length + TID Originator Length], the SONET NE may consider the TARP PDU invalid and discard it.

Table 8-1 TARP PDU Fields

TARP PDU Fields (within CLNP Data Part)	Abbreviation	Field Size (bytes)
TARP Lifetime	tar-lif	2
TARP Sequence Number	tar-seq	2
Protocol Address Type	tar-pro	1
Update Remote Cache (URC) and TARP Type Code	tar-tcd	1
TID Target Length	tar-tln	1
TID Originator Length	tar-oln	1
Protocol Address Length	tar-pln	1
TID of Target	tar-ttg	n = 0, 1, 2...
TID of Originator	tar-tor	n = 0, 1, 2...
Protocol Address of Originator	tar-por	n = 0, 1, 2...

8.7.2.1 TARP Lifetime (tar-lif)

The tar-lif field contains the TARP time-to-live in hops.

8.7.2.2 TARP Sequence Number (tar-seq)

The tar-seq field contains the TARP sequence number used for loop detection (see [Section 8.7.5.7](#)).

8.7.2.3 Protocol Address Type (tar-pro)

The tar-pro field is used to identify the type of protocol address to which the TID must be mapped. The value 'FE' (hex) will be used to identify the CLNP type of address (i.e., NET).

8.7.2.4 URC and TARP Type Code (tar-tcd)

The tar-tcd field consists of the URC bit (the first or most significant bit) and the TARP Type Code (the next 7 bits).

The value of the URC bit may be set to 0 or 1.¹²

12. The URC bit was originally intended to allow the PDU sender to signal the PDU receiver as to whether or not the receiver should update its local cache. Due to security concerns (i.e., the potential for fraudulent use of this feature), it is suggested that the URC feature no longer be used.

R8-123 [1042] The URC bit shall be ignored upon receipt of TARP PDUs.

The value of the TARP Type Code identifies the TARP Type of the PDU. Five TARP types are presently defined (see [Table 8-2](#)).

Table 8-2 TARP Types

TARP Type	Explanation
1	Request Protocol Address that matches tar-ttg; search Level 1 Routing Area
2	Same as Type 1, but also search Level 2 Routing Area
3	Response to a TARP request
4	Notification of TID or Protocol Address change
5	Request TID that matches Protocol Address (e.g., NET)

If a SONET NE receives a TARP PDU with a TARP Type Code value other than 1 through 5 (i.e., other than the standard values shown in [Table 8-2](#)), the SONET NE may consider the TARP PDU invalid and discard it.

8.7.2.5 TID Target Length (tar-tln)

The tar-tln field identifies the number of octets that are present in the tar-ttg field (see [Section 8.7.2.8](#)).

8.7.2.6 TID Originator Length (tar-oln)

The tar-oln field identifies the number of octets that are present in the tar-tor field (see [Section 8.7.2.9](#)).

8.7.2.7 Protocol Address Length (tar-pln)

The tar-pln field identifies the number of octets that are present in the tar-por field (see [Section 8.7.2.10](#)).

8.7.2.8 TID of Target (tar-ttg)

The tar-ttg field contains the TID value for the target NE.

8.7.2.9 TID of Originator (tar-tor)

The tar-tor field contains the TID value of the originator of the TARP PDU.

8.7.2.10 Protocol Address of Originator (tar-por)

The tar-por field contains the protocol address (for the protocol type identified in the tar-pro field) of the originator of the TARP PDU. When the tar-pro field is set to 'FE' (hex) (see [Section 8.7.2.3](#)), then tar-por will contain a CLNP address (i.e., the NET).

8.7.3 TARP Data Cache

A TARP Data Cache (TDC) may be provided in the SONET NE. When provided, the TDC consists of a set of values for the following triplet: (tar-pro, tar-tor, tar-por). A 4.1-kilobyte TDC could be used to hold approximately 100 entries. For the CLNP case, the TDC is essentially a database of TID-NET mappings.

8.7.4 NE Applications That Use the TARP Processor

Address Resolution and Address Change Notification are two NE applications that will use the TARP processor. These NE applications could be invoked from an OS or WS.

-
- R8-124 [887]** The NE shall process address resolution requests (from a higher-layer application in the NE) to find the NET that matches a given TID, according to the procedure given in [Section 8.7.4.1](#).
- R8-125 [888]** The NE shall process address resolution requests (from a higher-layer application in the NE) to find the TID that matches a given NET, according to the procedure given in [Section 8.7.4.2](#).
- R8-126 [889]** When a TID or Protocol Address change occurs at an NE, the NE shall notify other NEs of this change according to the procedure given in [Section 8.7.4.3](#).
-

8.7.4.1 Find NET That Matches TID

When an NE needs to find a NET for a given TID, the following procedure is used.

The TARP processor first checks its TDC for the match. If a match is found, the TARP processor would return the result to the requesting application. If no match is found, a TARP Type 1 PDU is originated per [Section 8.7.5.1](#). If Timer T1 (see [Table 8-3](#)) expires, a TARP Type 2 PDU is originated per [Section 8.7.5.2](#). Status information is passed back to the requesting application, indicating that the TARP Type 1 request has failed and that a TARP Type 2 request is being initiated. If Timer T2 expires, then Timer T4 is started, an error recovery routine is initiated, and status information is passed back to the requesting application indicating that error recovery is being initiated.

The error recovery routine is as follows. When Timer T4 expires, another TARP Type 2 PDU is originated and Timer T2 is again started. The *tar-seq* field of this PDU is set to zero, however, the sequence number at the NE is not reset. If Timer

T2 again expires, error information is passed back to the requesting application, indicating that the TID could not be resolved.

Table 8-3 TARP Timers

Timer	Description	Default (seconds)	Range (seconds)
T1	Waiting for response to TARP Type 1 request PDU	15	0 - 3600
T2	Waiting for response to TARP Type 2 request PDU	25	0 - 3600
T3	Waiting for response to Address Resolution request	40	0 - 3600
T4	Timer starts when T2 expires (used during Error Recovery)	20	0 - 3600

8.7.4.2 Find TID That Matches NET

When the NE has a NET and needs to find the matching TID, the following procedure takes place.

A TARP Type 5 PDU is originated per [Section 8.7.5.5](#). Timer T3 (see [Table 8-3](#)) is used. If this timer expires, no error recovery procedure occurs, and a status message is provided to indicate that the TID could not be found.

A scenario in which this may occur is where a Directory Server NE (DSNE) may want to populate its database. A DSNE would typically know which NETs it could communicate with and could then use TARP to learn the TIDs that correspond to those NETs.

8.7.4.3 Send Notification of TID or Protocol Address Change

When the NE needs to notify other NEs of a TID or Protocol Address change, the procedure is as follows.

The TARP Processor originates a TARP Type 4 PDU per [Section 8.7.5.4](#). The PDU's *tar-ttg* field contains the NE's TID value that existed prior to the change of TID or Protocol Address.

Note that there is no confirmation that other NEs have successfully received the address change information sent in the TARP Type 4 PDU.

8.7.5 TARP PDU Processing

The term “originate” is used below to refer to the generation of a TARP PDU by an NE in order to respond to a requesting application within that NE. This term is meant to exclude the propagation of TARP PDUs, which is separately described in [Section 8.7.5.8](#).

TARP PDU processing consists of originating and receiving TARP PDUs.

R8-127 [890] The NE shall provide the function of a TARP processor that is capable of originating and receiving TARP PDUs for all five TARP Types according to the descriptions given throughout [Section 8.7.5](#).

R8-128 [891v2] Each time an NE originates a TARP PDU, the NE shall increment the tar-seq field. The range of the tar-seq field shall be 0 to 65,535. If the tar-seq field reaches 65,535 (or if the NE is reset), a TARP Type 4 PDU shall be sent with a tar-seq field equal to zero, and the next TARP PDU shall be sent with the tar-seq field equal to one. A zero value will notify all other NEs that a reset has occurred.

Note that whenever tar-seq is reset to zero, a TARP Type 4 PDU is generated even if the TID and/or network address has not changed.

Various TARP PDUs must be disseminated to TARP processors on all neighboring systems. This implies that the NE is capable of identifying its neighbors. The list of neighboring network entities is obtained from the network layer Routing Information Base (RIB), which can also include entries created by provisioning.

The set of TARP adjacencies should contain an entry corresponding to each neighbor's NET or NSAP in the RIB. (For TARP purposes, the distinction between NETs and NSAPs is unimportant.) When transmitting a TARP PDU to a TARP adjacency, the PDU is sent using the N-UNITDATA Request primitive with an NSAP constructed by replacing the last octet of the neighbor system's NET or NSAP with the TARP N-SEL.

TARP adjacencies can also be provisioned. Note that a TARP adjacency is abstract and need not correspond to a specific data structure maintained by the TARP processor.

The term "adjacency" is used below to mean "TARP processor adjacency".

8.7.5.1 Origination of a TARP Type 1 PDU

When an NE originates a TARP Type 1 PDU, the PDU is sent to all adjacencies within the NE's routing area. Note that this implies that the NE is capable of identifying its adjacencies. Also note that the NE will typically have more adjacencies when a broadcast network is used (such as a LAN) than when a point-to-point network (such as the DCC) is used, and thus would need to send a greater number of PDUs.

R8-129 [1043] Inclusion of the tar-tor field in TARP Type 1 or Type 2 PDUs is optional (at the PDU sender's discretion). The receiver of TARP Type 1 or Type 2 PDUs shall ignore the contents of the tar-tor field (if present) and shall be capable of correctly processing the PDUs regardless of whether or not the tar-tor field is present.

8.7.5.2 Origination of a TARP Type 2 PDU

When an NE originates a TARP Type 2 PDU, the PDU is sent to adjacencies in all routing areas within the NE's routing domain. Note that typically only NEs that perform a Level 2 IS function will have adjacencies outside of their routing area.

8.7.5.3 Origination of a TARP Type 3 PDU

A TARP Type 3 PDU is a response to a TARP request PDU. The response is sent only to the originator of the request and thus does not use the TARP propagation procedure (e.g., the receiver of a TARP Type 3 PDU could ignore the *tar-lif* field). The *tar-ttg* field of the TARP Type 3 PDU is empty (i.e., zero length).

When an NE has multiple NETs (see **R8-119 [883v2]**) and is responding to a TARP Type 5 PDU for a NET other than the “designated NET”, the *tar-tor* field of the TARP Type 3 PDU is also empty (i.e., zero length).

8.7.5.4 Origination of a TARP Type 4 PDU

A TARP Type 4 PDU is a notification of a TID or Protocol Address change made at the NE that originates the notification (see [Section 8.7.4.3](#)). The PDU is sent to all adjacencies both within and outside the NE's routing area.

8.7.5.5 Origination of a TARP Type 5 PDU

When a TARP Type 5 PDU is sent, the CLNP destination address is known and thus the PDU is only sent to that address. Thus a TARP Type 5 PDU does not utilize the TARP propagation procedure (e.g., the receiver of a TARP Type 5 PDU could ignore the *tar-lif* field). The *tar-ttg* field of the TARP Type 5 PDU is empty (i.e., zero length).

8.7.5.6 Receipt of a TARP PDU

The following steps are taken by the TARP processor upon receipt of an incoming TARP PDU:

1. Check if *tar-lif* = 0; if so, discard TARP PDU
2. Check *tar-pro* to see if the Protocol Address Type is supported; if not supported, discard the TARP PDU
3. Check *tar-seq* and perform the Loop Detection Procedure (only when NE is an IS, see [Section 8.7.5.7](#)).

The next step depends on the TARP Type Code value and whether the NE is an ES or IS.

8.7.5.6.1 End Systems

4. If the TARP Type Code is 1 or 2, then check *tar-ttg*. If *tar-ttg* matches the ES's TID, then originate a TARP Type 3 PDU response.
5. If the TARP Type Code is 3, update TDC and pass response to the requesting application, unless the response is unsolicited and/or a duplicate response in which case the response should be discarded.
6. If the TARP Type Code is 4, then check to see if *tar-ttg* matches with TDC data. If there is a match, update TDC with the new information.

7. If the TARP Type Code is 5, originate a TARP Type 3 PDU response.
8. If the TARP Type Code is a value that is not supported by the NE, discard the PDU.

8.7.5.6.2 Level 1 Intermediate Systems

4. If the TARP Type Code is 1 or 2, check tar-ttg. If tar-ttg matches the IS's TID, originate a TARP Type 3 PDU response. If the tar-ttg does not match the IS's TID, perform Level 1 Propagation (see [Section 8.7.5.8](#)).
5. If the TARP Type Code is 3, update TDC and pass response to the requesting application, unless the response is unsolicited and/or a duplicate response in which case the response should be discarded.
6. If the TARP Type Code is 4, check to see if tar-ttg matches with TDC data. If there is a match, update TDC with the new information. In either case, perform Level 1 Propagation (see [Section 8.7.5.8](#)).
7. If the TARP Type Code is 5, originate a TARP Type 3 PDU response.
8. If the TARP Type Code is a value other than 1 through 5 inclusive, the PDU may be considered invalid and may be discarded.

8.7.5.6.3 Level 2 Intermediate Systems

4. If the TARP Type Code is 1, check tar-ttg. If tar-ttg matches the IS's TID, originate a TARP Type 3 PDU response. If tar-ttg does not match the IS's TID, perform Level 1 Propagation (see [Section 8.7.5.8](#)).
5. If the TARP Type Code is 2, check tar-ttg. If tar-ttg matches the IS's TID, originate a TARP Type 3 PDU response. If tar-ttg does not match the IS's TID, perform Level 1 and Level 2 Propagation (see [Section 8.7.5.8](#)).
6. If the TARP Type Code is 3, update TDC and pass response to the requesting application, unless the response is unsolicited and/or a duplicate response in which case the response should be discarded.
7. If the TARP Type Code is 4, check to see if tar-ttg matches with TDC data. If there is a match, update TDC with the new information. In either case, perform Level 1 and Level 2 Propagation (see [Section 8.7.5.8](#)).
8. If the TARP Type Code is 5, originate a TARP Type 3 PDU response.
9. If the TARP Type Code is a value other than 1 through 5 inclusive, the PDU may be considered invalid and may be discarded.

8.7.5.7 Loop Detection Procedure (Performed by ISs)

The Loop Detection Procedure (performed by ISs) is as follows.

Upon receipt of a TARP PDU other than Type 3 or Type 5, the NE checks its Loop Detection Buffer (LDB) for a *tar-por* match. If there is no match, the PDU will be processed and a new couplet entry (*tar-por*, *tar-seq*) is added to the LDB, and

if *tar-seq* is zero, a timer associated with the LDB entry is started using the provisionable LDB Entry Timer value. If there is a match, then *tar-seq* is compared to the LDB entry.

If *tar-seq* is non-zero and is \leq LDB entry, the PDU is discarded.

Otherwise, if *tar-seq* $>$ the LDB entry, the PDU is processed and the *tar-seq* field in the LDB entry is updated with the new value. The timer is not affected.

Otherwise, *tar-seq* must be zero. If the LDB entry timer is running, the PDU is discarded. If the timer is not running (i.e., expired), *tar-seq* in the LDB entry remains zero and the associated timer is started as described above.

When the LDB is being populated, only the System ID portion of the *tar-por* address needs to be used. A 4-kilobyte LDB could be used to hold approximately 500 entries. The LDB is flushed periodically in accordance with the LDB Flush Timer.

-
- R8-130 [893]** All NEs supporting an IS function shall maintain a circular (first-in first-out) TARP Loop Detection Buffer.
- R8-131 [1012]** All NEs supporting an IS function shall maintain a LDB Entry Timer for each LDB entry for which *tar-seq* = zero. The timer shall be settable within a range of 1 to 10 minutes. The default value shall be 5 minutes.
- R8-132 [894]** The LDB Flush Timer shall be settable within a range of 0 to 1440 minutes. The default value shall be 5 minutes.
-

8.7.5.8 Propagation Procedure (Performed by ISs)

The Propagation Procedure (performed by ISs) is as follows:

For *Level 1 Propagation*, PDUs are propagated to all adjacencies within the NE's routing area except as noted below.

For *Level 1 and 2 Propagation*, PDUs are propagated to all adjacencies both within and outside the NE's routing area within the NE's routing domain except as noted below.

The *Propagation Exception* is as follows. On either a point-to-point subnetwork or a broadcast subnetwork, PDUs are not propagated back to the NE from which the PDU was received.

For each NE to which a TARP PDU must be propagated, the NE constructs a new outgoing TARP PDU by decrementing the *tar-lif* field of the received PDU by one hop and providing new source and destination addresses in the appropriate CLNP header fields. If the decremented lifetime is 0, the system may discard the PDU without further processing, since it would be discarded by any receiving system. Note that the *tar-seq* field of the received PDU is not changed during propagation.

The conditions under which TARP PDUs are propagated are given in [Section 8.7.5.6.2](#) (for Level 1 ISSs) and in [Section 8.7.5.6.3](#) (for Level 2 ISSs). According to those conditions, either Level 1 Propagation or Level 1 and 2 Propagation is performed.

8.7.6 Management of the TARP Processor

At a minimum, SONET NEs need to provide the following capabilities to manage the TARP processor function within the NE.

-
- R8-133 [895]** The NE shall allow TARP propagation to be selectively disabled by link/adjacency.
 - R8-134 [896v2]** The TARP PDU fields listed in [Table 8-4](#) shall be provisionable in accordance with the default values and ranges specified by [Table 8-4](#). The values of other TARP PDU fields shall not be provisionable.
-

Table 8-4 Provisionable TARP PDU Fields

TARP PDU Field	Default	Range
tar-lif	100 hops	0 – 65,535
tar-pro	'FE' (hex)	1 byte

-
- R8-135 [897]** The NE shall allow the value of all TARP timers (as shown in [Table 8-3](#)) to be provisionable.
 - R8-136 [898]** The NE shall be capable of displaying (via the local WS at a minimum) the TDC, the LDB and the TARP Sequence Number in use.
 - R8-137 [899]** The NE shall provide a manual flush capability for the TDC and the LDB.
 - R8-138 [900v2]** The NE shall allow manual provisioning of entries for the TDC and the LDB.
 - R8-139 [901]** The NE shall allow the disabling of any of the following: all TARP functions, TARP propagation functions, TARP origination functions, or the TDC.
 - R8-140 [902]** The NE shall allow TARP requests to be manually generated.
-

8.7.7 TARP Echo Function

This section describes a TARP Echo Function (TEF) that may be used to aid in troubleshooting and can confirm layer-3 reachability of a CLNP address. The TEF may be invoked from either an OS or from the NE's WS. Invocation of the TEF will result in the NE sending a TARP Type 5 PDU.

When the TEF is invoked, the following information is supplied:

- Address (must be supplied as either a TID¹³, a System ID, or a NET)
- How many times to run the TEF (for this invocation)
- When to timeout (i.e., give up waiting for a response to the TEF)
- Format for returning results (e.g., results may include round-trip time(s), NET, TID, and TARP Lifetime count).

The response to a TEF invocation returns round-trip time(s) in milliseconds, and success statistics in percent successful. When the TEF is run more than once (for a given invocation), then round-trip times are returned as both individual times and the aggregate time.

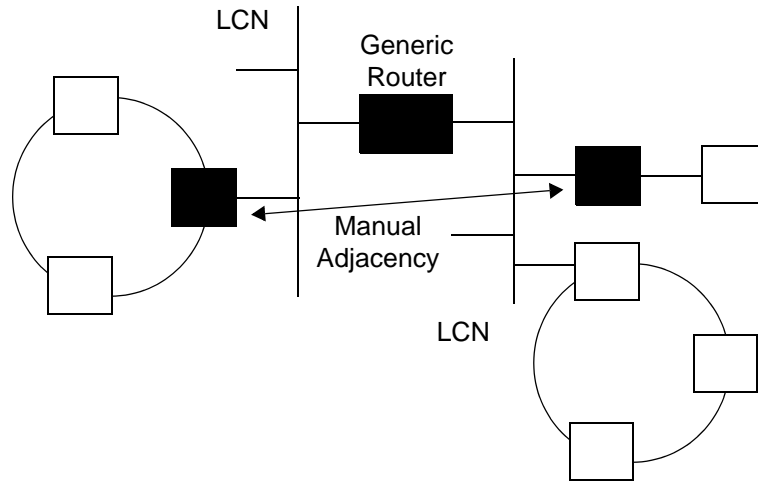
Note that if multiple TARP PDUs are outstanding, it is possible to mistake the response from a Type 1 or Type 2 TARP PDU as the response from the TEF. This could result in incorrect TEF information. For this reason, it is suggested that the number of outstanding TARP PDUs at any given time be limited to one.

8.7.8 Manual TARP Adjacencies

The use of non-SONET NEs without TARP capability (e.g., generic routers) could cause compatibility issues relating to TARP. Such devices may not have TIDs; however, TARP requests might need to cross a generic router. In such cases, the ability to provision a manual TARP adjacency in the SONET NE may be useful. This manual adjacency would in a sense cause a TARP request to hop through a generic router. This is depicted in [Figure 8-13](#).

¹³. If the address is supplied as a TID, the address resolution process described in [Section 8.7.4.1](#) would be performed before sending the TARP Type 5 PDU for TEF.

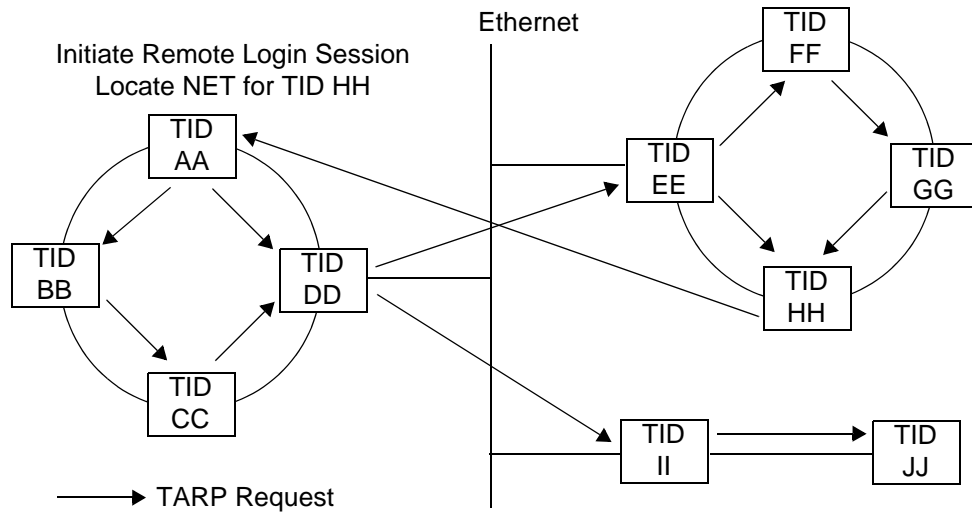
Figure 8-13 Manual TARP Adjacencies



8.7.9 TARP Example

Figure 8-14 illustrates an example of how TARP works. In this example, a Remote Login session is being initiated from the NE with TID AA to the NE with TID HH. The NE with TID AA originates a TARP Type 1 Request PDU. This request is propagated through the network until it reaches the NE with TID HH. At this point, the NE with TID HH originates a TARP Type 3 Response PDU, which is sent back to the NE with TID AA.

Figure 8-14 TARP Example



8.7.10 TARP Pseudocode

This section provides pseudocode as an aid to the reader in understanding the steps taken by the TARP processor upon receipt of an incoming TARP PDU. The normative description of these steps is provided in [Section 8.7.5.6](#).

BEGIN Pseudocode ALL

End System

Procedure EndSystem()

BEGIN

Check the TARP Lifetime

IF Lifetime has expired **THEN** Discard packet **END PROC**

ELSE (*Check to see if the protocol type is supported*)

IF Protocol Type is not supported **THEN** Discard Packet **END PROC**

ELSE (*Perform Type Analysis*)

CASE

TARP type is 1 and Target TID is my TID:

Construct Response (Type 3)

hand-off to Forward Process **END PROC**

TARP type is 2 and Target TID is my TID:

Construct Response (Type 3)

hand-off to Forward Process **END PROC**

TARP type is 3:

Add triplet to my cache and pass response to requesting application

Discard duplicate responses **END PROC**

TARP type is 4:

IF My cache has a TID match **THEN**

Update my cache with new information

Discard packet **END PROC**

ELSE My cache does not have a match

Discard packet **END PROC**

TARP type is 5:

Construct Response (Type 3)

hand-off to Forward Process **END PROC**

TARP type is not supported or Target TID is not my TID:

Discard packet **END PROC**

END;

END EndSystem;

Level 1 NE

Procedure Propagation(LEVEL)

LEVEL: (Level1, Level2, All)

BEGIN

Consult adjacency and routing database for adjacency information.
Decrement TARP packet lifetime by 1 hop.

IF tar-lif > 0 **THEN**

Construct packet(s) with new destination and source addresses
(PDU Header only, all adjacencies except adjacencies to NE from which
packet was received) and hand-off to Forwarding process.

END;

END Propagation()

Procedure Loop Detection()

BEGIN (*Loop Check*)

IF Tar 3 or 5 **THEN RETURN END PROC**

IF tar-por is a match **THEN** (*Check sequence number*)

BEGIN

IF pdu.tar-seq is non-zero and \leq ldb.tar-seq

OR pdu.tar-seq = 0 and LDB Entry Timer is running **THEN** discard pdu

ELSE IF pdu.tar-seq = 0 and LDB Entry Timer is not running **THEN**

ldb.tar-seq = 0

start LDB Entry Timer

ELSE pdu.tar-seq > ldb.tar-seq

ldb.tar-seq = pdu.tar-seq

END

ELSE tar-por is not a match.

Add couplet (tar-por, tar-seq) to LDB.

IF tar-seq = 0 **THEN** start LDB Entry Timer

END;

END Loop Detection()

Procedure Intermediate System1()

BEGIN

(*Check the TARP Lifetime and Run (Loop Detection Procedure)*)

IF Lifetime has expired and or LDB is a match **THEN** Discard packet

END PROC

ELSE (*Check to see if the protocol type is supported*)

IF Protocol Type is not supported **THEN** Discard Packet

END PROC

ELSE (*Perform Type Analysis*)

CASE

TARP type is 1 and Target TID is my TID:

Construct Response (Type 3)

hand-off to Forward Process **END PROC**

TARP type is 2 and Target TID is my TID:

Construct Response (Type 3)

hand-off to Forward Process **END PROC**

TARP type is 3:

Add triplet to my cache and pass response to requesting application.

Discard duplicate responses **END PROC**

TARP type is 4:

IF My cache has a match **THEN**

Update my cache with new information

BEGIN Propagation(All)

END Propagation **END PROC**

ELSE My cache does not have a match

Do not update cache

BEGIN Propagation(All)

END Propagation **END PROC**

TARP type is 5:

Construct Response (Type 3)

hand-off to Forward Process **END PROC**

TARP type is not supported or Target TID is not my TID:

BEGIN Propagation(Level1)

END Propagation **END PROC**

END;

END IntermediateSystem1;

Level 2 NEs

Procedure IntermediateSystem2()

BEGIN

Check the TARP Lifetime and Run (Loop Detection Procedure)

IF Lifetime has expired and or LDB is a match **THEN** Discard packet

END PROC

ELSE (*Check to see if the protocol type is supported*)

IF Protocol Type is not supported **THEN** Discard Packet

END PROC

ELSE (*Perform Type Analysis*)

CASE

TARP type is 1:


```
IF TID is my TID THEN  
    Construct Response (Type 3)  
    hand-off to Forward Process END PROC  
ELSE Target TID does not match  
    BEGIN Propagation(Level1)  
    END Propagation END PROC  
END;  
TARP type is 2:  
IF TID is my TID THEN  
    Construct Response (Type 3)  
    hand-off to Forward Process END PROC  
ELSE TID is not My TID  
    BEGIN Propagation(All)  
    END Propagation END PROC  
END;  
TARP type is 3:  
Add triplet to my cache and pass response to requesting application  
Discard duplicate responses END PROC  
TARP type is 4:  
IF My cache has a TID match THEN  
    Update my cache with new information  
    BEGIN Propagation(All)  
    END Propagation END PROC  
ELSE My cache does not have a TID match  
    Do not update cache  
    BEGIN Propagation(All)  
    END Propagation END PROC  
END;  
TARP type is 5:  
Construct Response (Type 3)  
hand-off to Forward Process END PROC  
TARP type is not supported:  
BEGIN Propagation(All)  
END Propagation END PROC  
END;  
END IntermediateSystem2;  
END Pseudocode ALL
```


Appendix A: Deleted Requirement-Object List

In Issues 1 and 2 of GR-253-CORE, this Appendix consisted of a list of all of the Requirement-Objects contained in the document. However, for a number of reasons it was decided that Requirement-Object Lists should no longer be included in Telcordia GRs. Therefore, the list that appeared as Appendix A in Issue 2 of this document was removed.

Subsequent to the release of Issue 3 of this GR, it was determined that a list of the Requirement-Objects that had been deleted from the document could be of use to the reader. Therefore, listed below are the numbers of the Requirement-Objects that were removed as of each Issue or Revision of the GR and, in the case of those removed as of Issue 4, the text of each Requirement-Object and an indication of the reason for its removal. Note that to avoid confusion between Requirement-Object numbers provided here and those in the rest of the document, only the absolute numbers (e.g., [48]) are included below. As discussed in Sections 1.2.1 and 1.3.3, each such number is permanently assigned to an object and is not changed if the object is moved in relation to other Requirement-Objects, or reused if the object is deleted. No local numbers (e.g., R3-1) are included here, as they would either be meaningless (if they indicated the deleted object's order of appearance in this Appendix) or duplicates of numbers that currently appear in other sections of the document (if they were set to indicate each deleted object's local number as of the last issue or revision of the GR in which it appeared).

A.1 Requirement-Objects Deleted Through Issue 3

CRx-x [48]	Deleted as of Issue 3.
CRx-x [50]	Deleted as of Issue 3.
Rx-x [51]	Deleted as of Issue 3.
CRx-x [143]	Deleted as of Issue 3.
Rx-x [218]	Deleted as of Issue 2.
Rx-x [219]	Deleted as of Issue 2.
Rx-x [220]	Deleted as of Issue 2.
Rx-x [221]	Deleted as of Issue 2.
Rx-x [228]	Deleted as of Issue 2.
Rx-x [229]	Deleted as of Issue 2.
Rx-x [230]	Deleted as of Issue 2.
Rx-x [231]	Deleted as of Issue 2.
Rx-x [232]	Deleted as of Issue 2.
Rx-x [233]	Deleted as of Issue 2.
CRx-x [234]	Deleted as of Issue 2.
CRx-x [235]	Deleted as of Issue 2.
Rx-x [236]	Deleted as of Issue 2.
Rx-x [237]	Deleted as of Issue 2.
Rx-x [238]	Deleted as of Issue 2.
CRx-x [244v2]	Deleted as of Issue 3.
CRx-x [245v2]	Deleted as of Issue 3.
CRx-x [251]	Deleted as of Issue 2, Revision 1.
CRx-x [252]	Deleted as of Issue 2, Revision 1.
Ox-x [263v2]	Deleted as of Issue 2, Revision 1.

Rx-x	[269]	Deleted as of Issue 2.
Rx-x	[270]	Deleted as of Issue 2.
Rx-x	[275]	Deleted as of Issue 2.
Rx-x	[299]	Deleted as of Issue 2.
Rx-x	[301]	Deleted as of Issue 2.
CRx-x	[303]	Deleted as of Issue 2.
CRx-x	[304]	Deleted as of Issue 2.
Rx-x	[306]	Deleted as of Issue 2.
Rx-x	[307]	Deleted as of Issue 2.
Rx-x	[308]	Deleted as of Issue 2.
Rx-x	[309]	Deleted as of Issue 2.
Ox-x	[344]	Deleted as of Issue 2.
Rx-x	[362]	Deleted as of Issue 2.
Rx-x	[363]	Deleted as of Issue 2.
Rx-x	[364]	Deleted as of Issue 2.
Rx-x	[365]	Deleted as of Issue 2.
Rx-x	[396]	Deleted as of Issue 2.
Rx-x	[401]	Deleted as of Issue 2.
Rx-x	[474]	Deleted as of Issue 2.
Rx-x	[475]	Deleted as of Issue 2.
Rx-x	[483]	Deleted as of Issue 2.
Rx-x	[484]	Deleted as of Issue 2.
Rx-x	[487]	Deleted as of Issue 2.
Rx-x	[490]	Deleted as of Issue 2.
Rx-x	[493]	Deleted as of Issue 2.
Rx-x	[494]	Deleted as of Issue 2.
Rx-x	[500]	Deleted as of Issue 2.
Rx-x	[503]	Deleted as of Issue 2.
Rx-x	[504]	Deleted as of Issue 2.
Rx-x	[507]	Deleted as of Issue 2.
Rx-x	[510]	Deleted as of Issue 2.
Rx-x	[511]	Deleted as of Issue 2.
Rx-x	[520]	Deleted as of Issue 2.
Rx-x	[522]	Deleted as of Issue 2.
Rx-x	[530]	Deleted as of Issue 2.
Rx-x	[532]	Deleted as of Issue 2.
Rx-x	[543]	Deleted as of Issue 2.
Rx-x	[545]	Deleted as of Issue 2.
Rx-x	[546]	Deleted as of Issue 2.
Rx-x	[563]	Deleted as of Issue 2.
Rx-x	[565]	Deleted as of Issue 2.
Rx-x	[575]	Deleted as of Issue 2.
Rx-x	[577]	Deleted as of Issue 2.
Rx-x	[590]	Deleted as of Issue 2.
Rx-x	[594]	Deleted as of Issue 2.
CRx-x	[600]	Deleted as of Issue 2.
Rx-x	[601]	Deleted as of Issue 2.
Rx-x	[602]	Deleted as of Issue 2.
Rx-x	[603]	Deleted as of Issue 2.
Rx-x	[604]	Deleted as of Issue 2.

Rx-x	[605]	Deleted as of Issue 2.
Ox-x	[606]	Deleted as of Issue 2.
Rx-x	[607]	Deleted as of Issue 2.
Ox-x	[608]	Deleted as of Issue 2.
Rx-x	[613]	Deleted as of Issue 2.
Rx-x	[614]	Deleted as of Issue 2.
Rx-x	[615]	Deleted as of Issue 2.
Rx-x	[616]	Deleted as of Issue 2.
Rx-x	[619]	Deleted as of Issue 2.
Rx-x	[624]	Deleted as of Issue 2.
Rx-x	[630]	Deleted as of Issue 2.
Rx-x	[631]	Deleted as of Issue 2.
Rx-x	[644]	Deleted as of Issue 2.
Ox-x	[646]	Deleted as of Issue 2, Revision 2.
Ox-x	[647]	Deleted as of Issue 2, Revision 2.
Rx-x	[648v2]	Deleted as of Issue 2, Revision 2.
Rx-x	[649]	Deleted as of Issue 2, Revision 2.
Ox-x	[650]	Deleted as of Issue 2, Revision 2.
Rx-x	[651]	Deleted as of Issue 2, Revision 2.
Rx-x	[652]	Deleted as of Issue 2.
Rx-x	[653]	Deleted as of Issue 2, Revision 2.
Rx-x	[654]	Deleted as of Issue 2.
Rx-x	[655]	Deleted as of Issue 2, Revision 2.
Ox-x	[665]	Deleted as of Issue 2, Revision 2.
Rx-x	[666]	Deleted as of Issue 2, Revision 2.
Rx-x	[671]	Deleted as of Issue 2, Revision 2.
Rx-x	[684]	Deleted as of Issue 2.
Rx-x	[692]	Deleted as of Issue 2.
Ox-x	[696]	Deleted as of Issue 2.
Rx-x	[697]	Deleted as of Issue 2.
Rx-x	[698]	Deleted as of Issue 2.
Ox-x	[699]	Deleted as of Issue 2.
Rx-x	[700]	Deleted as of Issue 2.
Rx-x	[701]	Deleted as of Issue 2.
Rx-x	[702]	Deleted as of Issue 2.
Rx-x	[704]	Deleted as of Issue 2.
CRx-x	[705]	Deleted as of Issue 2.
Ox-x	[708]	Deleted as of Issue 2.
Rx-x	[716]	Deleted as of Issue 3.
Rx-x	[726]	Deleted as of Issue 2.
Rx-x	[727]	Deleted as of Issue 2.
Rx-x	[733]	Deleted as of Issue 2.
Ox-x	[734]	Deleted as of Issue 2.
Rx-x	[739]	Deleted as of Issue 2, Revision 1.
Rx-x	[742]	Deleted as of Issue 2.
Rx-x	[752]	Deleted as of Issue 2.
Rx-x	[753]	Deleted as of Issue 2.
Rx-x	[754]	Deleted as of Issue 2.
Rx-x	[758]	Deleted as of Issue 2.
Rx-x	[769]	Deleted as of Issue 2.

Rx-x	[770]	Deleted as of Issue 2.
Ox-x	[775]	Deleted as of Issue 3.
Rx-x	[789]	Deleted as of Issue 3.
Rx-x	[815]	Deleted as of Issue 3.
Rx-x	[840]	Deleted as of Issue 2, Revision 1.
CRx-x	[892]	Deleted as of Issue 3.
Rx-x	[927]	Deleted as of Issue 2, Revision 2.
Ox-x	[929]	Deleted as of Issue 2, Revision 1.
Rx-x	[993]	Deleted as of Issue 2, Revision 2.
Ox-x	[998]	Deleted as of Issue 2, Revision 2.

A.2 Requirement-Objects Deleted as of Issue 4

A.2.1 Requirement-Objects Deleted From Section 3

The following requirements were merged into **R3-102 [72v2]**.

Rx-x [79] ATM cells shall be mapped into the STS-3c payload capacity by aligning the byte structure of every cell with the byte structure of the STS-3c SPE. The entire STS-3c payload capacity (i.e., 260 columns) shall be filled with cells, yielding a transfer capacity for ATM cells of 149.760 Mb/s.

Rx-x [84] ATM cells shall be mapped into the STS-12c payload capacity by aligning the byte structure of every cell with the byte structure of the STS-12c SPE. The entire STS-12c payload capacity (i.e., 1040 columns) shall be filled with cells, yielding a transfer capacity for ATM cells of 599.040 Mb/s.

Rx-x [1105] ATM cells shall be mapped into the STS-48c payload capacity by aligning the byte structure of every cell with the byte structure of the STS-48c SPE. The entire STS-48c payload capacity (i.e., 4160 columns) shall be filled with cells, yielding a transfer capacity for ATM cells of 2.396160 Gb/s.

A.2.2 Requirement-Objects Deleted From Section 4

The following objective was replaced by the discussion following **R4-5 [117v3]** (on page 4–8).

Ox-x [118] In Tables 4-10, 4-14 and 4-15, it is an objective for MLM transmitters to meet the narrower spectral width specifications for those applications that list two possible values for $\Delta\lambda_{\text{rms}}$.

A.2.3 Requirement-Objects Deleted From Section 5

The following requirement was combined with **R5-206 [298v4]** so that it will also apply to SONET NEs containing stratum 2, 3E or 3 clocks.

Rx-x [249v3] A SONET NE containing an SMC shall be capable of entering holdover when all of its timing references are determined to be failed (as per Section 5.4.6) or contain the “DON’T USE for Synchronization” synchronization status message.

The following requirements were considered redundant with other criteria.

Rx-x [250v2] Restoration from holdover shall be error-free.

Rx-x [290v2] The NE shall validate (as per Section 5.4.7.1) a degradation in the synchronization status message in the OC-N signal to or below the threshold, react as per Section 5.4.5.2.1, and generate AIS (if appropriate) on the derived DS1 timing output within 10 seconds.

Rx-x [293v2] When the synchronization status message in the OC-N used to derive the DS1 changes, the NE shall validate the change (as per Section 5.4.7.1), react as per Section 5.4.5.2.1, and insert the appropriate message in the derived DS1, both within 10 seconds.

The following criteria were replaced by references to equivalent criteria in GR-1244-CORE.

Rx-x [1014] For SONET NEs that contain stratum 2 internal clocks, the MTIE of the SONET outputs during the internal rearrangements listed above (i.e., the first six rearrangements listed) shall meet the “objective” mask in Figure 5-19.

CRx-x [1054] A SONET NE may be required to support a Forced Reference Switch command.

Rx-x [1055] If the Forced Reference Switch command is supported, the functionality of that command shall be clearly documented.

Rx-x [1056] If the Forced Reference Switch command is supported, a Reference Switch Clear command to clear a Forced Reference Switch shall also be supported.

CRx-x [1057] A SONET NE may be required to support a Lockout a Reference command.

Rx-x [1058] If the Lockout a Reference command is supported, its impact shall be to effectively cause the specified reference to be temporarily suspended from the NE’s provisioned timing reference list (i.e., until the corresponding Clear command is received).

Rx-x [1059] If the Lockout a Reference command is supported, a Reference Lockout Clear command shall also be supported. That command shall cause the NE to clear an existing Lockout a Reference for the specified reference.

Rx-x [1060] A SONET NE shall declare (and report to an OS) a standing condition when the first of one or more concurrent or consecutive Forced Reference Switch or Lockout a Reference commands is completed. The level of this condition shall be user-provisionable either as Not Alarmed or as a MN alarm. The condition shall be cleared (and a clear message sent to an OS) when all Forced Reference Switch and/or Lockout a Reference commands have been cleared.

The following objectives were merged into **R5-173 [1017v2]** and **R5-257 [343v4]**.

Ox-x [1112] Clocks that are timed from an incoming OC-12, OC-48 or OC-192 signal should meet the Category II jitter tolerance objective in Figure 5-28.

Ox-x [1121] A SONET NE's OC-12, OC-48 and OC-192 interfaces should tolerate, as a minimum, input jitter applied according to the "objective" mask in Figure 5-28 (with the parameter values specified in the figure for the particular rate signal).

A.2.4 Requirement-Objects Deleted From Section 6

The following requirement was deleted when **R6-450 [1004v2]** was changed from an objective to a requirement.

Rx-x [1003] As a minimum, an NE that compares the incoming path trace with a previously received path trace shall be capable of performing a case-sensitive comparison of ASCII-based path traces.

A.2.5 Requirement-Objects Deleted From Section 8

The following objective was considered unnecessary.

Ox-x [1035] A SONET NE should support the capability to communicate with an OS via a LAN interface.

The following criteria were merged into **R8-16 [790v3]**.

CRx-x [792v2] SONET NEs may, on an interim basis, be required to support TL1 as the application layer protocol for Interactive Class communications.

Rx-x [1136] If TL1 is supported as the application layer protocol for Interactive Class communications, it shall be supported in accordance to the criteria in [Section 8.3.7.5](#) of this document.

Appendix B: Fiber Optic Transmission System Design Worksheets

This appendix contains Worksheets 1 through 4, which are forms that can be used in gathering fiber optic transmission system design information. Worksheet 1 consists of general terminal equipment information the supplier provides. Worksheets 2 and 3 consist of terminal equipment parameters under normal operating and short-term emergency conditions, respectively, that the supplier and system integrator provide. Worksheet 4 summarizes the fiber optic cable transmission parameters to be provided by the system integrator and by the suppliers.

Some criteria are distance-related. If a supplier provides a complete system (NEs plus cable), then the average regenerator spacing may be used in verifying distance-related criteria. If a supplier provides only NEs (and is not given any specific distance information by the network provider), then the supplier is to show that the criteria are satisfied as the number of intermediate regenerators varies. For example, for certain criteria that are worded “up to 250 miles” the supplier is to verify that the criteria are satisfied for a number of intermediate regenerators (e.g., 10), spanning the range of applications up to 250 miles (assuming that the network provider provides the proper cable).

WORKSHEET 1: GENERAL INFORMATION

PAGE 1 OF 2

Supplier-Provided Information

System Information

Terminal Equipment Identification _____

Optical Line Rate (Mb/s) _____

Transmitter Information

General:

Identification _____

Optical Device Temperature Controller (e.g., TEC) _____

FDA Classification (e.g., Class I, Class II) _____

Product Change Designation (e.g., issue, revision) _____

Optical Source:

Type of Device (e.g., Laser, LED, etc.) _____

Material Composition of Source (e.g., In GaAs) _____

Generic Device Structure (e.g., DFB) _____

Type of Modulation (e.g., direct, MZ, EA) _____

Transmitter Connector:

Manufacturer _____

Type (e.g., SC, FC) _____

Model Number _____

Classification (Multimode, Single-mode) _____

Mating Connector Model Number _____

Transmitter Pigtail:

General Fiber Type _____

Class of Fiber _____

Mode Field Diameter _____ μm

Receiver Information

General:

Identification _____

Optical Device Temperature Controller (e.g., TEC) _____

Product Change Designation (e.g., issue, revision) _____

Optical Detector:

Type of Device (e.g., PIN, APD) _____

Material Composition of Detector (e.g., Ge, Si) _____

Receiver Pigtail:

General Fiber Type _____

Class of Fiber _____

Mode Field Diameter _____ μm

Receiver Connector:

Manufacturer _____

Type (e.g., SC, FC) _____

Model Number _____

Classification (Multimode, Single-mode) _____

Mating Connector Model Number _____

WORKSHEET 2: TERMINAL EQUIPMENT PARAMETERS PAGE 2 OF 2
STANDARD OPERATING CONDITIONS —
WORST-CASE VALUES

WDM Device:
 WDM Loss: $U_{WDM} = \underline{\hspace{2cm}} \text{ dB}$
 WDM Isolation: $Isol_{min} = \underline{\hspace{2cm}} \text{ dB}$

Passive Dispersion Compensator:
 Wavelength Range of Operation: $\lambda_{min}-\lambda_{max} = \underline{\hspace{2cm}} \text{ nm}$
 Maximum Dispersion Compensated for: $D_{max} = \underline{\hspace{2cm}} \text{ ps}$
 Minimum Dispersion Compensated for: $D_{min} = \underline{\hspace{2cm}} \text{ ps}$
 Attenuation or Gain: $U_{PDC} = \underline{\hspace{2cm}} \text{ dB}$
 $G_{PDC} = \underline{\hspace{2cm}} \text{ dB}$

Optical Fiber Amplifier:
 Wavelength Range of Operation: $\lambda_{min}-\lambda_{max} = \underline{\hspace{2cm}} \text{ nm}$
 Gain: $G_{OFA} = \underline{\hspace{2cm}} \text{ dB}$
 Minimum Input Power: $P_{in_{min}} = \underline{\hspace{2cm}} \text{ dBm}$
 Maximum Input Power: $P_{in_{max}} = \underline{\hspace{2cm}} \text{ dBm}$
 Minimum Output Power: $P_{out_{min}} = \underline{\hspace{2cm}} \text{ dBm}$
 Maximum Output Power: $P_{out_{max}} = \underline{\hspace{2cm}} \text{ dBm}$

Connectors:
 Loss: $U_{con} = \underline{\hspace{2cm}} \text{ dB}$
 Connector Variation: $\Delta U_{con} = \underline{\hspace{2cm}} \text{ dB}$
 Connector Reflectance: $OR_{con} = \underline{\hspace{2cm}} \text{ dB}$

Station Cable:
 Loss: $U_{SM} = \underline{\hspace{2cm}} \text{ dB/km}$
 Cutoff Wavelength: $\lambda_{cc} = \underline{\hspace{2cm}} \text{ nm}$
 PMD Coefficient: $PMD = \underline{\hspace{2cm}} \text{ ps}/(\text{km}^{0.5})$
 Chromatic Dispersion Coefficient: $D_{chrom} = \underline{\hspace{2cm}} \text{ ps}/(\text{nm}\cdot\text{km})$

System Integrator Provided Information

Nominal Central Wavelength: $\lambda_{Tnom} = \underline{\hspace{2cm}}$
 Safety Margin: $M = \underline{\hspace{2cm}} \text{ dB}$

Station Cable Length: Transmitter Receiver Total
 $\underline{\hspace{2cm}} \text{ km}$ $\underline{\hspace{2cm}} \text{ km}$ $I_{SM} = \underline{\hspace{2cm}} \text{ km}$

Number of Connectors: Transmitter³ Receiver⁴ Total
 $\underline{\hspace{2cm}}$ $\underline{\hspace{2cm}}$ $N_{con} = \underline{\hspace{2cm}}$

3. Excluding transmitter module connector.

4. Excluding receiver module connector.

WORKSHEET 3: TERMINAL EQUIPMENT PARAMETERS PAGE 2 OF 2
EXTENDED OPERATING CONDITIONS —
WORST-CASE VALUES

WDM Device:
 WDM Loss: $U_{WDM} = \underline{\hspace{2cm}} \text{ dB}$
 WDM Isolation: $Isol_{min} = \underline{\hspace{2cm}} \text{ dB}$

Passive Dispersion Compensator:
 Wavelength Range of Operation: $\lambda_{min} - \lambda_{max} = \underline{\hspace{2cm}} \text{ nm}$
 Maximum Dispersion Compensated for: $D_{max} = \underline{\hspace{2cm}} \text{ ps}$
 Minimum Dispersion Compensated for: $D_{min} = \underline{\hspace{2cm}} \text{ ps}$
 Attenuation or Gain: $U_{PDC} = \underline{\hspace{2cm}} \text{ dB}$
 $G_{PDC} = \underline{\hspace{2cm}} \text{ dB}$

Optical Fiber Amplifier:
 Wavelength Range of Operation: $\lambda_{min} - \lambda_{max} = \underline{\hspace{2cm}} \text{ nm}$
 Gain: $G_{OFA} = \underline{\hspace{2cm}} \text{ dB}$
 Minimum Input Power: $P_{in_{min}} = \underline{\hspace{2cm}} \text{ dBm}$
 Maximum Input Power: $P_{in_{max}} = \underline{\hspace{2cm}} \text{ dBm}$
 Minimum Output Power: $P_{out_{min}} = \underline{\hspace{2cm}} \text{ dBm}$
 Maximum Output Power: $P_{out_{max}} = \underline{\hspace{2cm}} \text{ dBm}$

Connectors:
 Loss: $U_{con} = \underline{\hspace{2cm}} \text{ dB}$
 Connector Variation: $\Delta U_{con} = \underline{\hspace{2cm}} \text{ dB}$
 Connector Reflectance: $OR_{con} = \underline{\hspace{2cm}} \text{ dB}$

Station Cable:
 Loss: $U_{SM} = \underline{\hspace{2cm}} \text{ dB/km}$
 Cutoff Wavelength: $\lambda_{cc} = \underline{\hspace{2cm}} \text{ nm}$
 PMD Coefficient: $PMD = \underline{\hspace{2cm}} \text{ ps}/(\text{km}^{0.5})$
 Chromatic Dispersion Coefficient: $D_{chrom} = \underline{\hspace{2cm}} \text{ ps}/(\text{nm}\cdot\text{km})$

System Integrator Provided Information

Nominal Central Wavelength: $\lambda_{Tnom} = \underline{\hspace{2cm}}$
 Safety Margin: $M = \underline{\hspace{2cm}} \text{ dB}$

Station Cable Length: Transmitter Receiver Total
 $\underline{\hspace{2cm}} \text{ km}$ $\underline{\hspace{2cm}} \text{ km}$ $I_{SM} = \underline{\hspace{2cm}} \text{ km}$

Number of Connectors: Transmitter⁷ Receiver⁸ Total
 $\underline{\hspace{2cm}}$ $\underline{\hspace{2cm}}$ $N_{con} = \underline{\hspace{2cm}}$

7. Excluding transmitter module connector.

8. Excluding receiver module connector.

**WORKSHEET 4: CABLE TRANSMISSION PARAMETERS
FOR A SPECIFIC APPLICATION**

PAGE 1 OF 1

System Integrator Provided Information

Application: Underground _____ Aerial _____ Buried _____

Temperature Range: _____ °C to _____ °C

Cabled Fiber Reel Length: $l_R =$ _____ km

Nominal Central Wavelength: $\lambda_{Tnom} =$ _____ nm

Central Wavelength Range: $\lambda_{Tmin} =$ _____ nm to $\lambda_{Tmax} =$ _____ nm

Splices:

Type of splice: _____

Splice loss at 23°C: $U_S =$ _____ dB/splice

Maximum additional splice loss due to temperature variation: $U_{ST} =$ _____ dB/splice

Supplier Provided Information

Cable Designation: _____

Maximum Cable Cutoff Wavelength (measured under the cable deployment conditions shown in [Figure 4-8](#):

$\lambda_{cc} =$ _____ nm

Cable Loss at λ_{Tnom} and 23°C:

$U_c =$ _____ dB/km

Maximum additional loss at 23°C due to wavelength variation:

$U_\lambda =$ _____ dB/km

Maximum additional loss at λ_T due to temperature variation:

$U_{CT} =$ _____ dB/km

Dispersion Parameters:

Zero-Dispersion Wavelength Range:

$\lambda_{0min} =$ _____ nm

to $\lambda_{0max} =$ _____ nm

Maximum Zero-Dispersion Slope:

$S_{0max} =$ _____ ps/(nm²·km)

Worst-Case Chromatic Dispersion Over Wavelength Range:

$D_{max} =$ _____ ps/(nm·km)

Worst-Case Polarization Mode Dispersion Coefficient Over Wavelength Range:

PMD = _____ ps/(km^{0.5})

Splice Related Parameters:

Mode Field Diameter:

Nominal: _____ μ m

Tolerance: _____ %

Cladding Diameter:

Nominal: _____ μ m

Tolerance: _____ %

Maximum Cladding Ovality:

_____ %

Maximum Core/Cladding Concentricity Error:

_____ μ m

Appendix C: SONET Operations Communications Lower Layers Protocol Profile

C.1 Introduction

The SONET profile specifies the implementation requirements for the SONET Operations Communications protocol stack. This lower-layer profile provides a set of tables for each of the protocols used in Layers 2 through 4 [i.e., LAPD (for the DCC), LLC (for the LAN), CLNP, IS-IS, ES-IS and TP4]. The structure and content of the tables within this Appendix are the same as those for each of the PICS proforma contained in the corresponding ISO or ITU-T standard. An additional SONET profile column has been added that specifies the Telcordia requirements for a given item. The support column in the tables can be used:

- By the protocol implementor, as a checklist to ensure more complete compliance with the SONET profile
- By the supplier and acquirer, or potential acquirer, of the implementation, as a detailed indication of the capabilities of the implementation
- By the user, or potential user, of the implementation, as a basis for initially checking the possibility of interworking with another implementation (note that while interworking can never be guaranteed, failure to interwork can often be predicted from incompatible profiles)
- By a protocol analyzer as the basis for selecting appropriate tests against which to assess the claim for conformance of the implementation.

C.1.1 Source Documents

When defining the SONET Lower Layers Profile, the following source documents were used.

C.1.1.1 Base Standards

- ITU-T Rec. Q.921
- ISO/IEC 8073: *Information technology - Open Systems Interconnection - Protocol for providing the connection-mode transport service*
- ISO/IEC 8473-1: *Information technology - Protocol for providing the connectionless-mode network service: Protocol specification*
- ISO/IEC 8802-2: *Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements - Part 2: Logical link control*
- ISO/IEC 8802-3: *Information technology - Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements - Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications*

- ISO 9542:1988: *Information processing systems - Telecommunications and information exchange between systems - End system to Intermediate system routing exchange protocol for use in conjunction with protocol for providing the connectionless-mode Network Service (ISO 8473)*
- ISO/IEC 10589:1992: *Information technology - Telecommunications and information exchange between systems - Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)*

C.1.1.2 International Standardized Profiles

- ISO/IEC ISP 10608-1:1992: *Information technology - International Standardized Profile TANnnn - Connection-mode Transport Service over Connectionless-mode Network Service - Part 1: General overview and subnetwork-independent requirements*
- ISO/IEC ISP 10608-2:1992: *Information technology - International Standardized Profile TANnnn - Connection-mode Transport Service over Connectionless-mode Network Service - Part 2: TA51 profile including subnetwork-dependent requirements for CSMA/CD Local Area Networks (LANs)*

C.1.1.3 American National Standards for Telecommunications

- ANSI T1.204

C.1.1.4 Telcordia Requirements

- [Section 8](#) of this GR
- GR-828-CORE

C.1.2 Notations Used in the SONET Lower Layer Profile

C.1.2.1 Status Symbols

Each PICS proforma has a set of symbols used in the status column of the PICS. These symbols appear (or are referred to) at the beginning of each protocol profile.

C.1.2.2 Profile Symbols

The profile symbols are the set of symbols used in the SONET profile column. These symbols present the view of Telcordia regarding a given item in the PICS proforma.

M mandatory support is required

- M:n** mandatory support is required given the condition n
- M(N)** mandatory support is required as clarified by additional comment N
- i** the item is out of scope for this profile
- i(N)** the item is out of scope for this profile as clarified by additional comment N
- X** the item is excluded (i.e., prohibited) for this profile
- X(N)** the item is excluded (i.e., prohibited) for this profile as clarified by additional comment N
- n/a** the item is not applicable for this profile
- O.n** the item is optional, but, if chosen, support is required for either at least one or only one of the options in the group labeled by the same numeral <n>
- GNE** the status following this symbol applies only when the NE is providing Gateway NE functionality as described in [Section 8.1.3](#)
- INE** the status following this symbol applies only when the NE is providing Intermediate NE functionality as described in [Section 8.1.4](#)
- ENE** the status following this symbol applies only when the NE is providing End NE functionality as described in [Section 8.1.5](#).

When an item is designated as 'i' (out of scope), implementation of that item is not precluded; however, for the purposes of the SONET Lower Layers Profile, the implementation or non-implementation of that item is ignored.

When an item is designated as 'n/a' (not applicable), that item is not used for SONET Operations Communications Applications. For the purpose of the SONET Lower Layers Profile the item is not used.

There are instances in the SONET profile where Telcordia has declared an item to be prohibited, while it is mandatory in the base standard. Telcordia is not stating that a supplier's OSI implementation should not have this feature; rather Telcordia is stating that the given item's use in SONET applications is prohibited.

Similarly, there are instances in the SONET profile where the comment "see note #", where # is a given note number, appears. These instances signify Telcordia objectives, and should be considered as such.

All conditionally mandatory functions in the protocol profile that are **always true** have been modified. The profile column no longer includes the conditional, and is made strictly mandatory. This is to reduce the confusion associated with searching to determine whether the condition is true. All truly conditional functions (where the conditional is out of scope for the profile) are left conditional in the profile column.

The SONET protocol profiles follow the standardized PICS proformas as closely as possible. There are places where there have been editorial modifications to the Standard PICS, allowing the profile to be more clearly understood. Some sections of the standardized PICS proformas have been deleted from the SONET protocol profiles. The sections of the standardized PICS proformas that are not contained within this document are to be viewed as either out of scope of the SONET protocol profile or not applicable for SONET applications. For example, the SONET DCC

uses only a subset of the LAPD capabilities. Thus much of the unused LAPD capabilities (e.g., Terminal End-point Identifier [TEI] management) are not applicable to SONET applications and are not included in this profile. Therefore, there may be parameters in each protocol profile that are out-of-sequence.

C.1.2.3 Support Symbols

The support symbols are the set of symbols used in the Support column. These symbols present the supplier's or implementor's view of the SONET profiles.

- Y** yes, the feature has been implemented;
- N** no, the feature has not been implemented;
- n/a** not applicable.

C.1.2.4 References

The reference columns in the following tables contain the reference or references to the material that specifies the item in the main body of the corresponding base standard.

C.2 LAPD SONET Protocol Profile

The following SONET profile specifies the implementation requirements for the Link Access Procedure on the D-Channel (LAPD). The protocol profile is based on the PICS proforma contained in ITU-T Rec. Q.921, with an additional SONET profile column.

C.2.1 Notations

C.2.1.1 Abbreviations

- APPX** Appendix
- DCC** Data Communications Channel
- DLCI** Data Link Connection Identifier, DLCI=(SAPI,TEI)
- DLE** Data Link Entity
- FR** prefix for Index of the Frames group
- IUT** Implementation Under Test
- NE** Network Element
- PC** prefix for the Index number of the Protocol Capabilities group
- SAPI** Service Access Point Identifier
- SP** prefix for the Index number of the System Parameter group

TEI Terminal End-point Identifier

C.2.1.2 Status Symbols

M mandatory

O optional

O.<n> optional, but, if chosen, support is required for either at least one or only one of the options in the group labeled by the same numeral <n>

X prohibited

C.2.1.3 Additional Symbols

<r> receive aspects of an item

<s> send aspects of an item

C.2.2 Protocol Capabilities (PC)

Index	Protocol Feature	Status	Ref.	Profile	Support
PC1.1	Is the <i>NE</i> of the non-automatic TEI assignment category?	O.1	3.3.4.2	M	
PC1.2	Is the <i>NE</i> of the automatic TEI assignment category?	O.1	3.3.4	i	
PC2	Does the <i>NE</i> support the broadcast data link?	M	5.2	M:1	
PC4	Does the <i>NE</i> support data link monitor function?	O	5.10	M	
PC5	Does the <i>NE</i> support reject retransmission procedure?	O	3.6.7, 5.8.1 Appx. I	i	
PC6.1	Does the DLE support automatic negotiation of data link layer parameters?	O.2	Appx. IV	i	
PC6.2	Does the DLE support internal parameter initialization?	O.2	5.4	M	
PC7	Does the <i>NE</i> permit concurrent LAPB data link connection within the D-channel?	O	2.3	n/a	
Service Access Point Identifier (SAPI)					
PC8	If the <i>NE</i> supports Layer 3 call control procedures, is SAPI=0 supported?	M	3.3.3	M:2	
PC9	If the <i>NE</i> supports X.25 Layer 3 packet procedures on the DCC, is SAPI=16 supported?	M	3.3.3	M:2	
PC10	Is SAPI=63 supported?	M	3.3.3	M:2	

O.1 = Support of at least one of these items is required.

O.2 = Support of at least one of these items is required.

Notes:

1. Support of the Broadcast Data Link functionality is mandatory, but the functionality is not used for SONET DCC applications.
2. SAPI subfields of the LAPD address field for LAPD D-channel applications do not apply to the SONET DCC applications. SAPI value of 62 is used for SONET DCC applications. The need to reserve additional SAPI values for specific purposes is FFS.

Index	Protocol Feature	Status	Ref.	Profile	Support
PC11.1	Does the implementation support the association of a given TEI with all SAPs which the <i>NE</i> supports?	O	3.3.4,5.3.1 (Q.920 3.4.3)	i	
PC11.2	If the <i>NE</i> is an X.31 type of packet mode terminal equipment, is a given TEI for point-to-point data link connection (<127) associated with all SAPs which the <i>NE</i> supports?	M	3.3.4,5.3.1 (Q.920 3.4.3)	M:3	
PC12	Does the implementation support modulus 128 for frames numbering?	M	3.5.2.1, 5.5.1	M	
Peer-to-Peer Procedures					
	<i>Unacknowledged Information Transfer</i>				
PC13	Does the <i>NE</i> support UI-command?	M	5.2.2	M	
PC14	Is the P/F bit set to 0?	M	5.1.1	M	
TEI Management					
PC15	Does the <i>NE</i> transmit management entity messages in UI frames with DLCI=(63,127)	O.3	5.3.1	i(4)	
Establishment and Release of Multiple Frame Operation					
PC32	Does the <i>NE</i> support multiple frame operation?	M	5.5	M	
	<i>Does the DLE initiate multiframe establishment</i>				
PC33.1	a) immediately after TEI assignment?	O.7	5.5	M	
PC33.2	b) when there is an incoming or an outgoing call?	O.7	5.5	n/a	
PC34.1	c) Does the DLE remain in TEI Assigned state when the multiple frame operation is released?	O.8	5.5.3	M	
PC34.2	d) Does the DLE initiate immediate re-establishment when the multiple frame operation is released?	O.8	5.5.3	M	

O.3 = Support of at least one of these items is required.

O.7 = Support of at least one of these items is required.

O.8 = Support of at least one of these items is required.

Notes:

3. See note 2.

4. Support of the TEI Management functionality is conditionally optional, but the functionality is not used for SONET DCC applications.

C.2.3 Frames - Protocol Data Units (FR)

Index	Protocol Feature	Status	Ref.	Profile	Support
Frame Format					
FR1	Format A	M	2.1	M	
FR2	Format B	M	2.1	M	
Flag Sequence					
FR3	Opening flag	M	2.2	M	
FR4	Closing flag	M	2.2	M	
Address Field					
FR5	Two octets	M	2.3	M	
FR6	If the DLE permits concurrent LAPB data link connection with the D-channel, is the one octet address field recognized?	M	2.3	n/a	
Control Field					
	<i>Unacknowledged operation</i>				
FR7	Single octet	M	2.4	M	
	<i>Multiple frame operation</i>				
FR8	Two octets	M	2.4	M	
FR9	Single octet (unnumbered frame)	M	2.4	M	
Order of Bit Transmission					
FR10	Ascending numerical order	M	2.8.2	M	
Field Mapping Convention					
FR11	Lowest bit number = Lowest order value	M	2.8.3	M	
	<i>Do all transmitted frames contain the following fields?</i>				
FR12.1	- Flag	M	2.2	M	
FR12.2	- Address	M	2.3	M	
FR12.3	- Control	M	2.4	M	
FR12.4	- FCS	M	2.7	M	
FR13	Is the <i>NE</i> capable of accepting the closing flag as the opening flag of the next frame?	M	2.2	M	
FR14	Does the <i>NE</i> generate a single flag as above?	O	2.2	i	
FR15	Does the <i>NE</i> ignore one flag, or two or more consecutive flags that do not delimit frames?	M	2.2	M	
FR16	Are all invalid frames discarded and no action taken?	M	2.9	M	
FR17	Are seven or more contiguous I bits interpreted as an abort and the associated frames ignored?	M	2.10	M	
FR18	If the <i>NE</i> supports the automatic negotiation of data link layer parameters, does the <i>NE</i> support XID frames?	M	Appx IV	n/a	

C.2.4 System Parameters (SP)

Index	Protocol Feature	Status	Ref.	Profile	Supported Range/ Default	Support
	<i>If the DLE supports multiple frame operation</i>					
SP1	Retransmission time (T200)	M	5.9.1	M	0.2 to 20 seconds/ 0.2 seconds	
SP2	Maximum number of retransmission	M	5.9.2	M	2 to 16 frames/ 3 frames	
	<i>Maximum number of octets in information field</i>					
SP3	for SAP supporting signaling	M	5.9.3	n/a	n/a	
SP4	for SAP supporting packet on the DCC	M	5.9.3	M	512 or greater octets/ 512 octets, see note 5	
	<i>Maximum number of outstanding I frames (k)</i>					
SP5	for SAP supporting basic access signaling	M	5.9.5	n/a	n/a	
SP6	for SAP supporting basic access packet on the DCC	M	5.9.5	M	1 to 127 frames/ 7 frames	
	<i>If the NE supports the data link monitor function,</i>					
SP9	Maximum time allowed without frames being exchanged (T203)	M	5.9.8	M	4 to 120 seconds/ 10 seconds	

Notes:

- For both I-Frame and UI-Frame operation

C.3 LLC Protocol Profile

The following SONET profile specifies the implementation requirements for the Logical Link Control Protocol (ISO/IEC 8802-2). The protocol profile is based on the PICS proforma contained in Annex C of ISO/IEC 8802-2:1989/Amd. 3, with an additional SONET profile column.

C.3.1 Abbreviations and Special Symbols

C.3.1.1 Status Symbols

- M** mandatory
- O** optional
- O.<n>** optional, but support of at least one of the group of options labeled by the same numeral <n> is required
- X** prohibited
- <item>** conditional symbol, status is dependent on the support marked for <item>

C.3.1.2 Item References

The following is a list of item references used in the PICS proforma:

- CLS** Class of LLC supported
- UI** UI PDUs
- XID** XID PDUs
- TES** TEST PDUs
- UIT** UI PDUs transmitted
- XDT** XID PDUs transmitted
- TST** TEST PDUs transmitted
- UIR** UI PDUs received
- XDR** XID PDUs received
- TSR** TEST PDUs received
- MIS** Miscellaneous

C.3.2 Claimed Conformance to ISO/IEC 8802-2:1989/Amd. 1, Amd. 2 and Amd. 4

Item	Protocol Feature	Ref.	Status	Profile	Support
CLS1a	Is Class 1 LLC supported?	4.2	O.1	M	
CLS1b	Are LLC Type 1 procedures supported?	4.2	CLS1a:M	M	
CLS2a	Is Class 2 LLC supported?	4.2	O.1	i	
CLS2b	Are LLC Type 1 and Type 2 procedures supported?	4.2	CLS2a:M	i	
CLS3a	Is Class 3 LLC supported?	4.2	O.1	i	
CLS3b	Are LLC Type 1 and Type 3 procedures supported?	4.2	CLS3a:M	i	
CLS4a	Is Class 4 LLC supported?	4.2	O.1	i	
CLS4b	Are LLC Type 1, Type 2 and Type 3 procedures supported?	4.2	CLS4a:M	i	

C.3.3 LLC Type 1 Operation - Unacknowledged Connectionless Mode

C.3.3.1 LLC Type 1 - Supported PDU Types

Item	Protocol Feature Supported PDU types	Ref.	Status	Profile	Support
UI/1	UI_CMD supported on transmission	6.1, 6.5.1	M	M	
UI/2	UI_CMD supported on receipt	6.1, 6.5.2	M	M	
XID/3	XID_CMD supported on transmission	6.6	O	i	
XID/4	XID_CMD supported on receipt	6.6	M	M	
XID/5	XID_RSP supported on transmission	6.6	M	M	
XID/6	XID_RSP supported on receipt	6.6	XID/3:M	XID/3:M	
TES/7	TEST_CMD supported on transmission	6.7	O	i	
TES/8	TEST_CMD supported on receipt	6.7	M	M	
TES/9	TEST_RSP supported on transmission	6.7	M	M	
TES/10	TEST_RSP supported on receipt	6.7	TES/7:M	TES/7:M	

C.3.3.2 LLC Type 1 - Supported Parameters in PDUs on Transmission

Item	Protocol Feature Supported PDU types	Ref.	Status	Profile	Support
UIT/11	UI_CMD - DSAP address	6.2	M	M	
UIT/12	UI_CMD - SSAP address	6.2	M	M	
UIT/13	UI_CMD - P-bit = 0	6.3	M	M	
UIT/14	UI_CMD - information	3.3	O	M	
XDT/15	XID_CMD - DSAP address	6.2, 6.6	XID/3:M	XID/3:M	
XDT/16	XID_CMD - SSAP address	6.2, 6.6	XID/3:M	XID/3:M	
XDT/17	XID_CMD - P-bit = 1	6.3	XID/3:O.2	XID/3:O.2	
XDT/18	XID_CMD - P-bit = 0	6.3	XID/3:O.2	XID/3:O.2	
XDT/19	XID_CMD - information	5.4.1.1.2, 6.6	XID/3:M	XID/3:M	
XDT/20	XID_RSP - DSAP address	6.2, 6.6	M	M	
XDT/21	XID_RSP - SSAP address	6.2, 6.6	M	M	
XDT/22	XID_RSP - F-bit= P-bit	6.3	M	M	
XDT/23	XID_RSP - information	5.4.1.2.1, 6.6	M	M	
TST/24	TEST_CMD - DSAP address	6.2	TES/7:M	TES/7:M	
TST/25	TEST_CMD - SSAP address	6.2	TES/7:M	TES/7:M	
TST/26	TEST_CMD - P-bit = 1	6.3	TES/7:O.3	TES/7:O.3	
TST/27	TEST_CMD - P-bit = 0	6.3	TES/7:O.3	TES/7:O.3	
TST/28	TEST_CMD - information	5.4.1.1.3, 6.7	TES/7:O	i	
TST/29	TEST_RSP - DSAP address	6.2	M	M	
TST/30	TEST_RSP - SSAP address	6.2	M	M	
TST/31	TEST_RSP - F-bit = P-bit	6.3	M	M	
TST/32	TEST_RSP - information	5.4.1.2.2, 6.7	M	M	

C.3.3.3 LLC Type 1 - Supported Parameters in PDUs on Receipt

Item	Protocol Feature Supported PDU types	References	Status	Profile	Support
UIR/33	UI_CMD - DSAP address	6.2	M	M	
UIR/34	UI_CMD - SSAP address	6.2	M	M	
UIR/35	UI_CMD - P-bit = 0	6.3	M	M	
UIR/36	UI_CMD - information	3.3	O	M	
XDR/37	XID_CMD - DSAP address	6.2, 6.6	M	M	
XDR/38	XID_CMD - SSAP address	6.2, 6.6	M	M	
XDR/39	XID_CMD - P-bit = 1	6.3	M	M	
XDR/40	XID_CMD - P-bit = 0	6.3	M	M	
XDR/41	XID_CMD - information	5.4.1.1.2, 6.6	M	M	
XDR/42	XID_RSP - DSAP address	6.2, 6.6	M	M:1	
XDR/43	XID_RSP - SSAP address	6.2, 6.6	M	M:1	
XDR/44	XID_RSP - F-bit = P-bit	6.3	M	M:1	
XDR/45	XID_RSP - information	5.4.1.2.1, 6.6	M	M:1	
TSR/46	TEST_CMD - DSAP address	6.2	M	M	
TSR/47	TEST_CMD - SSAP address	6.2	M	M	
TSR/48	TEST_CMD - P-bit = 1	6.3	M	M	
TSR/49	TEST_CMD - P-bit = 0	6.3	M	M	
TSR/50	TEST_CMD - information	5.4.1.1.3, 6.7	M	M	
TSR/51	TEST_RSP - DSAP address	6.2	M	M:2	
TSR/52	TEST_RSP - SSAP address	6.2	M	M:2	
TSR/53	TEST_RSP - F-bit = P-bit	6.3	M	M:2	
TSR/54	TEST_RSP - information	5.4.1.2.2, 6.7	TST/ 28:M	TST/ 28:M	

Notes:

1. Support is mandatory if XID_RSP on reception is supported (see XID/6).
2. Support is mandatory if TES_RSP on reception is supported (see TES/10).

C.3.3.4 LLC Type 1 - Miscellaneous

Item	Protocol Feature Supported PDU types	Ref.	Status	Profile	Support
MIS/55	Do all transmitted PDUs contain an integral number of octets	3.3	M	M	
	If the following PDUs are received from the MAC sub-layer are they treated as invalid and ignored:				
MIS/56	-contains a non-integral number of octets	3.3.5	M	M	
MIS/57	-has a length less than 3 octets	3.3.5	M	M	
	Which of the following addresses are supported in the DSAP address field of UI PDUs				

Item	Protocol Feature Supported PDU types	Ref.	Status	Profile	Support
MIS/58	-individual address	5.4.1.1.1	O.4	M	
MIS/59	-group address	5.4.1.1.1	O.4	M	
MIS/60	-global address	5.4.1.1.1	O.4	M	
MIS/61	-null address	5.4.1.1.1	O.4	M	
MIS/62	Is the address in the SSAP address field of a UI PDU the originator's individual address	5.4.1.1.1	M	M	
MIS/63	Are all UI PDU's transmitted as UI_CMD PDU's	6.5.1	M	M	
MIS/64	Are all UI_CMD PDU's transmitted with the P-bit = 0	6.5.1	M	M	
MIS/65	If a UI_CMD PDU is received with the P-bit = 1 is the PDU discarded?	6.3	O	i	
MIS/66	If a UI_RSP PDU is received is the frame discarded	6.5.2	M	M	
	Which of the following addresses are supported in the DSAP address field of XID_RSP PDUs				
MIS/73	-individual address	5.4.1.2.1	M	M	
MIS/74	-null address	5.4.1.2.1	M	M	
	Which of the following addresses are supported in the SSAP address field of XID_RSP PDUs				
MIS/75	-individual address	5.4.1.2.1	M	M	
MIS/76	-null address	5.4.1.2.1	M	M	
	Which of the following addresses are supported in the DSAP address field of TEST_RSP PDUs				
MIS/83	-individual address	5.4.1.2.2	M	M	
MIS/84	-null address	5.4.1.2.2	M	M	
	Which of the following addresses are supported in the SSAP address field of TEST_RSP PDUs				
MIS/85	-individual address	5.4.1.2.2	M	M	
MIS/86	-null address	5.4.1.2.2	M	M	
MIS/87	Is Duplicate Address Checking supported	6.9.2	O	i	
MIS/88	Is the ACK_TIMER function supported	6.9.2	MIS/87:M	MIS/87:M	
MIS/89	ACK_TIMER range			i	
MIS/90	Is the RETRY_COUNTER function supported	6.9.2	MIS/87:M	MIS/87:M	
MIS/91	RETRY_COUNTER range			i	
MIS/92	Is the XID_R_COUNTER function supported	6.9.2	MIS/87:M	MIS/87:M	

C.4 ISO 8473 Protocol Profile

The following SONET profile specifies the implementation requirements for the Connectionless-mode Network Layer Protocol (CLNP). The protocol profile is based on the PICS proforma contained in Annex A of ISO/IEC DIS 8473-1:1993, with an additional SONET profile column.

C.4.1 Notations

C.4.1.1 Status Symbols

M mandatory

O optional

O.<n> optional, but support of at least one of the group of options labeled by the same numeral <n> is required

X prohibited

<pred>:conditional-item symbol, including predicate identification

^ logical negation, applied to a conditional item's predicate

***** each item whose reference is used in a predicate or a predicate definition is indicated by an asterisk in the Item column

C.4.1.2 Additional Symbols

<r> receive aspects of an item

<s> send aspects of an item

C.4.2 Major Capabilities

Item	Capability	Ref.	Status	Profile	Support
*ES	End system		O.1	M	
*IS	Intermediate system		O.1	M:1	
FL-r	<r> Full protocol	6	M	M	
FL-s	<s> Full protocol	6	M	M	
NSS-r	<r> Non-segmenting subset	5.2	M	M	
*NSS-s	<s> Non-segmenting subset	5.2	IS:M, ^IS:O	X(2)	
*IAS-r	<r> Inactive subset	5.2	ES:O	M:3	
*IAS-s	<s> Inactive subset	5.2	IAS-r:M, ^IAS-r:X	X(4)	

Notes:

1. Support is mandatory for IS NEs.
2. Implementation of this function is conditionally mandatory (for IS implementations) or conditionally optional (for ^IS implementations) in the base PICS, but its **use** is **prohibited** in this profile.
3. Received NPDUs encoded with the inactive subset are discarded.
4. Implementation of this function is conditionally mandatory (for IAS-r implementations) or conditionally optional (for ^IAS-r implementations) in the base PICS, but its **use** is **prohibited** in this profile.

C.4.3 End Systems

C.4.3.1 Applicability

The profile items in [Section C.4.3](#) are applicable only to end system implementations; i.e., those in which item ES in [Section C.4.2](#) is supported.

C.4.3.2 Supported Functions

Item	Function	Ref.	Status	Profile	Support
ePDUC	PDU composition	6.1	M	M	
ePDUD	PDU decomposition	6.2	M	M	
ePHFA	Header format analysis	6.3	M	M	
ePDUL-s	<s> PDU lifetime control	6.4	M	M	
ePDUL-r	<r> PDU lifetime control	6.4	O	M	
eRout	Route PDU	6.5	M	M	
eForw	Forward PDU	6.6	M	M	
eSegm	Segment PDU	6.7	M	M	
eReas	Reassemble PDU	6.8	M	M	
eDisc	Discard PDU	6.9	M	M	
eErep	Error reporting	6.10	M	M	
eEdec-s	<s> Header error detection	6.11	M	M	
eEdec-r	<r> Header error detection	6.11	M	M	
*eSecu-s	<s> Security	6.13	O	i	
*eSecu-r	<r> Security	6.13	O	i	
*eCRR-s	<s> Complete route recording	6.15	O	i	
*eCRR-r	<r> Complete route recording	6.15	O	i	
*ePRR-s	<s> Partial route recording	6.15	O	i	
*ePRR-r	<r> Partial route recording	6.15	O	i	
*eCSR	Complete source routing	6.14	O	i	
*ePSR	Partial source routing	6.14	O	X	
*ePri-s	<s> Priority	6.17	O	i	
*ePri-r	<r> Priority	6.17	O	i	
*eQOSM-s	<s> QoS maintenance	6.16	O	M	
*eQOSM-r	<r> QoS maintenance	6.16	O	M	
*eCong-s	<s> Congestion notification	6.18	eQOSM-s:M	eQOSM-s:M, see note 5	
*eCong-r	<r> Congestion notification	6.18	O	see note 6	
*ePadd-s	<s> Padding	6.12	O	i	
ePadd-r	<r> Padding	6.12	M	M	
eEreq	Echo request	6.19	O	i(7)	
eErsp	Echo response	6.20	O	i(7)	
eSegS	Create segments smaller than necessary	6.8	O	i	

Notes:

5. Implementation of this function is conditionally mandatory (for eQOSM-s implementations) in the base PICS, but its **use** is a **Telcordia objective**.
6. This is a **Telcordia objective**. It is defined to be **out of scope** for this profile.
7. This is a new function introduced in the 1993 version of ISO/IEC 8473. To be compatible with ISO/IEC ISP 10608-1 (which is based on the 1988 version of ISO/IEC 8473), this new functionality is defined to be **out of scope** for this profile.

C.4.3.3 Supported PDUs

Item	NPDU	Ref.	Status	Profile	Support
eDT-t	DT (full protocol) transmit	7.7	M	M	
eDT-r	DT (full protocol) receive	7.7	M	M	
eDTNS-t	DT (non-segmenting) transmit	7.7	NSS-s:M	X	
eDTNS-r	DT (non-segmenting) receive	7.7	M	M	
eER-t	ER transmit	7.9	M	M	
eER-r	ER receive	7.9	M	M	
eIN-t	Inactive PDU transmit	7.8	IAS-s:M	X	
eIN-r	Inactive PDU receive	7.8	IAS-r:M	M(8)	
eERQ-t	ERQ transmit	7.10	eEreq:M	i(9)	
eERQ-r	ERQ receive	7.10	M	i(9)	
eERP-t	ERP transmit	7.11	eErsp:M	i(9)	
eERP-r	ERP receive	7.11	M	i(9)	

Notes:

8. See note 3.
9. See note 7. Note that implementation of this function may cause a backwards incompatibility with 1988 version of ISO/IEC 8473.

C.4.3.4 Supported Parameters

C.4.3.4.1 DT Parameters

Item	Parameter	Ref.	Status	Profile	Support
edFxPt-s	<s> Fixed part	7.2	M	M	
edFxPt-r	<r> Fixed part	7.2	M	M	
edAddr-s	<s> Addresses	7.3	M	M	
edAddr-r	<r> Addresses	7.3	M	M	
edSeg-s	<s> Segmentation part	7.4	M	M	
edSeg-r	<r> Segmentation part	7.4	M	M	
edPadd-s	<s> Padding	7.5.2	ePadd-s:M	ePadd-s:M	
edPadd-r	<r> Padding	7.5.2	M	M	
edSecu-s	<s> Security	7.5.3	eSecu-s:M	eSecu-s:M	
edSecu-r	<r> Security	7.5.3	eSecu-r:M	eSecu-r:M	
edCRR-s	<s> Complete route recording	7.5.5	eCRR-s:M	eCRR-s:M	
edCRR-r	<r> Complete route recording	7.5.5	eCRR-r:M	eCRR-r:M	
edPRR-s	<s> Partial route recording	7.5.5	ePRR-s:M	ePRR-s:M	
edPRR-r	<r> Partial route recording	7.5.5	ePRR-r:M	ePRR-r:M	
edCSR-s	<s> Complete source routing	7.5.4	eCSR:M	eCSR:M	
edPSR-s	<s> Partial source routing	7.5.4	ePSR:M	X	
edQOSM-s	<s> QoS maintenance	7.5.6	c1:M	M	
edQOSM-r	<r> QoS maintenance	7.5.6	c2:M	M	
edPri-s	<s> Priority	7.5.7	ePri-s:M	ePri-s:M	
edPri-r	<r> Priority	7.5.7	ePri-r:M	ePri-r:M	
edData-s	<s> Data	7.6	M	M	
edData-r	<r> Data	7.6	M	M	
edUnSup2	Are received PDUs containing parameters selecting unsupported Type 2 functions discarded and where appropriate an Error Report PDU generated?	6.21	M	M	
edUnSup3	Are parameters selecting unsupported Type 3 functions ignored?	6.21	M	M	

Definition of conditional status entries:

c1: eQOSM-s OR eCong-s

c2: eQOSM-r OR eCong-r

C.4.3.4.2 ER Parameters

Item	Parameter	Ref.	Status	Profile	Support
eeFxPt-s	<s> Fixed part	7.2	M	M	
eeFxPt-r	<r> Fixed part	7.2	M	M	
eeAddr-s	<s> Addresses	7.3	M	M	
eeAddr-r	<r> Addresses	7.3	M	M	
eePadd-s	<s> Padding	7.5.2	ePadd-s:M	ePadd-s:M	
eePadd-r	<r> Padding	7.5.2	M	M	
eeSecu-s	<s> Security	7.5.3	eSecu-s:M	eSecu-s:M	
eeSecu-r	<r> Security	7.5.3	eSecu-r:M	eSecu-r:M	
eeCRR-s	<s> Complete route recording	7.5.5	eCRR-s:M	eCRR-s:M	
eeCRR-r	<r> Complete route recording	7.5.5	eCRR-r:M	eCRR-r:M	
eePRR-s	<s> Partial route recording	7.5.5	ePRR-s:M	ePRR-s:M	
eePRR-r	<r> Partial route recording	7.5.5	ePRR-r:M	ePRR-r:M	
eeCSR-s	<s> Complete source routing	7.5.4	eCSR:M	eCSR:M	
eePSR-s	<s> Partial source routing	7.5.4	ePSR:M	X	
eeQOSM-s	<s> QoS maintenance	7.5.6	c1:M	M	
eeQOSM-r	<r> QoS maintenance	7.5.6	c2:M	M	
eePri-s	<s> Priority	7.5.7	ePri-s:M	ePri-s:M	
eePri-r	<r> Priority	7.5.7	ePri-r:M	ePri-r:M	
eeData-s	<s> Data	7.6	M	M	
eeData-r	<r> Data	7.6	M	M	
eeUnSup2	Are received PDUs containing parameters selecting unsupported Type 2 functions discarded?	6.21	M	M	
eeUnSup3	Are parameters selecting unsupported Type 3 functions ignored?	6.21	M	M	

Definition of conditional status entries:

c1: eQOSM-s OR eCong-s

c2: eQOSM-r OR eCong-r

C.4.3.4.3 Inactive Network Layer Protocol PDU Parameters

Item	Parameter	Ref.	Status	Profile	Support
eiNLPI-s	<s> Inactive network layer protocol identifier	7.8.2	IAS-s:M	X	
eiNLPI-r	<r> Inactive network layer protocol identifier	7.8.2	IAS-r:M	M(10)	
eiData-s	<s> Data	7.8.3	IAS-s:M	X	
eiData-r	<r> Data	7.8.3	IAS-r:M	M(10)	

Notes:

10. See note 3

C.4.3.5 Timers

Item	Timer	Ref.	Status	Values	Profile/ Supported Range & Default	Support/ Values supported
eLifReas	<i>Is reassembly timer \leq received derived PDU lifetime?</i>	6.8	M		M(11)	
eReasLim	<i>What values of the reassembly timer are supported?</i>	6.8		500 ms to 127.5 s	500 ms to 127.5 s/ Default of 12 s	

Notes:

- This corrects a defect in Table 9 of the 1993 version of T1.204.

C.4.4 Intermediate Systems

C.4.4.1 Applicability

The profile items in [Section C.4.4](#) are applicable only to intermediate system implementations; i.e., those in which item IS in [Section C.4.2](#) is supported.

C.4.4.2 Supported Functions

Item	Function	Ref.	Status	Profile	Support
iPDUC	PDU composition	6.1	M	M	
iPDUD	PDU decomposition	6.2	M	M	
iHFA	Header format analysis	6.3	M	M	
iPDUL	<s> PDU lifetime control	6.4	M	M	
iRout	Route PDU	6.5	M	M	
iForw	Forward PDU	6.6	M	M	
iSegm	Segment PDU	6.7	iDSNS:M	M	
iReas	Reassemble PDU	6.8	O	i	
iDisc	Discard PDU	6.9	M	M	
iErep	Error reporting	6.10	M	M	
iEdec	<s> Header error detection	6.11	M	M	
*iSecu	<s> Security	6.13	O	i	
*iCRR	<s> Complete route recording	6.15	O	i	
*iPRR	<s> Partial route recording	6.15	O	i	
*iCSR	Complete source routing	6.14	O	i	
*iPSR	Partial source routing	6.14	O	X	
*iPri	<s> Priority	6.17	O	i	
*iQOSM	<s> QoS maintenance	6.16	O	M	
*iCong	<s> Congestion notification	6.18	O	see note 12	
iPadd	<s> Padding	6.12	M	M	
iEreq	Echo request	6.19	O	i(13)	
iErsp	Echo response	6.20	O	i(13)	
iSegS	Create segments smaller than necessary	6.8	O	i	
iDSNS	Simultaneous support of subnetworks with different SN-User data sizes	Table 9 note 3	O	M	

Notes:

12. See note 6.

13. See note 7.

C.4.4.3 Supported PDUs

Item	NPDU	Ref.	Status	Profile	Support
iDT-t	DT (full protocol) transmit	7.7	M	M	
iDT-r	DT (full protocol) receive	7.7	M	M	
iDTNS-t	DT (non-segmenting) transmit	7.7	M	X(14)	
iDTNS-r	DT (non-segmenting) receive	7.7	M	M	
iER-t	ER transmit	7.9	M	M	
iER-r	ER receive	7.9	M	M	
iERQ-t	ERQ transmit	7.10	iEreq:M	i(15)	
iERQ-r	ERQ receive	7.10	M	i(15)	
iERP-t	ERP transmit	7.11	iErsp:M	i(15)	
iERP-r	ERP receive	7.11	M	i(15)	

Notes:

14. Implementation of this function is mandatory in the base PICS, but its **use** is **prohibited** in this profile.
15. See note 7. Note that implementation of this function may cause a backwards incompatibility with 1988 version of ISO/IEC 8473.

C.4.4.4 Supported Parameters

C.4.4.4.1 DT Parameters

Item	Parameter	Ref.	Status	Profile	Support
idFxFt-s	<s> Fixed part	7.2	M	M	
idFxFt-r	<r> Fixed part	7.2	M	M	
idAddr-s	<s> Addresses	7.3	M	M	
idAddr-r	<r> Addresses	7.3	M	M	
idSeg-s	<s> Segmentation part	7.4	M	M	
idSeg-r	<r> Segmentation part	7.4	M	M	
idPadd-s	<s> Padding	7.5.2	M	M	
idPadd-r	<r> Padding	7.5.2	M	M	
idSecu-s	<s> Security	7.5.3	iSecu:M	iSecu:M	
idSecu-r	<r> Security	7.5.3	iSecu:M	iSecu:M	
idCRR-s	<s> Complete route recording	7.5.5	iCRR:M	iCRR:M	
idCRR-r	<r> Complete route recording	7.5.5	iCRR:M	iCRR:M	
idPRR-s	<s> Partial route recording	7.5.5	M	M	
idPRR-r	<r> Partial route recording	7.5.5	iPRR:M	iPRR:M	
idCSR-s	<s> Complete source routing	7.5.4	iCSR:M	iCSR:M	
idCSR-r	<r> Complete source routing	7.5.4	iCSR:M	iCSR:M	
idPSR-s	<s> Partial source routing	7.5.4	M	X	
idPSR-r	<r> Partial source routing	7.5.4	iPSR:M	X	
idQOSM-s	<s> QoS maintenance	7.5.6	M	M	
idQOSM-r	<r> QoS maintenance	7.5.6	c1:M	M	
idPri-s	<s> Priority	7.5.7	M	M	
idPri-r	<r> Priority	7.5.7	iPri:M	iPri:M	
idData-s	<s> Data	7.6	M	M	
idData-r	<r> Data	7.6	M	M	
idUnSup2	Are received PDUs containing parameters selecting unsupported Type 2 functions discarded and where appropriate an Error Report PDU generated?	6.21	M	M	
idUnSup3	Are parameters selecting unsupported Type 3 functions ignored?	6.21	M	M	

Definition of conditional status entries: c1: iQOSM OR iCong

C.4.4.4.2 ER Parameters

Item	Parameter	Ref.	Status	Profile	Support
ieFxFt-s	<s> Fixed part	7.2	M	M	
ieFxFt-r	<r> Fixed part	7.2	M	M	
ieAddr-s	<s> Addresses	7.3	M	M	
ieAddr-r	<r> Addresses	7.3	M	M	
iePadd-s	<s> Padding	7.5.2	M	M	
iePadd-r	<r> Padding	7.5.2	M	M	
ieSecu-s	<s> Security	7.5.3	iSecu:M	iSecu:M	
ieSecu-r	<r> Security	7.5.3	iSecu:M	iSecu:M	
ieCRR-s	<s> Complete route recording	7.5.5	iCRR:M	iCRR:M	
ieCRR-r	<r> Complete route recording	7.5.5	iCRR:M	iCRR:M	
iePRR-s	<s> Partial route recording	7.5.5	M	M	
iePRR-r	<r> Partial route recording	7.5.5	iPRR:M	iPRR:M	
ieCSR-s	<s> Complete source routing	7.5.4	iCSR:M	iCSR:M	
ieCSR-r	<r> Complete source routing	7.5.4	iCSR:M	iCSR:M	
iePSR-s	<s> Partial source routing	7.5.4	M	X	
iePSR-r	<r> Partial source routing	7.5.4	iPSR:M	X	
ieQOSM-s	<s> QoS maintenance	7.5.6	M	M	
ieQOSM-r	<r> QoS maintenance	7.5.6	c1:M	M	
iePri-s	<s> Priority	7.5.7	M	M	
iePri-r	<r> Priority	7.5.7	iPri:M	iPri:M	
ieData-s	<s> Data	7.6	M	M	
ieData-r	<r> Data	7.6	M	M	
ieUnSup2	Are received PDUs containing parameters selecting unsupported Type 2 functions discarded?	6.21	M	M	
ieUnSup3	Are parameters selecting unsupported Type 3 functions ignored?	6.21	M	M	

Definition of conditional status entries: c1: iQOSM OR iCong

C.4.4.5 Timer and Parameter Values

Item	Timer	Ref.	Status	Values	Profile/Supported Range & Default	Support/Values supported
iLifReas	<i>Is reassembly timer ≤ received derived PDU lifetime?</i>	6.8	iReas:M		iReas:M	
iReasLim	<i>What values of the reassembly timer are supported?</i>	6.8		500 ms to 127.5 s	500 ms to 127.5 s / Default of 12 s	

C.5 ISO 9542 SONET Protocol Profile

The following SONET profile specifies the implementation requirements for the End system to Intermediate system routing exchange protocol for use in conjunction with the Protocol for providing the connectionless-mode network service (ISO 8473). The protocol profile is based on the PICS proforma contained in Annex A of ISO 9542: 1988, with an additional SONET profile column.

C.5.1 Notations

C.5.1.1 Status Symbols

M mandatory

O optional

X prohibited

CI: the status following this symbol applies only when the profile states that configuration information is supported.

RI: the status following this symbol applies only when the profile states that redirection information supported.

(CI∨RI): the status following this symbol applies only when the profile states that either configuration information or redirection information (or both) is supported.

C.5.1.2 Other Symbols

<r> receive aspects of an item

<s> send aspects of an item

C.5.2 PICS Proforma: ISO 9542(1988) – End System

Item	Protocol Function	Ref.	Status	Profile	Support
CI	Is configuration information supported?		O	ENE:M:1	
RI	Is redirection information supported?		O	ENE:M:2	

Notes:

1. Support is mandatory for both LAN and DCC operations communications.
2. Support is mandatory for LAN operations communications only.

C.5.2.1 Supported Functions

Item	Function	Ref.	Status	Profile	Support
CfRs	Configuration Response	6.6	M	M	
ErrP	Protocol Error Processing	6.13	(CI∨RI):M	M	
HCsV	PDU Header Checksum Validation	6.12	(CI∨RI):M	M	
HCsG	PDU Header Checksum Generation	6.12	O	i	
RpCf	Report Configuration	6.2, 6.2.1	CI:M	M	
RcCf	Record Configuration	6.3, 6.3.2	CI:M	M	
FICf	Flush Old Configuration	6.4	CI:M	M	
QyCf	Query Configuration	6.5	CI:M	M	
RcRd	Record Redirect	6.9	RI:M	RI:M	
FIRd	Flush Old Redirect	6.11	RI:M	RI:M	
RfRd	Refresh Redirect	6.10	RI:O	RI:M	
CfNt	Configuration Notification	6.7	CI:O	i	
CTPr	ESCT Processing	6.3.2	CI:O	M	
AMPr	Address Mask (only) Processing	7.4.5	RI:O	see note 3	
SMPr	Address Mask and SNPA Mask Processing	7.4.5, 7.4.6	RI:O	see note 3	

Notes:

3. This is a **Telcordia objective**. It is defined to be **out of scope** for this profile. This is different from the T1.204 profile.

C.5.2.2 Supported PDUs

Item	NPDU	Ref.	Status	Profile	Support
ESH-s	<s> End System Hello	7.1, 7.5	M	M	
ESH-r	<r> End System Hello	7.1, 7.5	CI:M	M	
ISH-r	<r> Intermediate System Hello	7.1, 7.5	CI:M	M	
RD-r	<r> Redirect	7.1, 7.5	RI:M	RI:M	

C.5.2.3 Supported Parameters

Item	Parameter	Ref.	Status	Profile	Support
FxPt-s	<s> Fixed Part	7.2.1-7.2.7	M	M	
FxPt-r	<r> Fixed Part	7.2.1-7.2.7	(CI∨RI): M	M	
SA-s1	<s> Source Address, one NSAP only	7.3.1, 7.3.2	O.1	M	
SA-r1	<r> Source Address, one NSAP only	7.3.1, 7.3.2	CI:M	M	
SA-sm	<s> Source Address, two or more NSAPs	7.3.3	O.1	M	
SA-rm	<r> Source Address, two or more NSAPs	7.3.3	CI:M	M	
NET-r	<r> Network Entity Title	7.3.1, 7.3.2, 7.34	(CI∨RI): M	M	
DA-r	<r> Destination Address	7.3.1, 7.3.2, 7.3.5	RI:M	RI:M	
BSNPA-r	<r> Subnetwork Address	7.3.1, 7.3.2, 7.3.6	RI:M	RI:M	
Scty-s	<s> Security	7.4.2	O	see note 4	
Scty-r	<r> Security	7.4.2	O	see note 4	
Pty-s	<s> Priority	7.4.4	O	i	
Pty-r	<r> Priority	7.4.4	O	i	
QoSM-r	<r> QoS Maintenance	7.4.3	RI:O	i	
AdMk-r	<r> Address Mask	7.4.5	RI:O	see note 5	
SNMK-r	<r> SNPA Mask	7.4.6	RI:O	see note 5	
ESCT-r	<r> Suggested ES Configuration Timer	7.4.7	CI:O	M	
OOpt-r	<r> (ignore) unsupported or unknown options		M	M	
OOpt-s	<s> Other options	7.4.1	X	X	

Notes:

4. This is a **Telcordia objective**. It is defined to be **out of scope** for this profile.
5. See note 3.

C.5.2.4 Supported Parameter Ranges

Item	Ranges	Ref.	Status	Profile/ Supported Range & Default	Support
HTv	What range of values can be set for the Holding Time field in transmitted PDUs?	6.1, 6.1.2	M	M/ 1 sec to 500 s & default of 105 s	
CTv	If configuration information is supported, what range of values can be set for the Configuration Timer?	6.1, 6.1.1	CI:M	M/ 1 sec to 200 s & default of 50 s	

C.5.3 PICS Proforma: ISO 9542(1988) – Intermediate System

Item	Protocol Function	Ref.	Status	Profile	Support
CI	Is configuration information supported?		O	INE:M:6 GNE:M:6	
RI	Is redirection information supported?		O	INE:M:7 GNE:M:7	

Notes:

- 6. See note 1.
- 7. See note 2.

C.5.3.1 Supported Functions

Item	Protocol Function	Ref.	Status	Profile	Support
ErrP	Protocol Error Processing	6.13	M	M	
HCsV	PDU Header Checksum Validation	6.12	M	M	
HCsG	PDU Header Checksum Generation	6.12	O	i	
RpCf	Report Configuration	6.2, 6.2.2	CI:M	M	
RpCf	Record Configuration	6.3, 6.3.1	CI:M	M	
FICf	Flush Old Configuration	6.4	CI:M	M	
RqRd	Request Redirect	6.8	RI:M	RI:M	
CfNt	Configuration Notification	6.7	CI:O	i	
CTGn	ESCT Generation	6.3.2	CI:O	M	
AMGn	Address Mask (only) Generation	6.8	RI:O	see note 8	
SMGn	Address Mask and SNPA Mask Generation	6.8	RI:O	see note 8	

Notes:

- 8. See note 3.

C.5.3.2 Supported PDUs

Item	NPDU	Ref.	Status	Profile	Support
ESH-r	<s> End System Hello	7.1, 7.5	CI:M	M	
ISH-r	<r> Intermediate System Hello	7.1, 7.6	CI:O	M	
ISH-s	<s> Intermediate System Hello	7.1, 7.6	CI:M	M	
RD-s	<s> Redirect	7.1, 7.7	RI:M	RI:M	
RD-r	<r> Redirect	6.9, 7.1, 7.7	M	M	

C.5.3.3 Supported Parameters

Item	Parameter	Ref.	Status	Profile	Support
FxPt-s	<s> Fixed Part	7.2.1-7.2.7	M	M	
FxPt-r	<r> Fixed Part	7.2.1-7.2.7	M	M	
SA-r	<s> Source Address, two or more NSAPs	7.3.1, 7.3.2, 7.3.3	CI:M	M	
NET-s	<s> Network Entity Title	7.3.1, 7.3.2, 7.3.4	M	M	
NET-r	<r> Network Entity Title	7.3.1, 7.3.2, 7.3.4	ISH-r:M	M	
DA-s	<s> Destination Address	7.3.1, 7.3.2, 7.3.5	RI:M	RI:M	
BSNPA-s	<s> Subnetwork Address	7.3.1, 7.3.2, 7.3.6	RI:M	RI:M	
Scty-s	<s> Security	7.4.2	O	see note 9	
Scty-r	<r> Security	7.4.2	O	see note 9	
Pty-s	<s> Priority	7.4.4	O	i	
Pty-r	<r> Priority	7.4.4	O	i	
QoSM-s	<s> QoS Maintenance	7.4.3	RI:O	i	
AdMk-s	<s> Address Mask	7.4.5	RI:O	see note 10	
SNMK-s	<s> SNPA Mask	7.4.6	RI:O	see note 10	
ESCT-s	<s> Suggested ES Configuration Timer	7.4.7	CI:O	M	
ESCT-r	<r>(ignore)Suggested ES Configuration Timer	7.4.7	ISH-r:M	M	
OOpt-r	<r> (ignore) unsupported or unknown options		M	M	
OOpt-s	<s> Other options	7.4.1	X	X	

Notes:

9. See note 4.

10. See note 3.

C.5.4 Supported Parameter Ranges

Item	Ranges	Ref.	Status	Profile/ Supported Range & Default	Support
HTv	What range of values can be set for the Holding Time field in transmitted PDUs?	6.1, 6.1.2	M	M/ 1 s to 500 s & default of 25 s	
CTv	If configuration information is supported, what range of values can be set for the Configuration Timer?	6.1, 6.1.1	CI:M	M/ 1 s to 200 s & default of 10 s	

C.6 ISO/IEC 10589 Protocol Profile

The following profile specifies the implementation requirements for the Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol providing the connectionless-mode network service (ISO 8473). The protocol profile is based on the PICS proforma contained in Annex A of ISO/IEC 10589:1992, with an additional SONET profile column.

C.6.1 Notations (Status Symbols)

M mandatory

O optional

O.<n> optional, but support of at least one of the group of options labeled by the same numeral <n> is required

X prohibited

c.<n> conditional requirement, according to condition <p>

- not applicable

* Items whose references are used in predicates are indicated by an asterisk in the Item column.

C.6.2 Protocol Summary: ISO/IEC 10589 General

Item	Functionality/Description	Ref.	Status	Profile	Support
AllIS	Are all basic IS-IS routing functions implemented?	12.1.2	M	M	
System Management	Is the system capable of being managed by the specified management information?	11	M	M(3)	
Authentication	Is PDU authentication based on passwords implemented?	7.3.7-7.3.10, 7.3.15.1-7.3.15.4, 8.2.3-8.2.4, 8.4.1.1	O	see note 1	
Default Metric	Is the default metric supported?	7.2.2, 7.2.6	M	M	
Delay Metric	Is the delay metric supported?	7.2.2, 7.2.6	O	i	
Expense Metric	Is the expense metric supported?	7.2.2, 7.2.6	O	i	
Error Metric	Is the error metric supported?	7.2.2, 7.2.6	O	i	
ID Field Length	What values of RouteingDomainIDLength (from the set 1-8) are supported by this implementation. Is the value configurable by system management?	7.1.3	M	M(2) NO	
Forwarding Rate	How many ISO 8473 PDUs can the implementation forward per second?	12.2.5.1.b	M	M	
Performance	Are the implementation performance criteria met?	12.2.5	M	M	

Notes:

1. This is a **Telcordia objective**. It is defined to be **out of scope** for this profile.
2. The value supported is 6 octets.
3. If TL1 is used as the application layer protocol for management, then the NE shall provide management functionality equivalent to that which CMISE provides.

C.6.2.1 System Environment: General

Item	Functionality/Description	Ref.	Status	Profile	Support
ISO 9542	Are the appropriate ISO 9542 operations implemented?	10.3, 8.2.1-8.2.2, 8.3.4, 8.4.5, 8.4.6	M	M	
Timer Jitter	Is jitter introduced in all periodic timers whose expiration causes transmission of a PDU?	10.1	M	M	

C.6.2.2 Subnetwork Dependent Functions: General

Item	Functionality/Description	Ref.	Status	Profile	Support
*LAN	Are the subnetwork dependent functions for broadcast subnetworks implemented?	8.4	O.1	M:4	
LAN IS Adjacencies	Are the LAN IS adjacency establishment operations implemented?	8.4.1-8.4.3	LAN:M	LAN:M	
LAN ES Adjacencies	Are the LAN ES adjacency establishment operations implemented?	8.4.6	LAN:M	LAN:M	
LAN DIS	Are the LAN designated IS operations implemented?	8.4.4, 8.4.5	LAN:M	LAN:M	
*8208 Static	Are the subnetwork dependent functions for ISO 8208 subnetworks implemented?	8.3	O.1	i	
8208 SNDCF	Are the ISO 8208 Subnetwork Dependent Convergence Functions implemented?	8.3.1, 8.3.2.1	C.1:M	i	
*PtPt	Are the subnetwork dependent functions for point-to-point subnetworks implemented?	8.2	O.1	M:5	
PtPt IS Adjacencies	Are the point-to-point IS adjacency establishment operations implemented?	8.2.2-8.2.5	C.2:M	PtPt:M	
PtPt ES Adjacencies	Are the point-to-point ES adjacency establishment operations implemented?	8.2.1	C.2:M	PtPt:M	
PtPt IIIH PDU	Are point-to-point IIIH PDUs correctly constructed and parsed?	9.7	C.2:M	PtPt:M	

C.1 if 8208 Static or 8208 DA then M else -

C.2 if PtPt or 8208 Static then M else -

Notes:

4. Support is **mandatory** for LAN.

5. Support is **mandatory** for DCC.

C.6.2.3 Update Process: General

Item	Functionality/Description	Ref.	Status	Profile	Support
LSP Periodic Generation	Is periodic generation of new local LSPs implemented?	7.3.2, 7.3.5, 7.3.13	M	M	
LSP Event Driven Generation	Is event driven generation of new local LSPs implemented?	7.3.6	M	M	
Pseudonode LSP Generation	Is generation of pseudonode LSPs implemented?	7.3.8, 7.3.10	LAN:M	LAN:M	
Multiple LSP Generation	Is multiple LSP generation implemented?	7.3.4	M	M	
LSP Propagation	Is propagation of LSPs implemented?	7.3.12, 7.3.14, 7.3.15.1, 7.3.15.5	M	M	
LSP Lifetime Control	Are the LSP lifetime control operations implemented?	7.3.16.4, 7.3.16.3	M	M	
CSNP Generation	Is the generation of CSNPs implemented?	7.3.15.3, 7.3.17	M	M	
PSNP Generation	Is the generation of PSNPs implemented?	7.3.15.4, 7.3.17	M	M	
SNP Processing	Are the sequence number PDU processing procedures implemented?	7.3.15.2, 7.3.17	M	M	
LSDB Overload	Are the LSP database overload operations implemented?	7.3.19	M	M	

C.6.2.4 Decision Process: General

Item	Functionality/Description	Ref.	Status	Profile	Support
Minimum Cost Path	Is computation of a single minimum cost path based upon each supported metric implemented?	7.2.6	M	M	
Equal Cost Paths	Is computation of equal minimum cost paths based upon each supported metric implemented?	7.2.6	O	i	
Downstream Paths	Is computation of downstream routes based upon each supported metric implemented?	7.2.6	O	i	
Multiple LSPs Recognition	Are multiple LSPs used only when a LSP with LSP#0 and remaining lifetime greater than 0 is present?	7.2.5	M	M	
Overloaded IS Exclusion	Are links to ISs with overloaded LSDBs ignored?	7.2.8.1	M	M	
Two Way Connectivity	Are links not reported by both end ISs ignored?	7.2.8.2	M	M	
Path Preference	Is the order of preference for path selection implemented?	7.2.12	M	M	
Excess Path Removal	Is removal of excess paths implemented?	7.2.7	M	M	
FIB Construction	Is the construction of ISO 8473 Forwarding Information Bases implemented?	7.2.9	M	M	

C.6.2.5 Forward/Receive Process: General

Item	Functionality/Description	Ref.	Status	Profile	Support
FIB Selection	Is selection of appropriate Forwarding Information Base implemented?	7.4.2	M	M	
NPDU Forwarding	Is forwarding of ISO 8473 PDUs implemented?	7.4.3.1, 7.4.3.3	M	M	
Receive Process	Are the basic receive process functions implemented?	7.4.4	M	M	

C.6.3 Protocol Summary: ISO/IEC 10589 Level 1 Specific Functions

Item	Functionality/Description	Ref.	Status	Profile/ Supported Range & Default	Support
*LIIS	Are Level 1 IS-IS routing functions implemented?	12.1.3	M	M	
Maximum Area Addresses	What values of maximumAreaAddresses are supported by this implementation?	7.1.5, 7.2.11	L1IS:M	M/ 0 to 12 & default of 3	
Area IS Count	How many ISs can this system support in a single area?	12.2.5	L1IS:M	M/ 1 to 512 & default of 512	
L1 Manual ES Adjacency	Are the manual ES adjacencies implemented?	7.3.31	L1IS:M	M	

C.6.3.1 Level 1 Subnetwork Dependent Functions

Item	Functionality/Description	Ref.	Status	Profile	Support
L1 LAN IIH PDU	Are L1 LAN IIH PDUs correctly constructed and parsed?	9.5	C.3:M	C.3:M	

C.3 if LIIS and LAN then M else -

C.6.3.2 Level 1 Update Process

Item	Functionality/Description	Ref.	Status	Profile	Support
L1 LS PDU	Are L1 LS PDUs correctly constructed and parsed?	9.8	L1IS:M	M	
L1 CSN PDU	Are L1 CSN PDUs correctly constructed and parsed?	9.10	L1IS:M	M	
L1 PSN PDU	Are L1 PSN PDUs correctly constructed and parsed?	9.12	L1IS:M	M	

C.6.3.3 Level 1 Decision Process

Item	Functionality/Description	Ref.	Status	Profile	Support
L1 Nearest L2 IS Identification	Is the identification of the nearest L2 IS implemented?	7.2.9.1	L1IS:M	M	
L1 Area Addresses Computation	Is the computation of area addresses implemented?	7.2.11	L1IS:M	M	

C.6.4 Protocol Summary: ISO/IEC 10589 Level 2 Specific Functions

Item	Functionality/Description	Ref.	Status	Profile/ Supported Range & Default	Support
*L2IS	Are level 2 IS-IS routing functions implemented?	12.1.4	O	M:6	
IS Count	What is the total number of ISs that this L2 IS can support?	12.2.5	L2IS:M	L2IS:M/ 1 to 512 & default of 512 (see note 7)	
L2IS Count	How many level 2 ISs does this implementation support?	12.2.5.1	L2IS:M	L2IS:M/ 1 to 512 & default of 256 (see note 7)	
*RA Prefix	Are Reachable Address Prefixes supported on circuits?	8.1, 7.3.3.2	L2IS:O	L2IS:M	
External Metrics	Are external metrics supported?	7.2.2, 7.2.12, 7.3.3.2	RA Prefix:M	RA Prefix:M	
*Partition	Is level 1 partition repair implemented?	7.2.10	L2IS:O	see note 8	

Notes:

6. This function is **mandatory** when the Level 2 functions are supported.
7. Note that these numbers are preliminary and are subject to further study and possible change.
8. This is a **Telcordia objective**. It is defined to be **out of scope** for this profile.

C.6.4.1 Level 2 Subnetwork Dependent Functions

Item	Functionality/Description	Ref.	Status	Profile	Support
L2 LAN IIIH PDU	Are L2 LAN IIIH PDUs correctly constructed and parsed?	9.6	C.4:M	C.4:M	
*8208 DA	Are ISO 8208 Dynamic Assignment circuits implemented?	8.3	O.1	i	
RA Adjacency Management	Are the reachable address adjacency management operations implemented?	8.3.2.2-8.3.5.6	8208 DA:M	8208 DA:M	
Call Establishment Metric Increment	Are non-zero values of the callEstablishment-MetricIncrement supported?	8.3.5	8208 DA:O	i	
Reverse Path Cache	Is 8208 reverse path cache implemented?	8.3.3	8208 DA:O	i	

C.4 if L2IS and LAN then M else -

C.6.4.2 Level 2 Update Process

Item	Functionality/Description	Ref.	Status	Profile	Support
L2 LS PDU	Are L2 LS PDUs correctly constructed and parsed?	9.9	L2IS:M	L2IS:M	
L2 CSN PDU	Are L2 CSN PDUs correctly constructed and parsed?	9.11	L2IS:M	L2IS:M	
L2 PSN PDU	Are L2 PSN PDUs correctly constructed and parsed?	9.13	L2IS:M	L2IS:M	

C.6.4.3 Level 2 Decision Process

Item	Functionality/Description	Ref.	Status	Profile	Support
L2 Attached Flag	Is the setting of the attached flag implemented?	7.2.9.2	L2IS:M	L2IS:M	
L2 Partition DIS election	Is the election of partition L2 DIS implemented?	7.2.10.2	Partition : M	Partition: M	
L2 Partition Area Addresses Computation	Is the computation of L1 partition area addresses implemented?	7.2.10.3	Partition : M	Partition: M	
L2 DIS Partition Repair	Is partition detection and repair via virtual L1 links implemented?	7.2.10.1	Partition : M	Partition: M	

C.6.4.4 Level 2 Forward/Receive Process

Item	Functionality/Description	Ref.	Status	Profile	Support
L2 NPDU Encapsulation	Is the encapsulation of NPDU implemented?	7.2.10.4, 7.4.3.2	Partition: M	Partition: M	
L2 NPDU Decapsulation	Is the decapsulation of NPDU implemented?	7.4.4	Partition: M	Partition: M	

C.7 ISO/IEC 8073 Protocol Profile

The following profile specifies the implementation requirements for the Connection-mode Transport Protocol. The protocol profile is based on the PICS proforma contained in Annex C of ISO/IEC 8073-1:1992, with an additional SONET profile column.

C.7.1 Notations

C.7.1.1 Status Symbols

M mandatory

O optional to implement. If implemented, the feature may or may not be used.

O.<n> optional, but support of at least one of the group of options labeled by the same numeral <n> is required.

<index>: this predicate symbol means that the status following it applies only when the profile states that the feature identified by the index is supported. In the simplest case, <index> is the identifying tag of a single profile item. <index> may also be a Boolean expression composed of several indices.

<index>:: when this group predicate is true, the associated clause should be completed.

C.7.2 Protocol Implementation for TP4/CLNS (C4L::)

C.7.2.1 Annex B – NCMS

Index	Class	Ref.	Status	Profile	Support
A1	Network connection management procedures	Annex B	O	i	

C.7.2.2 Classes Implemented

Index	Class	Ref.	Status	Profile	Support
C4L	Class 4 operation over CLNS	14	O	M	

C.7.3 Initiator/Responder Capability for Protocol Classes 0-4

Index	Item	Ref.	Status	Profile	Support
IR1	Initiating CR TPDU	14.5 a)	O.2	M	
IR2	Responding to CR TPDU	14.5 a)	O.2	M	

C.7.4 Supported Functions

C.7.4.1 Supported Functions for Class 4 (C4L::)

The following functions are mandatory.

Index	Function	Ref.	Status	Profile	Support
T4F1	TPDU transfer	6.2	M	M	
T4F2	Segmenting	6.3	M	M	
T4F3	Reassembling	6.3	M	M	
T4F4	Separation	6.4	M	M	
T4F5	Connection establishment	6.5	M	M	
T4F6	Connection refusal	6.6	M	M	
T4F7	Data TPDU numbering (normal)	6.10	M	M	
T4F8	Retention and acknowledgment of TPDU's Retention until acknowledgment of TPDU's (AK)	6.13.4.1	M	M	
T4F9	Explicit flow control	6.16	M	M	
T4F10	Checksum	6.17	M	M:1	
T4F11	Frozen references	6.18	M	M	
T4F12	Retransmission on time-out	6.19	M	M	
T4F13	Resequencing	6.20	M	M	
T4F14	Inactivity control	6.21	M	M	

Notes:

1. Checksum is **mandatory for CR TPDU only**.

The following functions are mandatory if class 4 is operated over CLNS.

Index	Function	Ref.	Status	Profile	Support
T4F23	Transmission over CLNS	6.1.2	M	M	
T4F24	Normal release when operating over CLNS (explicit)	6.7.2	M	M	
T4F25	Association of TPDUs with Transport connection when operating over CLNS	6.9.2	M	M	
T4F26	Expedited data transfer when operating over CLNS (Network normal)	6.11.2	M	M:2	
T4F27	Treatment of protocol errors when operating over CLNS	6.22.2	M	M	

Notes:

- The support of this function is mandatory in ISO/IEC 8073. T1.204 calls for support of the expedited data transfer service with non-use being the default negotiation. Therefore, in generating CR and CC TPDUs, the transport layer entity shall set bit 1 of the Additional Option Selection optional parameter to 0 (zero) to negotiate non-use of the transport expedited data transfer service.

The following functions are optional.

Index	Function	Ref.	Status	Profile	Support
T4F28	Data TPDU numbering (extended)	6.10	O	M:3	
T4F29	Non-use of checksum	6.17	O	M:4	
T4F30	Concatenation	6.4	O	i	
T4F31	Retention and acknowledgment of TPDUs - Use of selective acknowledgment	6.13.4.4	O	i	
T4F32	Retention and acknowledgment of TPDUs - Use of request acknowledgment	6.13.4.3	O	i	

Notes:

- EXTENDED must be supported, but NORMAL is the default setting.
- All TPDUs, except CR TPDU, shall negotiate “non-use” of the checksum. Initiators shall request and responders shall agree to “non-use” of the checksum.

C.7.5 Supported TPDUs

The following TPDUs and the parameters that constitute their fixed parts are mandatory if a corresponding predicate in the status column is true.

Index	TPDUs		Ref.	Status	Profile	Support
ST1	CR	supported on transmission	13.1	IR1:M	M	
ST2	CR	supported on receipt	13.1	IR2:M	M	
ST3	CC	supported on transmission	13.1	IR2:M	M	
ST4	CC	supported on receipt	13.1	IR1:M	M	
ST5	DR	supported on transmission	13.1	IR2:M	M	
ST6	DR	supported on receipt	13.1	IR1:M	M	
ST7	DC	supported on transmission	13.1	C1 or C2 or C3 or C4 or C4L:M	M	
ST8	DC	supported on receipt	13.1	C1 or C2 or C3 or C4 or C4L:M	M	
ST9	DT	supported on transmission	13.1	M	M	
ST10	DT	supported on receipt	13.1	M	M	
ST11	ED	supported on transmission	13.1	C1 or C2 or C3 or C4 or C4L:M	M	
ST12	ED	supported on receipt	13.1	C1 or C2 or C3 or C4 or C4L:M	M	
ST13	AK	supported on transmission	13.1	C1 or C2 or C3 or C4 or C4L:M	M	
ST14	AK	supported on receipt	13.1	C1 or C2 or C3 or C4 or C4L:M	M	
ST15	EA	supported on transmission	13.1	C1 or C2 or C3 or C4 or C4L:M	M	
ST16	EA	supported on receipt	13.1	C1 or C2 or C3 or C4 or C4L:M	M	
ST19	ER	supported on receipt	13.1	M	M	

Notes:

C1 Class 1

C2 Class 2

C3 Class 3

C4 Class 4 over CONS

State for which classes, if any, ER is supported on transmission

Index	Class	Ref.	Status	Profile	Support
SER4L	Class 4 over CLNS	6.22.2	O	M(5)	

Notes:

5. See special cases listed in the GR for handling certain errors in CR and CC TPDUs.

C.7.6 Supported Parameters of Issued TPDU

C.7.6.1 Parameter Values for CR TPDU (C4L::)

If the additional options selection parameter is issued in a CR TPDU it is mandatory that

Index		Ref.	Profile	Support
ICR1	Bits 8 and 7 shall be set to zero	13.3.4 g)	M	

C.7.6.2 Supported Parameters for Class 4 TPDU (C4L::)

The following parameters are optional if a CR TPDU is issued with preferred class 4.

Index	Supported parameter	Ref.	Status	Profile	Support
I4CR7	Called TSAP-ID	13.3.4 a)	O	M(6)	
I4CR8	Calling TSAP-ID	13.3.4 a)	O	M(6)	
I4CR9	TPDU size	13.3.4 b)	O	M	
I4CR10	Version number	13.3.4 d)	O	i	
I4CR11	Protection parameters	13.3.4 e)	O	i	
I4CR12	Additional option selection	13.3.4 g)	O	M(7)	
I4CR13	Throughput	13.3.4 k)	O	i	
I4CR14	Residual error rate	13.3.4 m)	O	i	
I4CR15	Priority	13.3.4 n)	O	i	
I4CR16	Transit delay	13.3.4 p)	O	i	
I4CR17	Acknowledge time	13.3.4 j)	O	i	
I4CR18	Preferred maximum TPDU size	13.3.4 c)	O	i(8)	
I4CR19	Inactivity time	13.3.4 r)	O	i(8)	

Notes:

- The value of the Called TSAP-ID and Calling TSAP-ID parameters shall be ASCII "TT" (i.e., "5454" hex) to indicate that the ISO Session Layer is being run over TP4.
- All TPDU, except CR TPDU, shall negotiate "non-use" of the checksum. Initiators shall request "non-use" of the checksum. See also Note 2.
- This is a new optional function introduced in the 1992 version of ISO/IEC 8073. To be compatible with ISO/IEC ISP 10608-1 (which is based on ISO/IEC 8073:1988/ Amd 3 1992), this new functionality is defined to be **out of scope** for this profile.

The following parameters are optional if a CC TPDU is issued in class 4.

Index	Supported parameters	Ref.	Status	Profile	Support
I4CC6	Called TSAP-ID	13.4.4	O	M(9)	
I4CC7	Calling TSAP-ID	13.4.4	O	M(9)	
I4CC8	TPDU size	13.4.4	O	M	
I4CC9	Protection parameters	13.4.4	O	i	
I4CC10	Additional option selection	13.4.4	O	M(10)	
I4CC11	Acknowledge time	13.4.4	O	i	
I4CC12	Throughput	13.4.4	O	i	
I4CC13	Residual error rate	13.4.4	O	i	
I4CC14	Priority	13.4.4	O	i	
I4CC15	Transit delay	13.4.4	O	i	
I4CC16	Preferred maximum TPDU size	13.4.4	O	i(11)	
I4CC17	Inactivity time	13.4.4	O	i(11)	

Notes:

9. See Note 6.

10. All TPDUs, except CR TPDU, shall negotiate “non-use” of the checksum.
 Responders shall agree to “non-use” of the checksum. See also Note 2.

11. See note 8.

The following parameter is optional if a DR TPDU is issued in class 4.

Index	Supported parameter	Ref.	Status	Profile	Support
I4DR4	Additional information	13.5.4 a)	O	i	

The following parameter is mandatory in a DT TPDU if request of acknowledgment has been selected.

Index	Supported parameter	Ref.	Status	Profile	Support
I4DT4	ROA	13.7.3 a)	O	i(12)	

Notes:

12. See note 7.

The following parameter is mandatory in an AK TPDU if issued in class 4.

Index	Supported parameter	Ref.	Status	Profile	Support
I4AK4	Flow control confirmation	13.9.4 c)	O	M	

If the implementation can reduce credit and does so in the manner outlined in 12.2.3.8.2 then subsequence number in AK TPDUs is mandatory. Otherwise complete item I4AK5.

Index	Supported parameter	Ref.	Status	Profile	Support
I4AK5	Subsequence number	13.9.4 b)	O	M	

The following parameter is optional in an AK TPDU if selective acknowledgment has been negotiated.

Index	Supported parameter	Ref.	Status	Profile	Support
I4AK6	Selective acknowledgment parameters	13.9.4 d)	O	i(13)	

Notes:

13. See note 8.

The following parameter is optional if a ER TPDU is issued in class 4.

Index	Supported parameter	Ref.	Status	Profile	Support
I4ER3	Invalid TPDU	13.12.4 a)	O	i	

C.7.7 Supported Parameters for Received TPDUs

Implementors should be aware that implementations shall be capable of receiving and processing all possible parameters for all possible TPDUs, dependent upon the class and optional functions implemented.

C.7.8 User Data in Issued TPDUs

A TS-user may issue data with a T-CONNECT request, T-CONNECT response or T-DISCONNECT request. Then it shall be possible to send user data as follows:

C.7.8.1 Class 4 (C4L::)

Index	User data	Ref.	Status	Profile	Support
D4ICR	User data of up to 32 octets in a CR with preferred class 4	13.3.5	O	X(14)	
D4ICC	User data of up to 32 octets in a CC	13.4.5	O	X(14)	
D4IDR	User data of up to 64 octets in a DR	13.5.5	O	i	

Notes:

14. Implementation of this function is optional in the base PICS, but its **use** is **prohibited** in the profile. No protocol implementations shall send user data in CR and CC TPDU. See discussion in GR-828-CORE.

C.7.9 User Data in Received TPDU

For classes 1 to 4, if it is possible to initiate a CR TPDU then it shall be possible to receive the following.

Index	User data	Ref.	Profile	Support
DRCC	32 octets of user data in a CC TPDU	13.4.5	M(15)	
DRDR	64 octets of user data in a DR TPDU	13.5.5	M	

Notes:

15. All protocol implementations shall be prepared to receive user data in CC TPDU, and all implementations may ignore user data, i.e., user data shall not cause a disconnect.

For classes 1 to 4, if it is possible to respond to a CR TPDU then it shall be possible to receive the following.

Index	User data	Ref.	Profile	Support
DRCR	32 octets of user data in a CR TPDU	13.3.5	M(16)	

Notes:

16. All protocol implementations shall be prepared to receive user data in CR TPDU, and all implementations may ignore user data, i.e., user data shall not cause a disconnect.

C.7.10 Negotiation

C.7.10.1 Class Negotiation - Initiator

What class(es) is (are) contained in the alternative class parameter if the preferred class is:

Index	Preferred class	Ref.	Allowed values	Profile values	Support
NAC5	Class 4 over CLNS	6.5.5 j)	None	None	

C.7.10.2 TPDU Size Negotiation

Index		Ref.	Status	Profile	Support
TS1	If maximum TPDU size is proposed in a CR TPDU then the initiator shall support all TPDU sizes from 128 octets to the maximum proposed.	14.6	M	M	
TS2	If the preferred maximum TPDU size parameter is used in a CR TPDU then the initiator shall support all TPDU sizes, except 0, that are multiples of 128 octets up to the preferred maximum proposed.	14.6 e)	I4CR18:M	i(17)	

Notes:

17. Implementation of this function is conditionally mandatory in the base PICS. To be compatible with ISO/IEC ISP 10608-1 (which is based on ISO/IEC 8073:1988/ Amd 3 1992), **use** of this functionality is defined to be **out of scope** for this profile.

Index	TPDU size	Ref.	Allowed values	Profile	Support
T4S1	What is the largest value of the maximum TPDU size parameter in a CR TPDU with preferred class 4?	14.6 e)	NOT I4CR18: One of 128, 256, 512, 1024, 2048, 4096, 8192 I4CR18: One of n128 with n = 1, 2, 3,...	1024, see note 18	
T4S2	What is the largest value of the maximum TPDU size parameter which may be sent in a CC TPDU when class 4 is selected?	14.6 e)	NOT I4CC16: One of 128, 256, 512, 1024, 2048, 4096, 8192 I4CC16: One of n128 with n = 1, 2, 3, ...	1024, see note 18	

Notes:

18. Note larger TPDU sizes (2048, 4096, and 8192) are optional.

C.7.10.3 Use of Extended Format

Index	Extended format	Ref.	Allowed values	Profile values	Support
NEF3	What formats can you propose in the CR TPDU in class 4?	6.5.5 n)	normal, extended	see note 19	
NEF6	What formats can you select in CC when extended has been proposed in CR class 4?	6.5.5 n)	normal, extended	see note 19	

Notes:

19. Extended format options shall be implemented. Non-use of extended format shall be negotiable. The responder shall honor the initiator's request whenever possible. Negotiation to other than what has been requested shall occur only under abnormal conditions: for example, severe congestion, as determined by the implementor. Initiators shall be prepared to operate in the mode confirmed by the responder. Normal is the default format.

C.7.10.4 Expedited Data Transport Service

Index		Ref.	Status	Profile	Support
TED1	Expedited data indication in CR and CC TPDU	6.5.5 r)	M	M	

C.7.10.5 Non-Use of Checksum (C4L AND T4F29::)

Index	Non-use checksum	Ref.	Allowed values	Profile values	Support
NUC1	What proposals can you make in the CR?	6.5.5 p)	non-use, use	non-use	
NUC2	What proposals can you make in CC when non-use of checksum has been proposed in CR?	6.5.5 p)	non-use, use	non-use	

C.7.10.6 Use of Selective Acknowledgment (See note 20)

Index	Selective acknowledgment	Ref.	Allowed values	Profile values	Support
USA1	Is use of selective acknowledgment proposed in CR TPDUs?	6.5.5 s)	Yes, No	i	
USA2	Is use of selective acknowledgment selected in a CC when it has been proposed in a CR?	6.5.5 s)	Yes, No	i	

Notes:

20. This is a new function in the 1992 base PICS. To be consistent with ISO/IEC ISP 10608-1 (which is based on ISO/IEC 8073:1988/Amd 3 1992), this functionality is defined to be **out of scope** for this profile.

C.7.10.7 Use of Request of Acknowledgment (See note 21)

Index	Request of acknowledgment	Ref.	Allowed values	Profile	Support
ROA1	Is use of request of acknowledgment proposed in CR TPDUs?	6.5.5 t)	Yes, No	i	
ROA2	Is use of request of acknowledgment selected in a CC when it has been proposed in a CR?	6.5.5 t)	Yes, No	i	

Notes:

21. See note 20.

C.7.11 Error Handling

C.7.11.1 Action on Receipt of a Protocol Error

Index	Item	Ref.	Allowed values	Profile	Support
PE4L	Class 4 over CLNS	6.22.2.3	C4L: ER, DR, Discard	ER, DR, Discard, see note 22	

Notes:

22. See note 5.

C.7.11.2 Actions on Receipt of an Invalid or Undefined Parameter in a CR TPDU

Index	Event	Ref.	Status	Profile	Support
RR1	A parameter not defined in ISO 8073 shall be ignored	13.2.3	M	M	
RR2	An invalid value in the alternative protocol class parameter shall be treated as a protocol error	13.2.3	M	M	
RR3	An invalid value in the class and option parameter shall be treated as a protocol error	13.2.3	M	M	
RR4	On receipt of the additional option selection parameter bits 8 to 5, and bits 4 to 1 if not meaningful for the proposed class shall be ignored.	13.3.4	M	M	
RR5	If non-use of explicit flow control is proposed and bit 1 of the additional option selection parameter equals 1, it shall be treated as a protocol error.	13.2.3	M	M	
RR6	On receipt of the class and option parameter bits 4 to 1 if not meaningful for the proposed class shall be ignored	13.3.3	M	M(23)	

Notes:

23. This entry is not contained in ISO/IEC 8073:1988, but is contained in ISO/IEC 8073:1992 and it is included in the ISO/IEC 8073: 1992 base PICS for clarity. There are no incompatibilities between the PICS due to this entry.

What action is supported on receipt of the following?

Index	Event	Ref.	Allowed actions	Profile	Support
RR7	A parameter defined in ISO/IEC 8073 (other than those covered above) and have an invalid value	13.2.3	ignore, protocol error	ignore	

C.7.11.3 Actions on Receipt of an Invalid or Undefined Parameter in a TPDU other than a CR TPDU

The following actions are mandatory.

Index	Event	Ref.	Status	Profile	Support
UI1	A parameter not defined in ISO/IEC 8073 shall be treated as a protocol error	13.2.3	M	M	
UI2	A parameter which has an invalid value as defined in ISO/IEC 8073 shall be treated as a protocol error	13.2.3	M	M	
UI3 (class 4 only)	A TPDU received with a checksum which does not satisfy the defined formula shall be discarded.	6.17.3	M	M	

C.7.12 Timers and Protocol Parameters

The following are mandatory if class 4 is supported.

Index		Ref.	Status	Profile/ Supported Range & Default	Support
TA1	T1	12.2.1	M	M/ 0.25 seconds to 64 seconds & default of 8 seconds	
TA2	N	12.2.1	M	M/ 2 to 15 & default of 2	
TA3	I _L	12.2.1	M	M/ 2 seconds to 512 seconds & default of 64 seconds	
TA4	W	12.2.1	M	M/ 1 second to 256 seconds & default of 16 seconds	
TA5	L	12.2.1	M	M/ 1 second to 256 seconds & default of 32 seconds	

Index		Ref.	Status	Profile	Support
OT9	Does IUT support optional timer <i>TS2</i> when operating in class 4?	6.22.2.3	O	i	

C.8 TARP Protocol Implementation Conformance Statement

The following tables serve two purposes. First, when the “Profile” column is removed, they serve as the PICS for the TARP protocol defined in [Section 8.7](#) of this GR. Second, when the “Profile” column is kept, they serve as the profile for the implementation requirements for the TARP protocol. The symbols used in the status column for this PICS are the same as those defined in [Section C.4.1](#) for the ISO CLNP PICS.

C.8.1 Major Function

Index	Functionality/Description	Ref.	Status	Profile	Support
MF1	Does the NE support TARP on the NE-NE interface according to the requirements in GR-253-CORE?	8.7	O	M:1	
MF2	Does the NE support the finding of the NET that matches a given TID?	8.7.4.1	M	M	
MF3	Does the NE support the finding of the TID that matches a given NET?	8.7.4.2	M	M	
MF4	Support of notification to other NEs of TID or protocol address changes	8.7.4.3	M	M	
MF5	Does the NE transmit TARP PDUs within ISO 8473 (CLNP) Data (DT) PDUs?	8.7.1	M	M	

1. If SONET NE supports TL1/OSI on the NE-NE interface (DCC or LAN)

C.8.2 Supported PDUs

The NE shall provide the function of a TARP processor that is capable of supporting:

Index	Functionality/Description	Ref.	Status	Profile	Support
SP1-s	TARP Type 1 PDUs on transmit	8.7.2.4, 8.7.5.1	M(2)	M(2)	
SP1-r	TARP Type 1 PDUs on receive	8.7.2.4, 8.7.5.1, 8.7.5.6	M(3)	M(3)	
SP2-s	TARP Type 2 PDUs on transmit	8.7.2.4, 8.7.5.1, 8.7.5.2	M(2)	M(2)	
SP2-r	TARP Type 2 PDUs on receive	8.7.2.4, 8.7.5.1, 8.7.5.6	M(3)	M(3)	
SP3-s	TARP Type 3 PDUs on transmit	8.7.2.4, 8.7.5.3	M	M	
SP3-r	TARP Type 3 PDUs on receive	8.7.2.4, 8.7.5.6	M	M	
SP4-s	TARP Type 4 PDUs on transmit	8.7.2.4, 8.7.5.4	M	M	
SP4-r	TARP Type 4 PDUs on receive	8.7.2.4, 8.7.5.6	M	M	
SP5-s	TARP Type 5 PDUs on transmit	8.7.2.4, 8.7.5.5	M	M	
SP5-r	TARP Type 5 PDUs on receive	8.7.2.4, 8.7.5.6	M	M	

2. The tar-tor field is optional.
3. The contents of the tar-tor field shall be ignored. The PDUs shall be processed correctly regardless of the presence or absence of the tar-tor field.

C.8.3 Protocol Specifications

C.8.3.1 TARP PDU CLNP Specifications

The following table defines the TARP PDU fields carried in the fixed part of the CLNP DT PDU.

Index	Functionality/Description	Ref.	Status	Profile	Support
FxPt1	Lifetime of the CLNP DT PDU set to a value of 25000 milliseconds	8.7.1	M	M	
FxPt2	Segmentation Permitted Flag set to a value of one	8.7.1	M	M	
FxPt3	Error Report Flag set to a value of zero	8.7.1	M	M	

C.8.3.2 TARP PDU Specifications

The following table defines the TARP PDU fields carried in the data part of the CLNP DT PDU.

Index	Functionality/Description	Ref.	Status	Profile	Support
Data1	TARP lifetime (tar-lif)	8.7.2.1	M	M	
Data2	TARP sequence number (tar-seq)	8.7.2.2	M	M	
Data3	Protocol address type (tar-pro)	8.7.2.3	M	M	
Data4	URC and TARP type code (tar-tcd)	8.7.2.4	M	M	
Data5	TID target length (tar-tln)	8.7.2.5	M	M	
Data6	TID originator length (tar-oln)	8.7.2.6	M	M	
Data7	Protocol address length (tar-pln)	8.7.2.7	M	M	
Data8	TID of target (tar-ttg)	8.7.2.8	M	M	
Data9	TID of originator (tar-tor)	8.7.2.9	M	M	
Data10	Protocol address of originator (tar-por)	8.7.2.10	M	M	
Data11	URC bit ignored upon receipt of TARP PDU	8.7.2.4	M	M	

C.8.3.3 Protocol Timer Specifications

Index	Timer	Ref.	Status	Profile/ Supported Range & default	Support
PTS1	Timer T1	8.7.4.1	M	M/ 0 to 3600 s & default of 15 s	
PTS2	Timer T2	8.7.4.1	M	M/ 0 to 3600 s & default of 25 s	
PTS3	Timer T3	8.7.4.1	M	M/ 0 to 3600 s & default of 40 s	
PTS4	Timer T4	8.7.4.1	M	M/ 0 to 3600 s & default of 20 s	

C.8.4 Major Capabilities

Index	Functionality/Description	Ref.	Status	Profile	Support
MC1	Does the NE have the capability of designating one NET as the NET to be used for the mapping purposes for all TARP-related TID/NET mappings?	8.7	O	M:4	
MC2	Does the NE return a TARP Type 3 PDU with the tar-oln field equal to zero on receipt of a TARP Type 5 PDU for an NET other than the "designated NET"?	8.7	MC1:M	M	
MC3	Support of a TARP Data Cache (TDC)	8.7.3	O	i	
MC4	Support the incrementing of the tar-seq field on the origination of a TARP PDU	8.7.5	M	M	
MC5	Generation of a TARP Type 4 PDU to ensure other NEs LDB's are reset to zero (i.e., broadcast Type 4 PDU with tar-seq=0 to all adjacencies)	8.7.5	M	M	
MC6	Support ES TARP PDU processing on receipt	8.7.5.6.1	O.1	M:5	
MC7	Support Level 1 IS TARP PDU processing on receipt	8.7.5.6.2, 8.7.5.8	O.1	M:6	
MC8	Support Level 2 IS TARP PDU processing on receipt	8.7.5.6.3, 8.7.5.8	O.1	M:7	
MC9	Support propagation procedures	8.7.5.8	M:8	M:8	
MC10	Support of a circular (first-in first-out)TARP Loop Detection Buffer (LDB)	8.7.5.7	M:8	M:8	
MC11	Support LDB Flush Timer	8.7.5.7	MC10:M	M:8	
MC12	Support of TARP Echo Function	8.7.7	O	i	
MC13	Support for LDB Entry Timer	8.7.5.7	M:8	M:8	

O.1 = One or more of these items must be supported.

Notes:

4. if an NE has multiple NETs
5. if an ES NE
6. if a Level 1 IS NE
7. if a Level 2 IS NE
8. if an IS NE

C.8.5 TARP Processor Management

Index	Functionality/Description	Ref.	Status	Profile	Support
PM1	Capable of selectively disabling the TARP propagation by link/adjacency	8.7.6	MC9:M	M	
PM2	Capable of provisioning tar-lif	8.7.6	M	M	
PM3	Capable of provisioning tar-pro	8.7.6	M	M	
PM4	Not capable of provisioning remaining TARP PDU fields	8.7.6	M	M	
PM5	Capable of provisioning TARP timers	8.7.6	M	M	
PM6	Capable of displaying TDC via the local WS at a minimum	8.7.6	MC3:M	MC3:M	
PM7	Capable of displaying LDB via the local WS at a minimum	8.7.6	MC10:M	MC10:M	
PM8	Capable of displaying TARP sequence number via the local WS at a minimum	8.7.6	M	M	
PM9	Capable of manually flushing TDC	8.7.6	MC3:M	MC3:M	
PM10	Capable of manually flushing LDB	8.7.6	MC10:M	MC10:M	
PM11	Capable of manually provisioning TDC	8.7.6	MC3:M	MC3:M	
PM12	Capable of manually provisioning LDB	8.7.6	MC10:M	MC10:M	
PM13	Capable of manually provisioning TARP sequence number	8.7.6	M	M	
PM14	Capable of disabling of all TARP functions	8.7.6	M	M	
PM15	Capable of disabling of TARP propagation function	8.7.6	MC9:M	MC9:M	
PM16	Capable of disabling of TARP origination functions	8.7.6	M	M	
PM17	Capable of disabling of the TDC	8.7.6	MC3:M	MC3:M	
PM18	Capable of manually generating TARP requests	8.7.6	M	M	
PM19	Capable of manually provisioning a TARP adjacency	8.7.8	O	i	
PM20	Capable of manually provisioning the LDB entry timer	8.7.5.7	M	M	
PM21	Capable of manually provisioning the LDB flush timer	8.7.5.7	M	M	

C.8.5.1 LDB Entry Timer Parameters

Index	Functionality/Description	Ref.	Status	Profile	Support
LET1	The LDB entry timer shall be settable within a range of 1 to 10 minutes	8.7.5.7	M	M	
LET2	The default value of the LDB entry timer shall be 5 minutes	8.7.5.7	M	M	

C.8.5.2 LDB Flush Timer Parameters

Index	Functionality/Description	Ref.	Status	Profile	Support
LFT1	The LDB flush timer shall be settable within a range of 0 to 1440 minutes	8.7.5.7	M	M	
LFT2	The default value of the LDB flush timer shall be 5 minutes	8.7.5.7	M	M	

C.8.5.3 Provisionable TARP PDU Fields

Index	PDU Field	Ref.	Status	Profile/ Supported Range & default	Support
PTF1	tar-lif	8.7.2.1, 8.7.6	M	M/ 0 to 65,535 hops & default of 100 hops	
PTF2	tar-pro	8.7.2.3, 8.7.6	M	M/ '00' to 'FF' (hex) & default of 'FE' (hex)	

Appendix D: SONET Operations Communications Upper Layers Protocol Profile

D.1 Introduction

This appendix provides an upper layers profile for Session Layer, Presentation Layer, and ACSE protocols to be used by SONET NEs across the SONET OS/NE and NE/NE interfaces. This profile is intended to be used together with the SONET Lower Layers Profile (see Appendix C) and with future SONET Application Service Element (ASE)-specific profiles. The specific applications that have been considered when defining the SONET Upper Layers Profile are Interactive Class (CMISE and TL1 ASEs) and File-oriented Class (FTAM ASE). While most of the SONET Upper Layers Profile is common across the above applications, there are some ASE-specific differences that have been noted.

D.2 Source Documents

When defining the SONET Upper Layers Profile, the source documents listed in the subsections were used.

D.2.1 Base Standards

- ISO/IEC 8650-1: *Information technology - Open Systems Interconnection - Connection-oriented protocol for the Association Control Service Element: Protocol specification*
- ISO/IEC 8823-1: *Information technology - Open Systems Interconnection - Connection oriented presentation protocol: Protocol specification*
- ISO/IEC 8327-1: *Information technology - Open Systems Interconnection - Connection oriented session protocol: Protocol specification*
- ISO/IEC 10040: *Information technology - Open Systems Interconnection - Systems management overview*

D.2.2 PICS Proforma

- ISO/IEC 8650-2: *Information technology - Open Systems Interconnection - Protocol specification for the Association Control Service Element: Protocol Implementation Conformance Statement (PICS) proforma*
- ISO/IEC 8823-2: *Information technology - Open Systems Interconnection - Connection oriented presentation protocol: Protocol Implementation Conformance Statement (PICS) proforma*
- ISO/IEC 8327-2: *Information technology - Open Systems Interconnection - Connection oriented Session protocol: Protocol Implementation Conformance Statement (PICS) proforma*

D.2.3 International Standardized Profiles

- ISO/IEC ISP 10607-1: *Information technology - International Standardized Profiles AFTnn - File Transfer, Access and Management - Part 1: Specification of ACSE, Presentation and Session Protocols for the use by FTAM*
- ISO/IEC ISP 11183-1: *Information technology - International Standardized Profiles AOM1n OSI Management- Management Communications - Part 1: Specification of ACSE, presentation and session protocols for the use by ROSE and CMISE*
- ISO/IEC ISP 11188-1: *Information technology - International Standardized Profile - Common upper layer requirements - Part 1: Basic connection oriented requirements*

D.2.4 Telcordia Requirements

- [Section 8](#) of this GR
- GR-828-CORE
- GR-1250-CORE

D.3 Goals of SONET Upper Layers Profile

The goals of the SONET Upper Layers Profile are as follows:

- Define an upper layers profile that is common to all SONET applications (whenever practical).
- Remain consistent with ISPs (when possible without violating the above goal).

These goals are consistent with the direction of ISO standards work as reflected in ISP 11188-1.

D.4 Structure of SONET Upper Layers Profile

The SONET Upper Layers Profile is contained in [Sections D.6](#) (ACSE), [D.7](#) (Presentation Layer), and [D.8](#) (Session Layer). Within each of these sections, the following material (specific to the given protocol layer) is contained:

- A subsection called “*Additions Beyond Existing ISP Requirements*” that points out the items for which conformance to ISPs 10607-1 or 11183-1 is not by itself sufficient to meet the SONET Upper Layers Profile
- A subsection with tables taken from the ISO PICS proforma (with their original table numbers as designated in the PICS proforma)
- A “profile” column added to the above tables indicating the constraints that are imposed by the SONET Upper Layers Profile.

Note that some conditional parts of the existing ISPs (whose conditions do not apply to SONET operations as defined by the Telcordia documents referenced in [Section D.2](#)) have not been included in the SONET Upper Layers Profile. For example, not all of the conditional FTAM abstract syntaxes given by ISP 10607-1 apply to SONET operations.

D.5 Notations Used in the SONET Upper Layers Profile

The notations used by the SONET Upper Layers Profile are essentially the same as those that are used in the ISO PICS Proforma documents. The following is taken directly from those documents (except for the “profile column” which is not present in the PICS Proforma).

D.5.1 Abbreviations

Sts - status column
Spt - support column
Sdr - sender
Rcv - receiver
Pfl - profile column

D.5.2 Status Column

This column indicates the level of support required for conformance to the ISO base standard (for the given protocol layer). The values are as follows:

- ‘m’ - mandatory support is required
- ‘o’ - optional support is permitted for conformance to the base standard. If implemented it must conform to the specifications and restrictions contained in the base standard. These restrictions may affect the optionality of other items
- ‘n/a’ - the item is not applicable
- ‘cn’ - the item is conditional (where *n* is the number which identifies the condition which is applicable)
- ‘o.n’ - the item is optional, but the optionality is qualified (where *n* is the number which identifies the qualification which is applicable).

D.5.3 Profile Column

This column indicates the level of support required for conformance to the SONET Upper Layers Profile (for the given protocol layer). The values are as follows:

- 'm' - mandatory support of the feature is required, however, it is not a requirement that the feature be used in all instances of communication, unless mandated by the base standard or stated otherwise in this profile
- 'm(n)'-mandatory support is required as clarified by Note n
- 'c(n)' -the item is conditional (where n is the Note number which identifies the condition which is applicable)
- 'i' - the item is out of scope for the profile, however, implementation of that item is not precluded; for the purposes of the this profile, the implementation or non-implementation of that item is ignored
- 'x' - the item is excluded (i.e., prohibited) for the profile.

D.5.4 Support Column

The 'Support' column can be completed by the supplier or implementor to indicate the level of implementation of each feature. The proforma has been designed such that the only entries required in the 'Support' column are:

- 'Y' - yes, the feature has been implemented
- 'N' - no, the feature has not been implemented
- '—' - not applicable.

D.5.5 PICS Numbers

Each line within the PICS proforma which requires implementation detail to be entered is numbered at the left-hand edge of the line. This numbering is included as a means of uniquely identifying all possible implementation details within the PICS proforma.

The means of referencing individual responses should be to specify the following sequence:

- A. a reference to the smallest subclause enclosing the relevant item
- B. a solidus character, '/'
- C. the reference number of the row in which the response appears
- D. if, and only if, more than one response occurs in the row identified by the reference number, then each possible entry is implicitly labeled a, b, c, etc., from left to right, and this letter is appended to the sequence.

D.6 SONET Upper Layers Profile: ACSE

D.6.1 Additions Beyond Existing ISP Requirements

The numbers used below refer to the PICS Proforma numbers used in [Section D.6.2](#).

A.5.1 Association establishment

This profile requires that both the Initiator and Responder roles always be supported. ISP 10607-1 and ISP 11183-1 do not always require support of both roles.

A.9 Supported APDU parameters

Application Entity Title

When used with the CMISE or FTAM ASEs, this profile requires support for sending the following parameters in the AARQ APDU: Calling AP title, Calling AE qualifier, Called AP title, and Called AE qualifier. Similarly, it requires (for the FTAM ASE only) sending the "Responding AP title" and the "Responding AE qualifier" parameters in the AARE APDU.

For ISP 10607-1, these parameters are also mandatory.

For ISP 11183-1, these parameters are optional.

A.10.1 AE title name form

When used with the CMISE or FTAM ASEs, this profile requires support for sending Form 1 AE title name forms. For the FTAM ASE, it also requires support for sending Form 2.

For ISP 10607-1, Form 2 is required and Form 1 is optional.

For ISP 11183-1, both forms are optional, i.e., use of AE title is optional.

D.6.2 Profile Tables

Note that the tables below use the PICS numbers of the ISO PICS Proforma. Only those tables from the PICS Proforma that apply to the SONET Upper Layers Profile are provided.

A.5 Supported roles

A.5.1 Association establishment

	Role	Sts	Pfl	Spt	Associated predicate
1	Initiator	o.01	m		A-CON/initiator
2	Responder	o.01	m		A-CON/responder

o.01: a conforming implementation shall support at least one of the roles.

A.5.2 Normal release

	Role	Sts	Pfl	Spt	Associated predicate
1	Requester	o.02	m		A-REL/requester
2	Acceptor	o.02	m		A-REL/acceptor

o.02: a conforming implementation shall support at least one of the roles.

A.6 Protocol mechanisms

	Protocol mechanism	Sts	Pfl	Spt	Associated predicate
1	Normal mode	o.03	m		
2	X.410-1984 mode	o.03	i		
3	Rules for extensibility	m	m		
4	Support operation of Session version 2	o	m		S-O-SESS-V2

o.03: either Normal mode or X.410-1984 mode or both shall be supported. If only X.410-1984 mode is supported, then the remainder of the proforma shall be ignored.

A.7 Functional units

	ACSE functional units	Sts	Pfl	Spt	Associated predicate
1	Kernel	m	m		
2	Authentication	o	i		A-FU(AU)

A.8 Supported APDUs

	APDU	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	A-associate-request (AARQ)	c01	m		c02	m	
2	A-associate-response (AARE)	c02	m		c01	m	
3	A-release-request (RLRQ)	c03	m		c04	m	
4	A-release-response (RLRE)	c04	m		c03	m	
5	A-abort (ABRT)	c05	m		c05	m	

c01:if [A-CON/initiator] then m else n/a
c02:if [A-CON/responder] then m else n/a
c03:if [A-REL/initiator] then m else n/a
c04:if [A-REL/responder] then m else n/a
c05:if [S-O-SESS-V2] then m else n/a

A.9 Supported APDU parameters

Note: Applications may place further constraints on the use of APDU "User information" parameters.

A.9.1 A-associate-request (AARQ)

	Parameter	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Protocol version	c06	m(2)		c02	m(2)	
2	Application context name	c01	m(6)		c02	m	
3	Calling AP title	c06	c(5)		c02	m(1)	
4	Calling AE qualifier	c06	c(5)		c02	m(1)	
5	Calling AP-invocation-identifier	c06	i		c02	m	
6	Calling AE-invocation-identifier	c06	i		c02	m	
7	Called AP title	c06	c(5)		c02	m(1)	
8	Called AE qualifier	c06	c(5)		c02	m(1)	
9	Called AP-invocation-identifier	c06	i		c02	m	
10	Called AE-invocation-identifier	c06	i		c02	m	
11	ACSE-requirements	c07	i		c08	m(3)	
12	Authentication-mechanism-name	c07	i		c08	m(3)	
13	Authentication-value	c07	i		c08	m(3)	
14	Implementation information	c06	i		c02	m	
15	User information	c06	c(5)		c02	m	

c01:if [A-CON/initiator] then m else n/a

c02:if [A-CON/responder] then m else n/a

c06:if [A-CON/initiator] then o else n/a

c07:if [A-CON/initiator and A-FU(AU)] then m else n/a

c08:if [A-CON/responder and A-FU(AU)] then m else n/a

(1)Both forms shall be static mandatory/dynamically optional for receiving (see Table A.10.1).

(2)The default value "version 1" is defined in the abstract syntax definition of ACSE APDUs in ISO/IES 8650. A sender may omit this parameter when this value is intended. A receiver shall interpret the omission of an explicit value as implying the default value.

(3)If the authentication FU is not supported, based on the extensibility rules, these tagged values shall be received and ignored. The "Authentication-mechanism-name" and "Authentication-value" fields shall only be present if the "ACSE-requirements" field includes the authentication FU. The "Authentication-mechanism-name" field shall be present if "Authentication-value" is of type ANY DEFINED BY.

(5)if CMISE or FTAM ASE then m else i

(6)Values for the application context name parameter are specified in the application layer profiles.

A.9.2 A-associate-response (AARE)

	Parameter	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Protocol version	c09	m(2)		c01	m(2)	
2	Application context name	c02	m(6)		c01	m	
3	Responding AP title	c09	c(7)		c01	m(1)	
4	Responding AE qualifier	c09	c(7)		c01	m(1)	
5	Responding AP-invocation-identifier	c09	i		c01	m	
6	Responding AE-invocation-identifier	c09	i		c01	m	
7	Result	c02	m		c01	m	
8	Result source diagnostic	c10	m		c11	m	
9	ACSE-requirements	c08	i		c07	m(3)	
10	Authentication-mechanism-name	c08	i		c07	m(3)	
11	Authentication-value	c08	i		c07	m(3)	
12	Implementation information	c09	i		c01	m	
13	User information	c09	c(5)		c01	m	

Note: notes from previous table (A.9.1) also apply to this table (A.9.2).

c09:if [A-CON/responder] then o else n/a

c10:if [A-CON/responder] then (if [A-FU(AU)] then m (with a value range of 11 to 14) else o (with a value range of 1 to 10)) else n/a

c11:if [A-CON/initiator] then (if [A-FU(AU)] then m (with a value range of 11 to 14) else o (with a value range of 1 to 10)) else n/a

(7)if FTAM ASE then m else i

A.9.3 A-release-request (RLRQ)

	Parameter	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Reason	c12	i		c04	m	
2	User information	c12	c(1)		c04	m	

c04:if [A-REL/acceptor] then m else n/a

c12:if [A-REL/requester] then o else n/a

(1)if FTAM ASE then m else i

A.9.4 A-release-response (RLRE)

	Parameter	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Reason	c13	i		c03	m	
2	User information	c13	c(1)		c03	m	

c03:if [A-REL/initiator] then m else n/a

c13:if [A-REL/acceptor] then o else n/a

(1)if FTAM ASE then m else i

A.9.5 A-abort (ABRT)

	Parameter	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Abort source	m	m		m	m	
2	Diagnostic	c6	i		c6	i	
3	User information	o	m		m	m	

c6:if [A-FU(AU)] then m else n/a

A.10 Supported parameter forms

A.10.1 AE title name form

	Syntax form	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Form 1 (Directory name)	o.04	m		m	m(1)	
2	Form 2 (Object identifier and integer)	o.04	c(2)		m	m(1)	

o.04: a conforming implementation shall support at least one of the forms.

(1)Both forms should be static mandatory/dynamically optional for receiving.

(2)if FTAM ASE then m else i

D.7 SONET Upper Layers Profile: Presentation Layer

D.7.1 Additions Beyond Existing ISP Requirements

The numbers used below refer to the PICS Proforma numbers used in Section D.7.2.

A.7.3 CPR PPDU

For this profile, implementations shall support sending the "Provider reason" parameter in the CPR PPDU.

For ISP 11183-1, the same requirement applies.

For ISP 10607-1, support of the "Provider reason" parameter is optional.

A.7.5 ARP PPDU

For this profile, implementations shall support sending the "Abort reason" parameter in the ARP PPDU. The "Event identifier" parameter shall be present if the "Abort reason" parameter is set with the value 2, 3, 4, 5 or 6.

For ISP 11183-1, the same requirement applies.

For ISP 10607-1, support of the "Abort reason" parameter in the ARP PPDU is optional and support of the "Event identifier" parameter is out of scope.

D.7.2 Profile Tables

Note that the tables below use the PICS numbers of the ISO PICS Proforma. Only those tables from the PICS Proforma that apply to the SONET Upper Layers Profile are provided.

A.5 Protocol mechanisms and functional units

A.5.1 Protocol mechanisms

	Protocol mechanism	Sts	Pfl	Spt	Associated predicate
1	X.410-1984 mode	o.01	i		P-MODE(X.410)
2	Normal	o.01	m		P-MODE(NORMAL)

o.01: either Normal mode or X.410 (1984) mode or both shall be supported.

A.5.2 Functional units

	Presentation functional units	Sts	Pfl	Spt	Associated predicate
1	Kernel	m	m		
2	Presentation Context Management	c00	i		P-FU(CM)
3	Presentation Context Restoration	c01	i		P-FU(CR)
	Pass through to Session functional units				
4	Negotiated Release	o	Note 1		S-FU(NR)
5	Half Duplex	o.02	Note 1		S-FU(HD)
6	Duplex	o.02	Note 1		S-FU(FD)
7	Expedited Data	o	Note 1		S-FU(EX)
8	Typed Data	o	Note 1		S-FU(TD)
9	Capability Data Exchange	c02	Note 1		S-FU(CD)
10	Minor Synchronize	o	Note 1		S-FU(SY)
11	Symmetric Synchronize	o	Note 1		S-FU(SS)
12	Major Synchronize	o	Note 1		S-FU(MA)
13	Resynchronize	o	Note 1		S-FU(RESYN)
14	Exceptions	c03	Note 1		S-FU(EXCEP)
15	Activity Management	o	Note 1		S-FU(ACT)

Note 1: See Section D.8 (Session Layer Profile).

o.02: pass through for at least one of the Session functional units Duplex and Half Duplex shall be supported.

c00:if [P-MODE(NORMAL)] then o else n/a

c01:if [P-FU(CM)] then o else n/a

c02:if [S-FU(ACT)] then o else n/a

c03:if [S-FU(HD)] then o else n/a

A.6 Elements of procedure related to the PICS

A.6.1 Kernel functional unit

A.6.1.1 Supported roles

A.6.1.1.1 Presentation connection

	Role	Sts	Pfl	Spt	Associated predicate
1	Initiator	o.03	m		P-CON/initiator
2	Responder	o.03	m		P-CON/responder

o.03: a conforming implementation shall support at least one of the roles.

A.6.1.1.2 Normal data

	Role	Sts	Pfl	Spt	Associated predicate
1	Requester	o.04	m		P-DATA/requester
2	Acceptor	o.04	m		P-DATA/acceptor

o.04: a conforming implementation shall support at least one of the roles.

A.6.1.1.3 Orderly release

	Role	Sts	Pfl	Spt	Associated predicate
1	Requester	o.05	m		P-REL/requester
2	Acceptor	o.05	m		P-REL/acceptor

o.05: a conforming implementation shall support at least one of the roles.

A.6.1.2 Supported PPDU's associated with the kernel services

	PPDU	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Connect presentation (CP)	c04	m(1)		c05	m(1)	
2	Connect presentation Accept(CPA)	c05	m		c04	m	
3	Connect presentation Reject (CPR)	c05	m		c04	m	
4	Abnormal release provider (ARP)	m	m		m	m	
5	Abnormal release user (ARU)	o	m		m	m	
6	Presentation data (TD)	c06	m		c07	m	

c04:if [P-CON/initiator] then m else n/a

c05:if [P-CON/responder] then m else n/a

c06:if [P-DATA/requester] then m else n/a

c07:if [A-DATA/acceptor] then m else n/a

(1)includes Cptype (see ISP 11183-1, Item B.3.1/2)

A.7 Supported PPDU parameters

Note: Applications may place constraints on the encoding of PPDU "User data parameters", e.g., see Section 8.3.7.5 for constraints specified for TL1.

A.7.1 Connect presentation (CP) PPDU

	Parameter	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Calling presentation selector	c10	i		c05	m	
2	Called presentation selector	c10	m		c05	m	
3	Mode selector	c04	m		c05	m	
4	Presentation context definition list	c10	m(1)		c05	m(1)	
5	Default context name	c10	i		c05	m	
6	Protocol version	c04	m(2)		c05	m(2)	
7	Presentation requirements	c10	i		c05	m	
8	User session requirements	c11	i		c05	m	
9	User data	c10	m		c05	m	

c04:if [P-CON/initiator] then m else n/a

c05:if [P-CON/responder] then m else n/a

c10:if [P-CON/initiator] then o else n/a

c11:if [P-CON/initiator and P-FU(CM)] then o else n/a

(1)A conforming implementation shall encode presentation context identifiers in the range 0 to 32,767. For selection of odd or even value, see ISO 8823, 6.2.2.7 and 6.5.2.1.

(2)The default value "version 1" is defined in the structure of SS-user data definition in clause 8.2 of ISO/IEC 8823. A sender may omit this parameter when this value is intended. A receiver shall interpret the omission of an explicit value as implying the default value.

A.7.2 Connect presentation accept (CPA) PPDU

	Parameter	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Responding presentation selector	c12	i		c04	m	
2	Mode selector	c05	m		c04	m	
3	Presentation context definition result list	c05	m(1)		c13	m(1)	
4	Protocol version	c05	m(2)		c04	m(2)	
5	Presentation requirements	c12	i		c14	m	
6	User session requirements	c12	i		c15	m	
7	User data	c12	m		c04	m	

c04:if [P-CON/initiator] then m else n/a

c05:if [P-CON/responder] then m else n/a

c12:if [P-CON/responder and P-FU(CM)] then o else n/a

c13:if [P-CON/initiator and A.7.1/4a] then m else n/a

c14:if [P-CON/initiator and A.7.1/7a] then m else n/a

c15:if [P-CON/initiator and A.7.1/8a] then m else n/a

(1) A conforming implementation shall encode presentation context identifiers in the range 0 to 32,767.

(2) The default value "version 1" is defined in the structure of SS-user data definition in clause 8.2 of ISO/IEC 8823. A sender may omit this parameter when this value is intended. A receiver shall interpret the omission of an explicit value as implying the default value.

A.7.3 Connect presentation reject (CPR) PPDU

	Parameter	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Responding presentation selector	c12	i		c04	m	
2	Presentation context definition result list	c12	m(1)		c13	m(1)	
3	Protocol version	c12	m(3)		c04	m(3)	
4	Default context result	c12	i		c16	m	
5	Provider reason	c12	m		c04	m	
6	User data	c12	m(2)		c04	m	

c04:if [P-CON/initiator] then m else n/a

c12:if [P-CON/responder and P-FU(CM)] then o else n/a

c16:if [P-CON/initiator and A.7.1/5a] then m else n/a

(1)The "Presentation context definition result list" is required if the "Provider reason" parameter is absent. If the "Provider reason" is present, then the "Presentation context definition result list" is optional.

(2)Is not present if the connection is rejected by the Presentation service provider.

(3)The default value "version 1" is defined in the structure of SS-user data definition in clause 8.2 of ISO/IEC 8823. A sender may omit this parameter when this value is intended. A receiver shall interpret the omission of an explicit value as implying the default value.

A.7.4 Abnormal release user (ARU) PPDU

	Parameter	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Presentation context identifier list	c17	m	m	m	m	m
2	User data	c18	m	m	m	m	m

c17:if [A.6.1.2/5a] then (if [P-FU(CM) and A.7.5/2a or A.7.2/4a or P-CON/responder] then m else o) else n/a

c18:if [A.6.1.2/5a] then o else n/a

A.7.5 Abnormal release provider (ARP) PPDU

	Parameter	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Abort reason	m	m		m	m	
2	Event identifier	o	m(1)		m	m(1)	

(1)Mandatory if "Provider reason" parameter is set with the following values: 2, 3, 4, 5, 6.

A.7.6 Presentation data (TD) PPDU

	Parameter	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	User data	c06	m		c07	m	

c06:if [P-DATA/requester] then m else n/a

c07:if [A-DATA/acceptor] then m else n/a

A.8 Support of syntaxes

A.8.1 Transfer syntaxes supported

	Type	Detail	Support		Reference to definition	Reference to restriction
			Pfl	Spt		
1	Object identifier	{joint-iso-ccitt asn1(1) basic-encoding(1)}	c(1)		ISO/IEC 8825	see note 3
2	Object identifier	{1 3 17 104 12 2 bellcoreSONETSyntax(1)}	c(2)		Section 8.3.7.5	

(1)if CMISE or FTAM ASE then m else n/a

(2)if TL1 ASE then m else n/a

(3)For further restrictions on the use of ASN.1 Basic Encoding, see ISO/IEC DISP 11188-1, Common Upper Layer Requirements.

A.8.2 Abstract syntaxes supported

	Type	Detail	Support		Reference to definition	Reference to restriction
			Pfl	Spt		
1	Object identifier	{joint-iso-ccitt association-control(2) abstract-syntax(1) apdus(0) version1(1)}	m		ISO 8650	
2	Object identifier	{joint-iso-ccitt ms(9) cmip(1) cmip-pci(1) abstractSyntax(4)}	c(1)		ISO 9596-1	
3	Object identifier	{joint-iso-ccitt ms(9) smo(0) negotiationAbstractSyntax(1) version1(1)}	c(1)		ISO 10040	
4	FTAM PCI	{iso standard 8571 abstract-syntax(2) ftam-pci(1)}	c(2)		ISP 10607-1	
5	FTAM unstructured text	{iso standard 8571 abstract-syntax(2) unstructured-text(3)}	c(2)		ISP 10607-1	
6	FTAM unstructured binary	{iso standard 8571 abstract-syntax(2) unstructured-binary(4)}	c(2)		ISP 10607-1	
7	Object identifier	{1 13 17 104 11 2 bellcoreSONETSyntax(1)}	c(3)		Section 8.3.7.5	

Note: ISP 10607-1 conditionally requires seven additional abstract syntaxes for the FTAM ASE (not listed here), however, their conditions are false for the SONET File Transfer application.

(1)if CMISE ASE then m else n/a

(2)if FTAM ASE then m else n/a

(3)if TL1 ASE then m else n/a

D.8 SONET Upper Layers Profile: Session Layer

D.8.1 Additions Beyond Existing ISP Requirements

The numbers used below refer to the PICS Proforma numbers used in Section D.8.2.

A.8.9 Abort (AB) SPDU

Reflect Parameter Values

For this profile, the "Reflect Parameter Values" parameter in the AB SPDU is conditionally present only if the "Transport Disconnect" value is "protocol error".

For ISP 11183-1, the same requirement applies.

For ISP 10607-1, the presence of "Reflect Parameter Values" parameter in the AB SPDU is optional.

D.8.2 Profile Tables

Note that the tables below use the PICS numbers of the ISO PICS Proforma. Only those tables from the PICS Proforma that apply to the SONET Upper Layers Profile are provided.

A.3 *ISO 8327 protocol versions implemented*

	Version	Sts	Pfl	Spt	Associated predicate
1	Version 1	o.1	i		S-V1
2	Version 2	o.1	m		

o.1:At least one version shall be implemented.

A.6 *Supported functional units and protocol mechanisms*

A.6.1 *Functional units*

	Functional Unit	Sts	Pfl	Spt	Associated predicate
1	Kernel	m	m		
2	Negotiated Release	o	i		S-FUN(NR)
3	Half Duplex	o.2	i		S-FU(HD)
4	Duplex	o.2	m		S-FU(FD)
5	Expedited Data	o	i		S-FU(EX)
6	Typed data	o	i		S-FU(TD)
7	Capability Data Exchange	c1	i		S-FU(CD)
8	Minor Synchronize	o	i		S-FU(SY)
9	Symmetric Synchronize	o	i		S-FU(SS)
10	Major Synchronize	o	i		S-FU(MA)
11	Resynchronize	o	i		S-FU(RESYN)
12	Exceptions	c2	i		S-FU(EXCEP)
13	Activity management	o	i		S-FU(ACT)

o.2:At least one of the functional units Duplex and Half Duplex shall be implemented.

c1:if [S-FU(ACT)] then o else n/a

c2:if [S-FU(HD)] then o else n/a

A.6.2 Protocol mechanisms

	Mechanism	Sts	Pfl	Spt	Associated predicate
1	Use of transport expedited data (Extended control Quality Of Service)	o	i		S-EXP/T
2	Reuse of transport connection	o	i		S-REUSE/T
3	Basic concatenation	m	m		
4	Extended concatenation (sending)	o	i		
5	Extended concatenation (receiving)	o	i		S-XCONC/RCV
6	Segmenting (sending)	o	i		S-SEG/SDR
7	Segmenting (receiving)	o	i		S-SEG/RCV
8	Max size of SS-user-data < or = 512	o	x		S-MAXSIZE/512
9	Max size of SS-user-data < or = 10240	o	see note 1		S-MAXSIZE/10240
10	Max size of SS-user-data < or = 9	o	x		S-MAXSIZE/9

(1) at least 10240 octets is the required Max size of SS-user-data. It is a Telcordia objective to be able to support a Max size of 65,535 octets (see Section 8.3.5).

A.7 Elements of procedure related to the PICS

A.7.1 Kernel functional unit

A.7.1.1 Supported roles for the Kernel functional unit services

A.7.1.1.1 Session Connection

Does the implementation support the Session Connection as:

	Role	Sts	Pfl	Spt	Associated predicate
1	Initiator	o.3	m		S-CON/initiator
2	Responder	o.3	m		S-CON/responder

o.3:a conforming implementation must support at least one of the above roles.

A.7.1.1.2 Orderly Release

Does the implementation support the Orderly Release as:

	Role	Sts	Pfl	Spt	Associated predicate
1	Requester	o.4	m		S-REL/requester
2	Acceptor	o.4	m		S-REL/acceptor

o.4:a conforming implementation must support at least one of the above roles.

A.7.1.1.3 Normal Data Transfer

Does the implementation support the Normal Data Transfer as:

	Role	Sts	Pfl	Spt	Associated predicate
1	Requester	o.5	m		S-DATA/requester
2	Acceptor	o.5	m		S-DATA/acceptor

o.5:a conforming implementation must support at least one of the above roles.

A.7.1.2 Support for the SPDUs associated with the Kernel services

	SPDU	Sending			Receiving			Associated predicate	
		Sts	Pfl	Spt	Sts	Pfl	Spt	Sending	Receiving
1	Connect (CN)	c3	m		c4	m			
2	Overflow Accept (OA)	c5	i		c6	i		S-OA/SDR	S-OA/RCV
3	Connect Data Overflow (CDO)	c6	i		c5	i		S-CDO/SDR	S-CDO/RCV
4	Accept (AC)	c4	m		c3	m			
5	Refuse (RF)	c4	m		c3	m			
6	Finish (FN)	c7	m		c8	m			
7	Disconnect (DN)	c8	m		c7	m			
8	Abort (AB)	m	m		m	m			
9	Abort Accept (AA)	o	i		o	i			
10	Data Transfer (DT)	c9	m		c10	m			
11	Prepare (PR)	c11	i		c11	i		S-PR/SDR	S-PR/RCV

c3:if [S-CON/initiator] then m else n/a

c4:if [S-CON/responder] then m else n/a

c5:if [S-V1 or (NOT S-CON/responder)] then n/a else if [NOT S-MAXSIZE/10240] then m else o

c6:if [S-V1 or (NOT S-CON/initiator)] then n/a else if [NOT S-MAXSIZE/10240] then m else o

c7:if [S-REL/requester] then m else n/a

c8:if [S-REL/acceptor] then m else n/a

c9:if [S-DATA/requester] then m else n/a

c10:if [S-DATA/acceptor] then m else n/a

c11:if [S-V1] then n/a else if [NOT S-MAXSIZE/9 and S-EXP/T] then m else o

A.7.1.3 Support for the SPDUs associated with Token Exchange

	SPDU	Sending			Receiving			Comment
		Sts	Pfl	Spt	Sts	Pfl	Spt	
1	Give Token (GT)	m	m		m	m		
2	Please Token (PT)	m	i (1)		m	m		

(1)According to the base standard basic concatenation rules, the conditions under which PT would be used are all out-of-scope for this profile.

Note: Tables A.7.2 through A.7.3 of the PICS Proforma are not applicable to this profile.

A.7.4 Duplex functional unit [Associated predicate: S-FU(FD)]

No additional SPDUs (this clause is present for completeness).

Note: Tables A.7.5 through A.7.13 of the PICS Proforma are not applicable to this profile

A.8 Supported SPDU-parameters

A.8.1 Connect (CN) SPDU

	PGI "Connection Identifier"	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Calling SS-user Reference	c46	i		c4	m	
2	Common Reference	c46	i		c4	m	
3	Additional Reference Information	c46	i		c4	m	

c4:if [S-CON/responder] then m else n/a

c46:if [S-CON/initiator] then o else n/a

A.8.1.2 Connect/Accept Item

A.8.1.2.1 Connect/Accept Item parameters

Important Remark: If presence of the PGI "Connect/accept Item" is supported (see clause A.8.1.2.2) then presence of Protocol Options and Version Number parameters must be supported.

	PGI "Connection/Accept Item"	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Protocol Options	c47	m		c48	m	
2	TSDU maximum size	c49	i		c50	m	
3	Version Number	c51	m		c52	m	
4	Initial Serial Number	c53	i		c54	m	
5	Token Setting Item	c46	i		c55	m	
6	Second Initial Serial Number	c56	i		c57	i	

c46:if [S-CON/initiator] then o else n/a
 c47:if [NOT S-CON/initiator] then n/a else if [S-XCONC/RCV] then m else o
 c48:if [NOT S-CON/responder] then n/a else if [S-XCONC/RCV] then m else o
 c49:if [NOT S-CON/initiator] then n/a else if [S-SEG/SDR or S-SEG/RCV] then m else o
 c50:if [NOT S-CON/responder] then n/a else if [S-SEG/SDR or S-SEG/RCV] then m else o
 c51:if [NOT S-CON/initiator] then n/a else if [NOT S-V1] then m else o
 c52:if [NOT S-CON/responder] then n/a else if [NOT S-V1] then m else o
 c53:if [NOT S-CON/initiator] then n/a else if [{S-FU(SY) or S-FU(MA) or S-FU(SS) or S-FU(RESYN)} and NOT S-FU(ACT)] then m else o
 c54:if [NOT S-CON/responder] then n/a else if [{S-FU(SY) or S-FU(MA) or S-FU(SS) or S-FU(RESYN)} and NOT S-FU(ACT)] then m else o
 c55:if [S-CON/responder] then o else n/a
 c56:if [NOT S-CON/initiator] then n/a else if [S-FU(SS) and NOT S-FU(ACT)] then m else o
 c57:if [NOT S-CON/responder] then n/a else if [S-FU(SS) and NOT S-FU(ACT)] then m else o

A.8.1.2.2 Presence of Connect/Accept Item

	Presence of "Connection/Accept Item"	Sts	Pfl	Spt
1	Sending	c58	m	
2	Receiving	c59	m	

c58:if [NOT S-CON/initiator] then n/a else if [A8.1.2.1/1a or A.8.1.2.1/2a or A.8.1.2.1/3a or A.8.1.2.1/4a or A.8.1.2.1/5a or A.8.1.2.1/6a] then m else o
 c59:if [NOT S-CON/responder] then n/a else if [A8.1.2.1/1b or A.8.1.2.1/2b or A.8.1.2.1/3b or A.8.1.2.1/4b or A.8.1.2.1/5b or A.8.1.2.1/6b] then m else o

A.8.1.3 Single Items

	Single Items	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Session User Requirements	c60	m		c61	m	
2	Calling Session Selector	c46	i		c4	m	
3	Called Session Selector	c3	m(1)		c4	m	
4	Data Overflow	c6	i		c5	i	
5	User Data	c3	m		c4	m	
6	Extended User Data	c62	m		c63	m	

c3:if [S-CON/initiator] then m else n/a

c4:if [S-CON/responder] then m else n/a

c5:if [S-V1 or (NOT S-CON/responder)] then n/a
else if [NOT S-MAXSIZE/10240] then m else o

c6:if [S-V1 or (NOT S-CON/initiator)] then n/a
else if [NOT S-MAXSIZE/10240] then m else o

c46:if [S-CON/initiator] then o else n/a

c60:if [NOT S-CON/initiator] then n/a
else if [S-FU(HD) and S-FU(SY) and S-FU(ACT) and S-FU(CD) and S-FU(EXCEP)] then o
else m

c61:if [NOT S-CON/responder] then n/a
else if [S-FU(HD) and S-FU(SY) and S-FU(ACT) and S-FU(CD) and S-FU(EXCEP)] then o
else m

c62:if [S-V1 or (NOT S-CON/initiator)] then n/a
else if [NOT S-MAXSIZE/512] then m else o

c63:if [S-V1 or (NOT S-CON/responder)] then n/a
else if [NOT S-MAXSIZE/512] then m else o

(1) The value of the Called Session Selector parameter shall be ASCII "SS" (i.e., "5353" hex) to indicate that the ISO Presentation Layer is being run over the ISO Session Layer.

Note: Tables A.8.2 through A.8.3 of the PICS Proforma are not applicable to this profile.

A.8.4 Accept (AC) SPDU

A.8.4.1 Connect Identifier

	PGI "Connection Identifier"	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Calling SS-user Reference	c55	i		c3	m	
2	Common Reference	c55	i		c3	m	
3	Additional Reference Information	c55	i		c3	m	

c3:if [S-CON/initiator] then m else n/a

c55:if [S-CON/responder] then o else n/a

A.8.4.2 Connect/Accept Item

A.8.4.2.1 Connect/Accept Item parameters

Important Remark: If presence of the PGI "Connect/accept Item" is supported (see clause A.8.4.2.2) then presence of Protocol Options and Version Number parameters must be supported.

	PGI "Connection/Accept Item"	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Protocol Options	c48	m		c47	m	
2	TSDU maximum size	c50	i		c49	m	
3	Version Number	c52	m		c51	m	
4	Initial Serial Number	c54	i		c53	m	
5	Token Setting Item	c55	i		c46	m	
6	Second Initial Serial Number	c57	i		c56	i	

A.8.4.2.2 Presence of Connect/Accept Item

	Presence of "Connection/Accept Item"	Sts	Pfl	Spt
1	Sending	c70	m	
2	Receiving	c71	m	

c70:if [NOT S-CON/responder] then n/a
 else if [A8.4.2.1/1a or A.8.4.2.1/2a or A.8.4.2.1/3a or A.8.4.2.1/4a or A.8.4.2.1/5a or A.8.4.2.1/6a]
 then m else o

c71:if [NOT S-CON/initiator] then m
 else if [A8.4.2.1/1b or A.8.4.2.1/2b or A.8.4.2.1/3b or A.8.4.2.1/4b or A.8.4.2.1/5b or A.8.4.2.1/6b]
 then m else o

A.8.4.3 Single Items

	Single Items	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Token Item	c55	i		c3	m	
2	Session User Requirements	c61	m		c60	m	
3	Enclosure Item	c72	i		c73	i	
4	Calling Session Selector	c4	i		c46	m	
5	Responding Session Selector	c55	i		c3	m	
6	User Data	c4	m		c3	m	

Note: The notes used on Table A.8.1.3 also apply to this table.

c55:if [S-CON/responder] then o else n/a
 c72:if [S-CON/responder and NOT S-V1] then m else n/a
 c73:if [S-CON/initiator and NOT S-V1] then m else n/a

A.8.5 Refuse (RF) SPDU

A.8.5.1 Connection Identifier

	PGI "Connection Identifier"	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Calling SS-user Reference	c55	i		c3	m	
2	Common Reference	c55	i		c3	m	
3	Additional Reference Information	c55	i		c3	m	

c3:if [S-CON/initiator] then m else n/a

c55:if [S-CON/responder] then o else n/a

A.8.5.2 Single Items

	Single Items	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Transport Disconnect	c74	i		c75	m	
2	Session User Requirements	c(1)	c(1)		c(1)	c(1)	
3	Version Number	c52	m		c51	m	
4	Enclosure Item	c72	i		c73	i	
5	Reason Code	c(2)	c(2)		c(2)	c(2)	

c51:if [NOT S-CON/initiator] then n/a else if [NOT S-V1] then m else o

c52:if [NOT S-CON/responder] then n/a else if [NOT S-V1] then m else o

c72:if [S-CON/responder and NOT S-V1] then m else n/a

c73:if [S-CON/initiator and NOT S-V1] then m else n/a

c74:if [NOT S-CON/responder] then n/a else if [S-REUSE/T] then m else o

c75:if [NOT S-CON/initiator] then n/a else if [S-REUSE/T] then m else o

(1) Shall only be present if the "Reason Code" value = 2.

(2) Mandatory if "Enclosure Item" parameter is present.

A.8.6 Finish (FN) SPDU

	Single Items	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Transport Disconnect	c76	i		c79	m	
2	Enclosure Item	c77	i		c80	i	
3	User Data	c78	m		c8	m	

c8:if [S-REL/acceptor] then m else n/a

c76:if [NOT S-REL/requester] then n/a else if [S-REUSE/T] then m else o

c77:if [S-REL/requester and NOT S-V1] then m else n/a

c78:if [S-REL/requester] then o else n/a

c79:if [NOT S-REL/acceptor] then n/a else if [S-REUSE/T] then m else o

c80:if [S-REL/acceptor and NOT S-V1] then m else n/a

A.8.7 Disconnect (DN) SPDU

	Single Items	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Enclosure Item	c80	i		c77	i	
2	User Data	c81	m		c7	m	

c7:if [S-REL/requester] then m else n/a
 c77:if [S-REL/requester and NOT S-V1] then m else n/a
 c80:if [S-REL/acceptor and NOT S-V1] then m else n/a
 c81:if [S-REL/acceptor] then o else n/a

A.8.8 Not Finish (NF) SPDU

Not applicable to this profile.

A.8.9 Abort (AB) SPDU

	Single Items	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Transport Disconnect	m	m		m	m	
2	Enclosure Item	c83	i		c83	i	
3	Reflect Parameter Values	c(1)	c(1)		o	m	
4	User Data	o	m		m	m	

c83:if [NOT S-V1] then m else n/a
 (1)Only sent if Transport disconnect = protocol error.

A.8.10 Abort Accept (AA) SPDU

No Parameter field.

A.8.11 Data Transfer (DT) SPDU

	Single Items	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Enclosure Item	c84	i		c85	i	
2	User Information Field	c9	m		c10	m	

c9:if [S-DATA/requester] then m else n/a
 c10:if [S-DATA/acceptor] then m else n/a
 c84:if [S-DATA/requester and S-SEG/SDR] then m else n/a
 c85:if [S-DATA/acceptor and S-SEG/RCV] then m else n/a

Note: Tables A.8.12 through A.8.15 of the PICS Proforma are not applicable to this profile

A.8.16 Give Tokens (GT) SPDU

	Single Items	Sending			Receiving		
		Sts	Pfl	Spt	Sts	Pfl	Spt
1	Token Item	c98	c(1)		c99	c(1)	
2	Enclosure Item	c83	i		c83	i	
3	User Data	c100	m		c83	m	

c83:if [NOT S-V1] then m else n/a

c98:if [S-FU(NR) or S-FU(HD) or S-FU(SY) or S-FU(MA) or S-FU(ACT)] then o
 else n/a

c99:if [S-FU(NR) or S-FU(HD) or S-FU(SY) or S-FU(MA) or S-FU(ACT)] then m
 else n/a

c100: if [NOT S-V1 and A.8.16/1a] then o else n/a

(1)if VT ASE then m else i

Note: Tables A.8.17 through A.8.36 of the PICS Proforma are not applicable to this profile

References

Telcordia Publications

Generic Requirements

- GR-20-CORE, *Generic Requirements for Optical Fiber and Optical Fiber Cable*.
- GR-57-CORE, *Functional Criteria for Digital Loop Carrier (DLC) Systems (a module of TSGR, FR-440)*.
- GR-63-CORE, *NEBS™ Requirements: Physical Protection (a module of LSSGR, FR-64, TSGR, FR-440, and NEBSFR, FR-2063)*.
- GR-78-CORE, *Generic Requirements for the Physical Design and Manufacture of Telecommunications Products and Equipment (a module of RQGR, FR-796 and NEBSFR, FR-2063)*.
- GR-199-CORE, *OTGR Section 12.2: Operations Application Messages – Memory Administration Messages (a module of OTGR, FR-439)*.
- GR-303-CORE, *Integrated Digital Loop Carrier System Generic Requirements, Objectives, and Interface (a module of TSGR, FR-440)*.
- GR-326-CORE, *Generic Requirements for Singlemode Optical Connectors and Jumper Assemblies*.
- GR-357-CORE, *Generic Requirements for Assuring the Reliability of Components Used in Telecommunications Equipment (a module of RQGR, FR-796)*.
- GR-409-CORE, *Generic Requirements for Premises Fiber Optic Cable*.
- GR-418-CORE, *Generic Reliability Assurance Requirements for Fiber Optic Transport Systems (a module of RQGR, FR-796)*.
- GR-436-CORE, *Digital Network Synchronization Plan*.
- GR-449-CORE, *Generic Requirements and Design Consideration for Fiber Distributing Frames*.
- GR-454-CORE, *Generic Requirements for Supplier Provided Documentation (a module of OTGR, FR-439, LSSGR, FR-64, and TSGR, FR-440)*.
- GR-468-CORE, *Generic Reliability Assurance Requirements for Optoelectronic Devices Used in Telecommunications Equipment (a module of RQGR, FR-796)*.
- GR-472-CORE, *OTGR Section 2.1: Network Element Configuration Management (a module of OTGR, FR-439)*.
- GR-474-CORE, *OTGR Section 4: Network Maintenance: Alarm and Control for Network Elements (a module of OTGR, FR-439)*.
- GR-487-CORE, *Generic Requirements for Electronic Equipment Cabinets*.

- GR-496-CORE, *SONET Add-Drop Multiplexer (SONET ADM) Generic Criteria (a module of TSGR, FR-440)*.
- GR-499-CORE, *Transport Systems Generic Requirements (TSGR): Common Requirements (a module of TSGR, FR-440)*.
- GR-765-CORE, *Generic Requirements for Single Fiber Single-Mode Optical Splices and Splicing Systems*.
- GR-815-CORE, *Generic Requirements for Network Element/Network System (NE/NS) Security (a module of LSSGR, FR-64, and OTGR, FR-439)*.
- GR-818-CORE, *Network Maintenance: Access and Testing – Generic Test Architecture (a module of OTGR, FR-439)*.
- GR-819-CORE, *Network Maintenance: Access and Testing – Special Services (SS) and SS-Like Networks (a module of OTGR, FR-439)*.
- GR-820-CORE, *OTGR Section 5.1: Generic Digital Transmission Surveillance (a module of OTGR, FR-439)*.
- GR-826-CORE, *OTGR Section 10.2: User Interface Generic Requirements for Supporting Network Element Operations (a module of OTGR, FR-439)*.
- GR-828-CORE, *OTGR Section 11.2: Generic Operations Interface – OSI Communications Architecture (a module of OTGR, FR-439)*.
- GR-831-CORE, *OTGR Section 12.1: Operations Application Messages – Language for Operations Application Messages (a module of OTGR, FR-439)*.
- GR-833-CORE, *OTGR Section 12.3: Network Maintenance: Network Element and Transport Surveillance Messages (a module of OTGR, FR-439)*.
- GR-834-CORE, *OTGR Section 12.4: Network Maintenance: Access and Testing Messages (a module of OTGR, FR-439)*.
- GR-836-CORE, *OTGR Section 15.2: Generic Operations Interfaces Using OSI Tools – Information Model Overview: Transport Configuration and Surveillance for Network Elements (a module of OTGR, FR-439)*.
- GR-836-IMD, *OTGR Section 15.2: Generic Operations Interfaces Using OSI Tools – Information Model Details: Transport Configuration and Surveillance for Network Elements (a module of OTGR, FR-439)*.
- GR-839-CORE, *Generic Requirements for Supplier-Provided Training (a module of LSSGR, FR-64, TSGR, FR-440, and OTGR, FR-439)*.
- GR-910-CORE, *Generic Requirements for Fiber Optic Attenuators*.
- GR-1031-CORE, *OTGR Section 15.6: Operations Interfaces Using OSI Tools: Test Access Management*.
- GR-1042-CORE, *Generic Requirements for Operations Interfaces Using OSI Tools – Information Model Overview: Synchronous Optical Network (SONET) Transport Information Model*.
- GR-1042-IMD, *Generic Requirements for Operations Interfaces Using OSI Tools – Information Model Details: Synchronous Optical Network (SONET) Transport Information Model*.

- GR-1089-CORE, *Electromagnetic Compatibility and Electrical Safety - Generic Criteria for Network Telecommunications Equipment (a module of LSSGR, FR-64, TSGR, FR-440 and NEBSFR, FR-2063).*
- GR-1093-CORE, *Generic State Requirements for Network Elements.*
- GR-1209-CORE, *Generic Requirements for Passive Optical Components.*
- GR-1221-CORE, *Generic Reliability Assurance Requirements for Passive Optical Components (a module of RQGR, FR-796).*
- GR-1230-CORE, *SONET Bidirectional Line-Switched Ring Equipment Generic Criteria (a module of TSGR, FR-440).*
- GR-1244-CORE, *Clocks for the Synchronized Network: Common Generic Criteria.*
- GR-1250-CORE, *Generic Requirements for Synchronous Optical Network (SONET) File Transfer.*
- GR-1253-CORE, *Generic Requirements for Operations Interfaces Using OSI Tools: Telecommunications Management Network Security Administration.*
- GR-1309-CORE, *TSC/RTU and OTAU Generic Requirements for Remote Optical Fiber Testing (a module of OTGR, FR-439).*
- GR-1312-CORE, *Generic Requirements for Optical Fiber Amplifiers and Proprietary Dense Wavelength-Division Multiplexed Systems.*
- GR-1332-CORE, *Generic Requirements for Data Communications Network Security.*
- GR-1377-CORE, *SONET OC-192 Transport Systems Generic Criteria (a module of TSGR, FR-440) – replaced by GR-253-CORE, Issue 3, and therefore no longer available.*
- GR-1400-CORE, *SONET Dual-Fed Unidirectional Path Switched Ring (UPSR) Equipment Generic Criteria (a module of TSGR, FR-440).*
- GR-1469-CORE, *Security for OSI-Based Telecommunications Management Network Interfaces.*
- GR-2854-CORE, *Generic Requirements for Fiber Optic Dispersion Compensators.*
- GR-2915-CORE, *Application of Software Management OSI Information Model to Software Download and Memory Restoration.*
- GR-2947-CORE, *Generic Requirements for Portable Polarization Mode Dispersion (PMD) Test Sets.*
- GR-3000-CORE, *Generic Requirements for SONET Element Management Systems (EMSs).*

Technical References

- TR-NWT-000170, *Digital Cross-Connect System (DSC 1/0) Generic Criteria.*

- TR-TSY-000824, *OTGR Section 10.1: User System Interface - User System Access (a module of OTGR, FR-439)*.
- TR-TSY-000825, *OTGR Section 10.A: User System Interface - User System Language (a module of OTGR, FR-439)*.
- TR-TSY-000827, *OTGR Section 11.1: Generic Operations Interfaces: Non-OSI Communications Architecture (a module of OTGR, FR-439)*.
- TR-NWT-000835, *OTGR Section 12.5: Network Element and Network System Security Administration Messages (a module of OTGR, FR-439)*.
- TR-NWT-000917, *SONET Regenerator (SONET RGTR) Equipment Generic Criteria (a module of TSGR, FR-440)*.
- TR-NWT-000930, *Generic Requirements for Hybrid Microcircuits Used in Telecommunications Equipment (a module of RQGR, FR-796)*.
- TR-NWT-001112, *Broadband-ISDN User to Network Interface and Network Node Interface Physical Layer Generic Criteria*.

Special Reports

- SR-NWT-002723, *Applicable TL1 Messages for SONET Network Elements*.

Note

All Telcordia documents are subject to change, and their citations in this document reflect the most current information available at the time of this printing. Readers are advised to check current status and availability of all documents.

To Contact Telcordia Customer Service or to Order Documents

Telcordia Customer Service
8 Corporate Place, Room 3B-176
Piscataway, NJ 08854-4156
1.866.672.6997 (USA)
+ 1.732.699.6700 (Worldwide)
+ 1.732.336.2226 (FAX)
e-mail: document-info@telcordia.com
<http://telecom-info.telcordia.com>

To Order Documents Online From the Telcordia Information SuperStore

1. **Go to the SuperStore Web site:** <http://telecom-info.telcordia.com>
(At this site, the **Search** and **Browse** selections provide access to the Telcordia catalog of technical information.)
2. If you know the document number, enter it in the **Product Search** box in the left margin and click **Go!**
3. Click on the desired product match.

If you do not know the document number, do the following:

1. Click on **Search** located on the top bar.
2. In the **Keywords field**, enter the keywords (or document number), then click **Submit Search**.
3. Click on the desired product match.

To Order Documents From Within Telcordia (Employees Only)

1. Access the Telcordia Internal Home Page (InSite).
2. Click on **Self-Service** on the Navigation Bar.
3. Click on **Document Database**.
4. Click on **Basic Search** in the bullet list to obtain the Basic Search Criteria box.
5. In the Search by Document Number field, enter the document number (e.g., GR-454), then scroll down to click on **Submit Search**.
6. In the Basic Search Navigation List, select **Click for Abstract** to order an available document, or select **Click for Document** to view an available document.

For New Information on Telcordia Technologies Technical Products

- **Digest** (Free Download or E-mail Distribution List)

Telcordia offers an electronic-only monthly publication — the *Digest of Technical Information*, which provides the latest information on GRs and other products. At the *Digest* Web site, you may view recent document releases and download the monthly *Digest* PDF files at no charge. For more information visit:

<http://www.telcordia.com/digest>

- **Digest Plus** (Free Subscription)

Telcordia also offers a printed publication — the *Digest Plus*, which includes a quarterly compilation of new document releases and other items that appeared in the monthly issues of the *Digest*. For more information visit:

<http://telecom-info.telcordia.com> (Click on Digest on the top.)

EIA/TIA Documents

- TIA-455-127, *FOTP-127, Spectral Characterization of Multimode Laser Diodes.*
- TIA-4920000B, *Generic Specification for Optical Waveguide Fibers.*
- TIA-526.2, *OFSTP-2, Effective Transmitter Output Power Coupled into Single-Mode Fiber Optic Cable.*
- TIA-526.3, *OFSTP-3, Fiber Optic Terminal Receiver Sensitivity and Maximum Receiver Input Power.*
- TIA-526.10, *OFSTP-10, Measurement of Dispersion Power Penalty in Single-Mode Systems.*
- TIA-559, *Single-Mode Fiber Optic System Transmission Design.*

These publications are available at www.tiaonline.org.

American National Standards Institute (ANSI) Documents

- T1.101-1999, *Synchronization Interface Standard*
- T1.102-1993 (R2005), *Digital Hierarchy – Electrical Interfaces*
- T1.105-2001, *Synchronous Optical Network (SONET) – Basic Description including Multiplex Structure, Rates and Formats*
- T1.105.01-2000 (R2005), *Synchronous Optical Network (SONET) – Automatic Protection Switching*
- T1.105.02-2001, *Synchronous Optical Network (SONET) – Payload Mappings*
- T1.105.03-2003, *Synchronous Optical Network (SONET) – Jitter and Wander at Network and Equipment Interfaces*
- T1.105.04-1995 (R2001), *Synchronous Optical Network (SONET): Data Communication Channel Protocols and Architectures*
- T1.105.05-2002, *Synchronous Optical Network (SONET): Tandem Connection Maintenance*
- T1.105.06-2002, *Synchronous Optical Network (SONET): Physical Layer Specifications*
- T1.105.08-2001, *Synchronous Optical Network (SONET) – In-band Forward Error Correction Code Specification*
- T1.105.09-1996 (R2002), *Synchronous Optical Network (SONET): Network Element Timing and Synchronization*
- T1.107-2002, *Digital Hierarchy – Formats Specifications*
- T1.204-1997 (R2003), *OAM&P – Lower Layer Protocols for Telecommunication Management Network (TMN) Interfaces, Q3 and X Interfaces*
- T1.231-2003, *Layer 1 In-Service Digital Transmission Performance Monitoring*

- T1.231.04-2003, *SONET Layer 1 In-Service Digital Transmission Performance Monitoring*
- T1.245-1997 (R2003), *Directory Service for Telecommunications Management Network (TMN) and Synchronous Optical Network (SONET)*
- T1.269-2000, *Information Interchange – Structure and Representation of Trace Message Formats for the North American Telecommunications System*
- T1.403-1999, *Network and Customer Installation Interfaces – DS1 Electrical Interface*
- T1.506-1997 (R2001), *Network Performance – Switched Exchange Access Network Transmission Specifications*
- T1.508-2003, *Loss Plan for Evolving Digital Networks*
- T1.646-2003, *Network and Customer Installation Interfaces – Broadband ISDN: Common Criteria*

These publications are available at www.atis.org/atis/docstore.

U.S. Government Publications

U.S. Department of Health, Education, and Welfare; Bureau of Radiological Health.

- 21 CFR 1040.10, *Performance Standard for Light-Emitting Products – Laser Products*.

This publication is available at www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfcfr/CFRSearch.cfm.

ITU-T Recommendations

- E.164, *The international public telecommunication numbering plan*, February 2005
- G.651, *Characteristics of a 50/125 μm multimode graded index optical fibre cable*, February 1998
- G.655, *Characteristics of a non-zero dispersion-shifted single-mode optical fibre and cable*, March 2003
- G.691, *Optical interfaces for single channel STM-64 and other SDH systems with optical amplifiers*, December 2003
- G.693, *Optical interfaces for intra-office systems*, January 2005
- G.707, *Network node interfaces for the synchronous digital hierarchy (SDH)*, December 2003
- G.783, *Characteristics of synchronous digital hierarchy (SDH) equipment functional blocks*, February 2004
- G.812, *Timing requirements of slave clocks suitable for use as node clocks in synchronization networks*, June 2004
- G.813, *Timing characteristics of SDH equipment slave clocks (SEC)*, March 2003

- G.825, *The control of jitter and wander within digital networks which are based on the synchronous digital hierarchy (SDH)*, March 2000
- G.831, *Management capabilities of transport networks based on the synchronous digital hierarchy (SDH)*, March 2000
- G.841, *Types and characteristics of SDH network protection architectures*, October 1998
- G.957, *Optical interfaces for equipments and systems relating to the synchronous digital hierarchy*, July 1999
- G.958, *Digital line systems based on the synchronous digital hierarchy for use on optical fiber cables*, November 1994 (withdrawn)
- G.959.1, *Optical transport network physical layer interfaces*, December 2003
- G.7041, *Generic framing procedure (GFP)*, August 2005
- G.7042, *Link capacity adjustment scheme (LCAS) for virtual concatenated signals*, February 2004
- G.8251, *The control of jitter and wander within the optical transport network (OTN)*, November 2001
- M.3010, *Principles for a telecommunications management network*, February 2000
- M.3100, *Generic network information model*, April 2005
- Rec. O.181, *Equipment to assess error performance on STM-N interfaces*, May 2002
- Q.811, *Lower Layer protocol profiles for the Q and X interfaces*, February 2004
- Q.812, *Upper layer protocol profiles for the Q and X interfaces*, February 2004
- Q.921, *ISDN user-network interface - Data link layer specification*, September 1997
- T.50, *International Reference Alphabet (IRA) (Formerly International Alphabet No. 5 or IA5) – Information technology – 7-bit coded character set for information interchange*, September 1992
- X.121, *International numbering plan for public data networks*, October 2000

These publications are available at www.itu.int/rec/recommendation.asp.

ISO Documents

- ISO 3166-1:1997, *Codes for the representation of names of countries and their subdivisions – Part 1: Country codes*
- ISO/IEC 7498-1, *Information technology – Open System Interconnection – Basic Reference Model: The Basic Model*
- ISO/IEC 8073: *Information technology – Open Systems Interconnection – Protocol for providing the connection-mode transport service*

- ISO/IEC 8208: *Information technology – Data communications – X.25 Packet Layer Protocol for Data Terminal Equipment*
- ISO/IEC 8327-1: *Information technology – Open Systems Interconnection – Connection oriented session protocol: Protocol specification*
- ISO/IEC 8327-2: *Information technology – Open Systems Interconnection – Connection oriented session protocol: Protocol Implementation Conformance Statement (PICS) proforma*
- ISO/IEC 8348, *Information technology – Open Systems Interconnection – Network service definition*
- ISO/IEC 8473-1: *Information technology – Protocol for providing the connectionless-mode network service: Protocol specification*
- ISO/IEC 8473-2: *Information technology – Protocol for providing the connectionless-mode network service: Provision of the underlying service by an ISO/IEC 8802 subnetwork*
- ISO/IEC 8473-3: *Information technology – Protocol for providing the connectionless-mode network service: Provision of the underlying service by an X.25 subnetwork*
- ISO/IEC 8473-4: *Information technology – Protocol for providing the connectionless-mode network service: Provision of the underlying service by a subnetwork that provides the OSI data link service*
- ISO/IEC 8650-1: *Information technology – Open Systems Interconnection – Connection-oriented protocol for the Association Control Service Element: Protocol specification*
- ISO/IEC 8650-2: *Information technology – Open Systems Interconnection – Protocol specification for the Association Control Service Element: Protocol Implementation Conformance Statement (PICS) proforma*
- ISO/IEC 8802-2, *Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements – Part 2: Logical link control*
- ISO/IEC 8802-3, *Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements – Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications*
- ISO/IEC 8823-1: *Information technology – Open Systems Interconnection – Connection oriented presentation protocol: Protocol specification*
- ISO/IEC 8823-2: *Information technology – Open Systems Interconnection – Connection oriented presentation protocol: Protocol Implementation Conformance Statement (PICS) proforma*
- ISO 9542:1988: *Information processing systems – Telecommunications and information exchange between systems – End system to Intermediate system routing exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)*

- ISO/IEC TR 9577: *Information technology – Protocol identification in the network layer*
- ISO/IEC 10040: *Information technology – Open Systems Interconnection – Systems management overview*
- ISO/IEC TR 10172:1991: *Information technology – Telecommunications and information exchange between systems – Network/Transport Protocol interworking specification*
- ISO/IEC 10589:1992: *Information technology – Telecommunications and information exchange between systems – Intermediate system to Intermediate system intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)*
- ISO/IEC ISP 10607-1: *Information technology – International Standardized Profiles AFTnn – File Transfer, Access and Management – Part 1: Specification of ACSE, Presentation and Session Protocols for the use by FTAM*
- ISO/IEC ISP 10608-1:1992: *Information technology – International Standardized Profile TAnnn – Connection-mode Transport Service over Connectionless-mode Network Service – Part 1: General overview and subnetwork-independent requirements*
- ISO/IEC ISP 10608-2:1992: *Information technology – International Standardized Profile TAnnn – Connection-mode Transport Service over Connectionless-mode Network Service – Part 2: TA51 profile including subnetwork-dependent requirements for CSMA/CD Local Area Networks (LANs)*
- ISO/IEC 10747:1994, *Information Technology – Telecommunications and information exchange between systems – Protocol for exchange of inter-domain routing information among intermediate systems to support forwarding of ISO 8473 PDUs*
- ISO/IEC ISP 11183-1: *Information technology – International Standardized Profiles AOMIn OSI Management – Management Communications – Part 1: Specification of ACSE, presentation and session protocols for the use by ROSE and CMISE*
- ISO/IEC ISP 11188-1: *Information technology – International Standardized Profile – Common upper layer requirements – Part 1: Basic connection oriented requirements*

These publications are available at <http://www.iso.org/iso/en/prods-services/ISOstore/store.html>.

IEEE Documents

- IEEE 802.3ae, *Standard for Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications - Media Access Control (MAC) Parameters, Physical Layer and Management Parameters for 10 Gb/s Operation, 2002*

- IEEE 802.6, *Distributed Queue Dual Bus (DQDB) Subnetwork of a Metropolitan Area Network (MAN)*, September 1990.

IEEE publications are available at <http://shop.ieee.org/ieeestore>.

SIF/NSIF Documents

- SIF-009-1997, *NE-NE Remote Login Implementation Requirements Specification*
- NSIF-033-1999, *Requirements for the TCP/IP Protocol Suite on the SONET Access DCN*
- NSIF-038-2000, *NSIF Requirements for a Centralized Security Server*

NSIF (formerly SIF) is an Alliance for Telecommunications Industry Solutions (ATIS) sponsored committee. SIF/NSIF-approved documents are available at www.atis.org/atis/sif/sifdoc.htm.

Internet Standards

- RFC 793, *Transmission Control Protocol*
- RFC 2126, *ISO Transport Service on top of TCP (ITOT)*

These documents are available from various sources, including www.ietf.org.

Glossary

This section includes definitions for a number of terms that are not defined elsewhere in this document, along with a list of acronyms. Note that change bars are not used in this section to indicate where definitions or acronyms that appeared in Issue 3 of this document have been removed, or to indicate where new acronyms have been added. Instead, they are used only to indicate where a definition has been added or (in the opinion of Telcordia) significantly revised.

Definitions

Add-Drop Multiplexer (ADM) – An NE that provides access to some or all of the STS and/or VT paths contained within the OC-N optical signals at one or two “high-speed” OC-N interfaces. If two high-speed OC-N interfaces are provided, the path signals are added (inserted) to and/or dropped (extracted) from OC-N signals as they pass through the ADM. [Note that in addition to its high-speed OC-N interfaces, an ADM may support multiple OC-M ($M \leq N$) and/or STS-N electrical “tributary” interfaces (in addition to any non-SONET tributary interfaces).]

Asynchronous Transfer Mode (ATM) – A multiplexing/switching technique in which information is organized into fixed-length cells with each cell consisting of an identification header field and an information field. The transfer mode is asynchronous in the sense that the use of the cells depends on the required or instantaneous bit rate.

Bit Interleaved Parity N (BIP-N) – A method of error monitoring. If “even parity” is used, the transmitting equipment generates an N-bit code over a specified portion of the signal in such a manner that the first bit of the code provides even parity over the first bit of all N-bit sequences in the covered portion of the signal, the second bit provides even parity over the second bits of all N-bit sequences within the specified portion, etc. Even parity is generated by setting the BIP-N bits so that there are an even number of ones in each of all N-bit sequences, including the BIP-N.

Digital Cross-Connect System (DCS) – An NE that terminates standard digital signals and facilities operating at a standard digital signal rate, and automatically cross-connects constituent (tributary) signals according to an electronically alterable memory map. In contrast to an ADM, a SONET DCS provides access to the STS and/or VT paths contained in multiple (rather than one or two) high-speed OC-N interfaces.

Distributed Queue Dual Bus (DQDB) – A multiplexing/switching technique similar to ATM that uses IEEE 802.6 PLCP.

Eye Diagram – A graphic presentation formed by the superimposition of the waveforms of all possible pulse sequences.

Fixed Stuff (R-Bits/Bytes) – Fixed stuff (R) bits and bytes are used to compensate for the differences between the bandwidth available in STS and VT SPEs and the bandwidth required for the actual payload mappings (e.g., DS1, DS1C, DS2, and DS3). R-bits and bytes have no defined value. The receiver is required to ignore the value of these bits/bytes (except for BIP-8 calculation/verification).

Gateway Communications Functions – Functions to facilitate operations communications between two communicating entities across dissimilar subnetworks. Examples include concentration, message routing and relaying (Network Layer), and application layer protocol conversion and/or message translation. See Mediation Functions.

Information Processing Functions – Functions centralized in a subnetwork to perform common application processing and management for NEs within the subnetwork. Example NE functions include performance data storage and failure condition threshold crossing detection for alarm reporting. Examples of management functions include the migration of “OS-like” functions out to the network for subnetwork trouble sectionalization, configuration (e.g., cross-connection) management, and Customer Network Management. Also see Mediation Functions.

Inter-Carrier Interface (ICI) – The interface between two networks that belong to different network providers or carriers.

Mediation Functions – Usually consist of Gateway Communications Functions, but can also include Information Processing Functions. When the two functions are combined, they are often contained in a stand-alone Mediation Device (MD), or packaged as an added module to an NE or equipment frame. Gateway Communications Functions often exist alone in a Gateway NE or Intermediate NE.

Pigtail – A length of optical fiber with one end terminated at a connector and the other end attached to a light source or detector. It is used to couple light from a source to a connectorized fiber cable or from a fiber cable to a detector.

Regenerator (RGTR) – A unidirectional device that can receive a digital signal and retransmit it in a form in which the amplitude, waveforms, and timing characteristics of the signal are constrained within specified limits. Typically, a SONET RGTR will be STE; however, it may relay some or all of the SOH bytes through (rather than terminating those bytes carried in the incoming signal and originating new bytes for transmission downstream). In cases where all of the SOH bytes are relayed through, the device is referred to as a Physical Layer RGTR. Also, in some cases “RGTR” is used to refer to a unidirectional device (as defined here), while in other cases it is used to refer to an NE consisting of two or more devices supporting signals in both directions of transmission.

STS-N Tandem Connection – A group of N STS path signals that are transported and maintained together through one or more tandem line systems, with the constituent SPE payload capacities unaltered. The STS tandem connection sub-layer falls between the SONET Line and STS path layers.

Super-rate Payload – A payload that is carried in an STS-Nc, STS-1-Xv or STS-3c-Xv SPE.

Synchronous – The essential characteristic of time scales or signals such that their corresponding significant instants occur (and are defined to occur) at precisely the same, or multiples of the same, average rate.

Terminal Multiplex (TM) – An ADM with a single high-speed OC-N interface.

Unequipped Channel – A portion of an STS-N, such as an STS or VT SPE, that is intentionally unoccupied.

User Channel – A channel that is allocated to the user for input (in a proprietary fashion) of information such as data communications for use in maintenance activities and remoting of alarms external to the span equipment. Unless stated otherwise, “user” in this document refers to the service provider.

VT-N Tandem Connection – A group of N VT path signals that are transported and maintained together through one or more tandem line systems, with the constituent SPE payload capacities unaltered. The VT tandem connection sub-layer falls between the SONET STS and VT path layers.

Work Station (WS) – Any one of a variety of Visual Display Terminals (VDTs), ranging from a simple keyboard and display to an intelligent, processor-controlled VDT.

Acronyms

ACI	Application Context Identifier
ACSE	Association Control Service Element
ADM	Add-Drop Multiplex
AFI	Authority and Format Identifier
AIS	Alarm Indication Signal
AIS-CI	AIS-Customer Installation
AIS-L	Line AIS
AIS-P	STS Path AIS
AIS-V	VT Path AIS
AISP	Automatic In-Service Provisioning
AISS-L	Line AIS Second
AISS-LFE	Far End Line AIS Second
AIMS	Acknowledged Information Transfer Service
AMI	Alternate Mark Inversion
ANSI	American National Standards Institute
APD	Avalanche Photodiode
API	Access Point Identifier
APS	Automatic Protection Switching
ASCII	American Standard Code for Information Interchange
ASE	Application Service Element
ASN.1	Abstract Syntax Notation 1
ATM	Asynchronous Transfer Mode
AU-n	Administrative Unit-n (n = 3 or 4)
AUG-N	Administrative Unit Group-N (N = 1, 4, 16, 64 or 256)
AUI	Attachment Unit Interface
B-ISDN	Broadband Integrated Services Digital Network
B3ZS	Bipolar with Three-Zero Substitution
B8ZS	Bipolar with Eight-Zero Substitution
BA	Booster Amplifier

BCD	Binary Coded Decimal
BER	Basic Encoding Rules
BER	Bit Error Ratio
BIP	Bit Interleaved Parity
BITS	Building Integrated Timing Supply
BLSR	Bidirectional Line Switched Ring
C-SMF	Conventional Single Mode Fiber
CC	Composite Clock
CC	Connection Confirm
CEV	Controlled Environmental Vault
CGA	Carrier Group Alarm
CLNP	Connectionless-mode Network Layer Protocol
CLNS	Connectionless-mode Network Service
CMI	Coded Mark Inversion
CMISE	Common Management Information Service Element
CO	Central Office
CONS	Connection-mode Network Service
CPE	Customer Premise Equipment
CPU	Central Processor Unit
<CR>	Carriage Return
CR	Conditional Requirement
CR	Connection Request
CR	Critical alarm
CRC	Cyclic Redundancy Check
CSMA/CD	Carrier Sense Multiple Access with Collision Detection
CSN	Complete Sequence Number
CSNP	Complete Sequence Number Packet
CTRL	LCAS Control
CV	Coding Violation
CV-L	Line Coding Violation
CV-LFE	Far End Line Coding Violation
CV-P	STS Path Coding Violation
CV-PFE	Far End STS Path Coding Violation
CV-S	Section Coding Violation
CV-V	VT Path Coding Violation
CV-VFE	Far End VT Path Coding Violation
DAP	Directory Access Protocol
DC	Direct Current
DCC	Data Communications Channel
DCF	Dispersion Compensating Fiber
DCN	Data Communications Network
DCS	Digital Cross-Connect System
DFB	Distributed Feedback

DFI	DSP Format Identifier
DGD	Digital Group Delay
DNR	Do Not Revert
DQDB	Distributed Queue Dual Bus
DR	Disconnection Request
DSA	Directory System Agent
DSAP	Destination Service Access Point
DSF	Dispersion Shifted Fiber
DSn	Digital Signal at level n (n = 0, 1, 1A, 1C, 2, 3 or 4NA)
DSNE	Directory Server NE
DSP	Domain Specific Part
DST	Dispersion Supported Transmission
DSX-n	Digital Signal Cross-connect at level n
DTAU	Digital Test Access Unit
DUA	Directory User Agent
DUS	DON'T USE for Synchronization
EA	Electro-Absorption
EC-1	Electrical Carrier - level 1
EDSX	Electronic Digital Signal Cross-connect
EIA	Electronic Industries Association
EMC	Electromagnetic Compatibility
EML	Element Management Layer
EMS	Element Management System
ENE	End NE
EOC	Embedded Operations Channel
EOW	Express Orderwire
ERDI	Enhanced Remote Defect Indication
ERDI-P	STS Path ERDI-P
ERDI-V	VT Path ERDI-P
ERFI	Enhanced Remote Failure Indication
ERFI-P	STS Path ERFI-P
ERFI-V	VT Path ERFI-P
ES	End System
ES	Errored Second
ES-L	Line ES
ES-LFE	Far End Line ES
ES-P	STS Path ES
ES-PFE	Far End STS Path ES
ES-S	Section ES
ES-V	VT Path ES
ES-VFE	Far End VT Path ES
ESA	ES Type A
ESB	ES Type B

ESD	Electrostatic Discharge
ESF	Extended Superframe
FC	Failure Count
FC-L	Line FC
FC-LFE	Far End Line FC
FC-P	STS Path FC
FC-PFE	Far End STS Path FC
FC-V	VT Path FC
FC-VFE	Far End VT Path FC
FD	Family of Documents
FDDI	Fiber Distributed Data Interface
FEAC	Far End Alarm and Control
FEBE	Far End Block Error (replaced with REI)
FEC	Forward Error Correction
FERF	Far End Receive Failure (replaced with RDI)
FFS	For Further Study
FR	Family of Requirements
FTAM	File Transfer Access and Management
FTP	File Transfer Protocol
G	Gigabits per second
GFP	Generic Framing Procedure
GID	Group Identification
GNE	Gateway NE
GR	Generic Requirements document
HDLC	High-level Data Link Control
ICI	Inter-Carrier Interface
ID	Identifier
IDI	Initial Domain Identifier
IDLC	Integrated Digital Loop Carrier
IDP	Initial Domain Part
IDRP	Inter Domain Routing Protocol
IDT	Integrated Digital Terminal
IEEE	Institute of Electrical and Electronics Engineers
IM	Indirect Modulation
INE	Intermediate NE
IR	Intermediate Reach
IS	Intermediate System
ISI	Intersymbol Interference
ISO	International Organization for Standardization
ISO DCC	ISO Data Country Code
ISP	International Standardized Profile
ITU-T	International Telecommunication Union – Telecommunication Standardization Sector

LAN	Local Area Network
LAPD	Link Access Protocol on the D-Channel
LBC	Laser Bias Current
LBO	Line Build Out
LCAS	Link Capacity Adjustment Scheme
LCD-P	STS Path Loss of Cell Delineation
LCN	Local Communications Network
LDB	Loop Detection Buffer
LED	Light-Emitting Diode
<LF>	Line Feed
LLC	Logical Link Control
LOF	Loss Of Frame
LOH	Line Overhead
LOP	Loss Of Pointer
LOP-P	STS Path LOP
LOP-V	VT Path LOP
LOS	Loss Of Signal
LOSS	Loss Of Signal Seconds
LOW	Local Orderwire
LR	Long Reach
LSB	Least Significant Bit
LSS	Link Status Signal
LTE	Line Terminating Equipment
MAC	Media Access Control
MAN	Metropolitan Area Network
MD	Mediation Device
MFAS	Multiframe Alignment Signal
MF11	1st Multiframe Indication
MF12	2nd Multiframe Indication
MJ	Major alarm
MLM	Multi-Longitudinal Mode
MN	Minor alarm
MPN	Mode Partition Noise
MSB	Most Significant Bit
MTAU	Metallic Test Access Unit
MTBF	Mean Time Between Failure
MTIE	Maximum Time Interval Error
MZ	Mach Zehnder
N-SEL	NSAP Selector
NA	Not Alarmed
NA	Not Applicable
NDF	New Data Flag
NE	Network Element

NET	Network Entity Title
NPDU	Network Protocol Data Unit
NPJC	Negative Pointer Justification Count
NR	Not Required
NRZ	Non-Return to Zero
NSA	Non-Service Affecting
NSAP	Network Service Access Point
NSIF	Network and Services Integration Forum
NZ-DSF	Non-Zero Dispersion Shifted Fiber
O	Objective
OAM&P	Operations, Administration, Maintenance, & Provisioning
OAR	Optically Amplified Receiver
OAT	Optically Amplified Transmitter
OC-N	Optical Carrier at level N (N = 1, 3, 12, 24, 48, 192 or 768)
OFA	Optical Fiber Amplifier
OOF	Out Of Frame
OOS	Out Of Service
OPR	Optical Power Received
OPT	Optical Power Transmitted
OR	Optical Reflectance
ORA	Out-of-Range Alert
ORG	Organization Identifier
ORL	Optical Return Loss
OS	Operations System
OSI	Open Systems Interconnection
OTU_k	Optical channel Transport Unit k (k = 1, 2 or 3)
OVTG	Optical VT Group
P-SEL	PSAP Selector
PA	Pre-Amplifier
PBX	Private Branch Exchange
PCH	Pre-Chirp
PCM	Pulse Code Modulation
PD	Payload Defect
PDC	Passive Dispersion Compensation
PDI	Payload Defect Indication
PDI-P	STS Path PDI
PDI-V	VT Path PDI
PDL	Polarization Dependent Loss
PDU	Protocol Data Unit
PIN	Positive-Intrinsic-Negative (photodiode)
PJ	Pointer Justification
PJCDiff	Pointer Justification Count Difference
PJCS	Pointer Justification Count Seconds

PKI	Public Key Infrastructure
PLCP	Physical Layer Convergence Procedure
PLL	Phase Lock Loop
PLM	Payload Label Mismatch
PLM-P	STS Path PLM
PLM-V	VT Path PLM
PM	Performance Monitoring
PMD	Polarization Mode Dispersion
PNO	Provisionable by the Network Operator
POH	Path Overhead
PPDU	Presentation Protocol Data Unit
PPJC	Positive Pointer Justification Count
PPP	Point-to-Point Protocol
PRS	Primary Reference Source or Stratum 1 Traceable
PSC	Protection Switching Count
PSC-L	Line PSC
PSC-P	STS Path PSC
PSC-V	VT Path PSC
PSD	Protection Switching Duration
PSD-L	Line PSD
PSD-P	STS Path PSD
PSD-V	VT Path PSD
PSN	Partial Sequence Number
PSNP	Partial Sequence Number Packet
PTE	Path Terminating Equipment
PVC	Permanent Virtual Circuit
QoS	Quality of Service
R	Requirement
RAI	Remote Alarm Indication
RAI-CI	RAI-Customer Installation
RBS	Robbed Bit Signaling
RDI	Remote Defect Indication
RDI-L	Line RDI
RDI-P	STS Path RDI
RDI-V	VT Path RDI
RDT	Remote Digital Terminal
REI	Remote Error Indication
REI-L	Line REI
REI-P	STS Path REI
REI-V	VT Path REI
RES	Reserved for Network Synchronization Use
RFI	Remote Failure Indication
RFI-L	Line RFI

RFI-P	STS Path RFI
RFI-V	VT Path RFI
RGTR	Regenerator
RIB	Routing Information Base
RNE	Remote NE
ROSE	Remote Operations Service Element
RPDU	Route Protocol Data Unit
RS-ACK	Re-Sequence Acknowledge
RTU	Remote Test Unit
RZ	Return to Zero
SA	Service Affecting
SAPI	Service Access Point Identifier
SAW	Surface Acoustic Wave
SBS	Stimulated Brillouin Scattering
SD	Signal Degrade
SDH	Synchronous Digital Hierarchy
SDU	Service Data Unit
SEF	Severely Errored Framing
SEFS	Severely Errored Framing Second
SEFS-S	Section SEFS
SES	Severely Errored Second
SES-L	Line SES
SES-LFE	Far End Line SES
SES-P	STS Path SES
SES-PFE	Far End STS Path SES
SES-S	Section SES
SES-V	VT Path SES
SES-VFE	Far End VT Path SES
SF	Signal Fail
SF	Superframe
SIF	SONET Interoperability Forum
SLM	Single Longitudinal Mode
SMC	SONET Minimum Clock (traceable)
SMF	Single Mode Fiber
SMSR	Side Mode Suppression Ratio
SNDCF	Subnetwork Dependent Convergence Function
SOH	Section Overhead
SONET	Synchronous Optical Network
SP	Segmentation Permitted
SPD	Spectral Power Density
SPE	Synchronous Payload Envelope
SPM	Self Phase Modulation
SQ	Sequence indicator

SR	Short Reach
SR	Special Report
SRS	Stimulated Raman Scattering
SS	Session Service
SS	Special Services
SSAP	Source Service Access Point
ST2	Stratum 2 (traceable)
ST3	Stratum 3 (traceable)
ST3E	Stratum 3E (traceable)
ST4	Stratum 4 (traceable)
STE	Section Terminating Equipment
STM-N	Synchronous Transport Module level N (N = 0, 1, 4, 16, 64 or 256)
STS	Synchronous Transport Signal
STS-1-Xv	X virtually concatenated STS-1s
STS-3c-Xv	X virtually concatenated STS-3cs
STS-N	STS level N (N = 1, 3, 12, 24, 48, 192 or 768)
STS-Nc	N contiguously concatenated STS-1s
STSX-N	STS electrical Cross-connect at level N (N = 1 or 3)
STU	Synchronized - Traceability Unknown
TARP	TID Address Resolution Protocol
TCA	Threshold Crossing Alert
TCP/IP	Transport Control Protocol/Internet Protocol
TDEV	Time Deviation
TDC	TARP Data Cache
TEF	TARP Echo Function
TEI	Terminal Endpoint Identifier
TEST-P	STS Supervisory-unequipped
TEST-V	VT Supervisory-unequipped
TIA	Telecommunications Industries Association
TID	Target Identification
TIM	Trace Identifier Mismatch
TIM-P	STS Path TIM
TIM-V	VT Path TIM
TL1	Transaction Language 1
TM	Terminal Multiplexer
TMN	Telecommunications Management Network
TNC	Transit Node Clock (traceable)
TP4	Transport Protocol Class 4
TPDU	Transport Protocol Data Unit
TR	Technical Reference document
TSAP	Transport Service Access Point
TSC	Test System Controller
TSG	Timing Signal Generator

TU-11	Tributary Unit-11
TUG-3	Tributary Unit Group-3
UAS	Unavailable Second
UAS-L	Line UAS
UAS-LFE	Far End Line UAS
UAS-P	STS Path UAS
UAS-PFE	Far End STS Path UAS
UAS-V	VT Path UAS
UAS-VFE	Far End VT Path UAS
UI	Unit Interval
UI_{pp}	UI peak-to-peak
UI_{rms}	UI root-mean-squared
UITS	Unacknowledged Information Transfer Service
UNEQ	Unequipped
UNEQ-P	STS Path UNEQ
UNEQ-V	VT Path UNEQ
UNI	User Network Interface
UPSR	Unidirectional Path Switched Ring
UR	Ultra-long Reach
URC	Update Remote Cache
USL	User System Language
UTC	Universal Time Coordinated
VC	Virtual Circuit
VC-n	Virtual Container-n (n = 11, 12, 2, 3 or 4)
VR	Very-long Reach
VSR	Very Short Reach
VT	Virtual Tributary
VTn	Virtual Container level n (n = 1.5, 2, 3 or 6)
VTn-Xv	X virtually concatenated VTns
WAN	Wide Area Network
WCh	Working Channel
WDM	Wavelength Division Multiplexing
WS	Workstation
WTR	Wait To Restore
ZCS	Zero Code Suppression

Requirement-Object Index

[1]	CR2-2	p. 2-3	[51]	Rx-x	p. A-1	[101]	R3-123	p. 3-91
[2]	CR2-3	p. 2-3	[52v2]	R3-63	p. 3-63	[102]	R3-124	p. 3-91
[3v2]	R2-4	p. 2-3	[53]	R3-64	p. 3-63	[103]	O3-125	p. 3-91
[4]	R3-1	p. 3-2	[54]	R3-65	p. 3-64	[104]	R3-126	p. 3-91
[5]	O3-2	p. 3-2	[55]	R3-66	p. 3-64	[105]	R3-127	p. 3-92
[6]	R3-3	p. 3-3	[56]	R3-68	p. 3-64	[106]	R3-128	p. 3-92
[7]	R3-4	p. 3-3	[57]	R3-70	p. 3-64	[107]	R3-129	p. 3-92
[8]	R3-5	p. 3-3	[58]	R3-71	p. 3-64	[108v2]	R3-130	p. 3-92
[9]	R3-6	p. 3-5	[59]	R3-72	p. 3-65	[109]	R3-132	p. 3-93
[10]	R3-7	p. 3-5	[60]	R3-73	p. 3-65	[110]	R4-1	p. 4-6
[11]	R3-8	p. 3-8	[61]	R3-74	p. 3-65	[111v2]	R4-17	p. 4-17
[12]	R3-9	p. 3-9	[62]	R3-75	p. 3-66	[112]	R4-13	p. 4-15
[13]	R3-13	p. 3-14	[63]	R3-77	p. 3-66	[113v2]	R4-15	p. 4-16
[14]	R3-14	p. 3-14	[64]	R3-78	p. 3-66	[114v2]	O4-18	p. 4-18
[15]	R3-15	p. 3-14	[65]	R3-79	p. 3-67	[115v3]	R4-2	p. 4-8
[16v2]	R3-19	p. 3-35	[66]	R3-80	p. 3-67	[116v2]	R4-4	p. 4-8
[17v2]	O3-21	p. 3-35	[67]	R3-81	p. 3-68	[117v3]	R4-5	p. 4-8
[18v3]	R3-23	p. 3-36	[68]	R3-82	p. 3-68	[118]	Ox-x	p. A-4
[19]	R3-26	p. 3-37	[69]	R3-87	p. 3-72	[119v3]	R4-8	p. 4-9
[20]	R3-29	p. 3-38	[70]	R3-88	p. 3-72	[120v2]	R4-10	p. 4-14
[21v2]	R3-30	p. 3-40	[71]	R3-89	p. 3-72	[121]	R4-14	p. 4-16
[22v2]	R3-31	p. 3-40	[72v2]	R3-102	p. 3-78	[122v2]	R4-11	p. 4-14
[23]	R3-33	p. 3-42	[73]	R3-92	p. 3-73	[123]	R4-19	p. 4-47
[24]	R3-34	p. 3-43	[74]	R3-93	p. 3-74	[124]	CR4-20	p. 4-47
[25v2]	R3-35	p. 3-43	[75]	R3-94	p. 3-74	[125]	R4-21	p. 4-47
[26]	R3-36	p. 3-43	[76]	R3-95	p. 3-75	[126]	R4-22	p. 4-47
[27v2]	R3-38	p. 3-47	[77]	R3-96	p. 3-75	[127v2]	R4-23	p. 4-47
[28]	R3-39	p. 3-49	[78]	R3-97	p. 3-75	[128v2]	CR4-24	p. 4-64
[29v2]	R3-40	p. 3-50	[79]	Rx-x	p. A-4	[129v2]	R4-26	p. 4-64
[30v3]	R3-41	p. 3-50	[80]	R3-98	p. 3-77	[130v2]	R4-27	p. 4-68
[31]	R3-42	p. 3-50	[81]	R3-99	p. 3-77	[131]	R5-1	p. 5-1
[32]	R3-45	p. 3-55	[82]	R3-100	p. 3-77	[132]	R5-2	p. 5-2
[33]	R3-46	p. 3-57	[83]	R3-101	p. 3-77	[133]	R5-3	p. 5-2
[34]	R3-48	p. 3-58	[84]	Rx-x	p. A-4	[134v2]	R5-6	p. 5-6
[35]	R3-49	p. 3-59	[85]	R3-106	p. 3-81	[135]	R5-7	p. 5-10
[36]	R3-50	p. 3-60	[86]	R3-107	p. 3-82	[136]	CR5-9	p. 5-14
[37]	O3-51	p. 3-60	[87]	R3-108	p. 3-83	[137]	O5-10	p. 5-14
[38]	CR3-52	p. 3-61	[88]	R3-109	p. 3-83	[138]	O5-11	p. 5-15
[39v2]	R3-53	p. 3-61	[89]	R3-110	p. 3-83	[139]	R5-12	p. 5-15
[40]	CR3-54	p. 3-61	[90]	O3-111	p. 3-83	[140]	R5-13	p. 5-15
[41]	CR3-55	p. 3-61	[91]	R3-112	p. 3-86	[141]	R5-14	p. 5-15
[42]	R3-56	p. 3-62	[92]	R3-113	p. 3-86	[142]	R5-15	p. 5-16
[43]	R3-57	p. 3-62	[93]	R3-114	p. 3-86	[143]	CRx-x	p. A-1
[44]	CR3-58	p. 3-62	[94]	R3-115	p. 3-86	[144]	CR5-16	p. 5-16
[45]	CR3-59	p. 3-62	[95]	R3-116	p. 3-86	[145]	CR5-17	p. 5-16
[46]	R3-60	p. 3-62	[96v2]	R3-117	p. 3-86	[146]	R5-18	p. 5-16
[47]	R3-61	p. 3-63	[97]	R3-119	p. 3-88	[147]	CR5-19	p. 5-16
[48]	CRx-x	p. A-1	[98]	R3-120	p. 3-89	[148]	CR5-20	p. 5-17
[49]	CR3-62	p. 3-63	[99]	R3-121	p. 3-90	[149]	CR5-21	p. 5-17
[50]	CRx-x	p. A-1	[100]	R3-122	p. 3-90	[150]	CR5-39	p. 5-25

[151] . . . R5-40 . . . p. 5-25	[203] . . . R5-98. . . p. 5-46	[255v3] . . R5-149 . . . p. 5-71
[152v2] . . R5-41 . . . p. 5-25	[204] . . . R5-99. . . p. 5-46	[256v3] . . R5-150 . . . p. 5-71
[153] . . . CR5-42 . . . p. 5-26	[205v2] . . R5-100 . . . p. 5-46	[257v2] . . R5-152 . . . p. 5-74
[154] . . . R5-43 . . . p. 5-26	[206] . . . R5-101 . . . p. 5-47	[258v2] . . R5-153 . . . p. 5-74
[155] . . . CR5-44 . . . p. 5-27	[207] . . . R5-103 . . . p. 5-47	[259v4] . . R5-154 . . . p. 5-76
[156] . . . CR5-45 . . . p. 5-27	[208] . . . O5-104 . . . p. 5-48	[260v2] . . O5-155 . . . p. 5-77
[157] . . . CR5-46 . . . p. 5-27	[209] . . . R5-105 . . . p. 5-49	[261v2] . . R5-166 . . . p. 5-81
[158] . . . CR5-47 . . . p. 5-27	[210] . . . R5-106 . . . p. 5-49	[262v2] . . O5-167 . . . p. 5-82
[159] . . . CR5-48 . . . p. 5-27	[211] . . . R5-107 . . . p. 5-49	[263v2] . . Ox-x. . . . p. A-1
[160] . . . R5-49 . . . p. 5-27	[212] . . . R5-108 . . . p. 5-50	[264v3] . . R5-151 . . . p. 5-74
[161] . . . O5-50 . . . p. 5-27	[213] . . . CR5-109 . . p. 5-50	[265v2] . . R5-160 . . . p. 5-81
[162] . . . R5-51 . . . p. 5-28	[214] . . . R5-110 . . . p. 5-50	[266] . . . CR5-161. . p. 5-81
[163] . . . CR5-52 . . . p. 5-28	[215] . . . R5-111 . . . p. 5-50	[267v3] . . R5-162 . . . p. 5-81
[164] . . . R5-53 . . . p. 5-28	[216] . . . R5-112 . . . p. 5-50	[268] . . . R5-163 . . . p. 5-81
[165] . . . R5-54 . . . p. 5-28	[217v2] . . R5-113 . . . p. 5-51	[269] . . . Rx-x. . . . p. A-2
[166] . . . CR5-55 . . . p. 5-29	[218] . . . Rx-x p. A-1	[270] . . . Rx-x. . . . p. A-2
[167] . . . R5-56 . . . p. 5-29	[219] . . . Rx-x p. A-1	[271v3] . . R5-168 . . . p. 5-82
[168] . . . R5-57 . . . p. 5-29	[220] . . . Rx-x p. A-1	[272v3] . . R5-169 . . . p. 5-82
[169] . . . R5-58 . . . p. 5-29	[221] . . . Rx-x p. A-1	[273v4] . . R5-175 . . . p. 5-84
[170] . . . O5-59 . . . p. 5-29	[222v2] . . R5-123 . . . p. 5-63	[274v3] . . CR5-176. . p. 5-85
[171] . . . R5-60 . . . p. 5-30	[223v2] . . R5-124 . . . p. 5-63	[275] . . . Rx-x. . . . p. A-2
[172] . . . O5-61 . . . p. 5-30	[224v2] . . R5-115 . . . p. 5-61	[276] . . . O5-186 . . . p. 5-87
[173v2] . . R5-62 . . . p. 5-30	[225v4] . . R5-116 . . . p. 5-61	[277] . . . CR5-187. . p. 5-87
[174] . . . R5-63 . . . p. 5-31	[226] . . . R5-125 . . . p. 5-63	[278] . . . R5-188 . . . p. 5-87
[175] . . . R5-64 . . . p. 5-31	[227] . . . R5-126 . . . p. 5-63	[279v3] . . R5-189 . . . p. 5-87
[176] . . . R5-66 . . . p. 5-32	[228] . . . Rx-x p. A-1	[280v3] . . R5-190 . . . p. 5-88
[177v3] . . O5-69 . . . p. 5-33	[229] . . . Rx-x p. A-1	[281v2] . . R5-191 . . . p. 5-88
[178v2] . . R5-73 . . . p. 5-35	[230] . . . Rx-x p. A-1	[282] . . . R5-192 . . . p. 5-88
[179] . . . R5-74 . . . p. 5-36	[231] . . . Rx-x p. A-1	[283] . . . R5-193 . . . p. 5-88
[180] . . . R5-75 . . . p. 5-36	[232] . . . Rx-x p. A-1	[284v2] . . R5-194 . . . p. 5-89
[181] . . . R5-76 . . . p. 5-38	[233] . . . Rx-x p. A-1	[285v3] . . R5-195 . . . p. 5-90
[182] . . . R5-77 . . . p. 5-39	[234] . . . CRx-x p. A-1	[286] . . . R5-196 . . . p. 5-91
[183] . . . R5-78 . . . p. 5-39	[235] . . . CRx-x p. A-1	[287v3] . . CR5-197. . p. 5-91
[184] . . . R5-79 . . . p. 5-40	[236] . . . Rx-x p. A-1	[288] . . . R5-198 . . . p. 5-91
[185] . . . R5-80 . . . p. 5-40	[237] . . . Rx-x p. A-1	[289] . . . R5-199 . . . p. 5-91
[186] . . . R5-81 . . . p. 5-40	[238] . . . Rx-x p. A-1	[290v2] . . Rx-x. . . . p. A-5
[187] . . . R5-82 . . . p. 5-40	[239] . . . CR5-130 . . p. 5-65	[291] . . . R5-200 . . . p. 5-92
[188] . . . R5-83 . . . p. 5-40	[240v2] . . R5-131 . . . p. 5-66	[292v2] . . R5-201 . . . p. 5-92
[189] . . . R5-84 . . . p. 5-41	[241] . . . R5-132 . . . p. 5-66	[293v2] . . Rx-x. . . . p. A-5
[190v2] . . R5-85 . . . p. 5-41	[242] . . . R5-133 . . . p. 5-66	[294] . . . R5-202 . . . p. 5-92
[191] . . . R5-86 . . . p. 5-41	[243v2] . . CR5-135 . . . p. 5-67	[295] . . . CR5-203. . p. 5-92
[192] . . . R5-87 . . . p. 5-42	[244v2] . . CRx-x p. A-1	[296] . . . R5-204 . . . p. 5-92
[193] . . . R5-88 . . . p. 5-42	[245v2] . . CRx-x p. A-1	[297] . . . R5-205 . . . p. 5-93
[194] . . . R5-89 . . . p. 5-42	[246v2] . . R5-137 . . . p. 5-68	[298v4] . . R5-206 . . . p. 5-93
[195] . . . R5-90 . . . p. 5-43	[247v2] . . R5-138 . . . p. 5-68	[299] . . . Rx-x. . . . p. A-2
[196v2] . . R5-91 . . . p. 5-43	[248] . . . R5-140 . . . p. 5-69	[300v5] . . R5-208 . . . p. 5-94
[197v3] . . R5-92 . . . p. 5-44	[249v3] . . Rx-x p. A-5	[301] . . . Rx-x. . . . p. A-2
[198] . . . R5-93 . . . p. 5-44	[250v2] . . Rx-x p. A-5	[302v2] . . R5-215 . . . p. 5-95
[199] . . . R5-94 . . . p. 5-44	[251] . . . CRx-x p. A-1	[303] . . . CRx-x p. A-2
[200] . . . R5-95 . . . p. 5-45	[252] . . . CRx-x p. A-1	[304] . . . CRx-x p. A-2
[201v2] . . R5-96 . . . p. 5-45	[253] . . . R5-144 . . . p. 5-69	[305] . . . R5-207 . . . p. 5-93
[202] . . . R5-97 . . . p. 5-46	[254] . . . R5-145 . . . p. 5-70	[306] . . . Rx-x. . . . p. A-2

[307]	Rx-x	p. A-2	[359]	R5-281	p. 5-134	[411]	R6-48	p. 6-12
[308]	Rx-x	p. A-2	[360v3]	R5-285	p. 5-137	[412v2]	O6-49	p. 6-12
[309]	Rx-x	p. A-2	[361v2]	R6-2	p. 6-2	[413]	R6-50	p. 6-12
[310]	R5-217	p. 5-97	[362]	Rx-x	p. A-2	[414]	O6-51	p. 6-12
[311]	R5-218	p. 5-97	[363]	Rx-x	p. A-2	[415v2]	R6-52	p. 6-12
[312]	R5-220	p. 5-98	[364]	Rx-x	p. A-2	[416v2]	R6-54	p. 6-16
[313v2]	R5-221	p. 5-99	[365]	Rx-x	p. A-2	[417v2]	R6-57	p. 6-17
[314v2]	R5-222	p. 5-99	[366]	R6-1	p. 6-2	[418v3]	R6-58	p. 6-17
[315v3]	R5-213	p. 5-95	[367]	R6-5	p. 6-4	[419]	R6-59	p. 6-18
[316]	R5-223	p. 5-99	[368]	R6-6	p. 6-4	[420]	R6-60	p. 6-18
[317v3]	R5-214	p. 5-95	[369v2]	R6-7	p. 6-4	[421v2]	R6-61	p. 6-18
[318]	R5-224	p. 5-99	[370]	R6-8	p. 6-5	[422v2]	R6-62	p. 6-18
[319v2]	R5-225	p. 5-99	[371]	R6-9	p. 6-5	[423]	R6-63	p. 6-19
[320v4]	O5-226	p. 5-100	[372]	O6-10	p. 6-5	[424]	R6-64	p. 6-19
[321v2]	R5-227	p. 5-100	[373]	R6-11	p. 6-5	[425]	O6-65	p. 6-19
[322v2]	R5-230	p. 5-101	[374v2]	R6-12	p. 6-5	[426v3]	R6-66	p. 6-19
[323v2]	R5-231	p. 5-101	[375]	R6-13	p. 6-5	[427]	R6-67	p. 6-19
[324]	R5-232	p. 5-101	[376]	R6-14	p. 6-5	[428v2]	R6-68	p. 6-19
[325]	R5-233	p. 5-102	[377]	CR6-15	p. 6-6	[429v2]	R6-69	p. 6-20
[326v3]	R5-234	p. 5-102	[378v2]	R6-16	p. 6-6	[430v4]	R6-73	p. 6-21
[327v4]	R5-235	p. 5-102	[379]	R6-17	p. 6-6	[431v3]	R6-77	p. 6-24
[328v2]	R5-240	p. 5-103	[380]	R6-18	p. 6-7	[432v2]	R6-78	p. 6-24
[329]	R5-241	p. 5-104	[381]	R6-19	p. 6-7	[433v2]	R6-83	p. 6-24
[330]	R5-242	p. 5-104	[382]	R6-20	p. 6-7	[434]	R6-85	p. 6-24
[331]	R5-243	p. 5-104	[383]	R6-21	p. 6-8	[435v2]	R6-86	p. 6-24
[332]	R5-244	p. 5-105	[384]	R6-22	p. 6-8	[436]	R6-88	p. 6-25
[333v2]	R5-245	p. 5-105	[385]	R6-23	p. 6-8	[437]	R6-89	p. 6-25
[334]	R5-246	p. 5-105	[386]	R6-24	p. 6-9	[438]	R6-90	p. 6-25
[335]	R5-247	p. 5-106	[387]	R6-26	p. 6-9	[439v2]	R6-91	p. 6-25
[336]	R5-248	p. 5-106	[388]	R6-27	p. 6-9	[440]	R6-92	p. 6-25
[337v2]	R5-251	p. 5-108	[389]	R6-28	p. 6-9	[441]	R6-93	p. 6-26
[338]	R5-253	p. 5-113	[390]	R6-29	p. 6-9	[442]	CR6-94	p. 6-26
[339]	R5-252	p. 5-112	[391]	R6-30	p. 6-9	[443]	O6-95	p. 6-26
[340v2]	R5-254	p. 5-113	[392]	R6-31	p. 6-9	[444v3]	R6-96	p. 6-26
[341]	R5-255	p. 5-114	[393v2]	R6-32	p. 6-10	[445v2]	R6-97	p. 6-26
[342]	R5-256	p. 5-115	[394]	R6-33	p. 6-10	[446v2]	R6-99	p. 6-28
[343v4]	R5-257	p. 5-115	[395v2]	R6-34	p. 6-10	[447]	R6-100	p. 6-28
[344]	Ox-x	p. A-2	[396]	Rx-x	p. A-2	[448v2]	R6-101	p. 6-28
[345]	R5-259	p. 5-117	[397v2]	R6-35	p. 6-10	[449]	R6-102	p. 6-28
[346]	R5-260	p. 5-118	[398]	R6-36	p. 6-10	[450]	R6-103	p. 6-28
[347]	R5-261	p. 5-119	[399]	CR6-37	p. 6-11	[451]	R6-104	p. 6-28
[348]	R5-262	p. 5-119	[400]	R6-38	p. 6-11	[452]	R6-105	p. 6-29
[349]	O5-263	p. 5-119	[401]	Rx-x	p. A-2	[453]	O6-106	p. 6-29
[350]	R5-266	p. 5-122	[402]	R6-39	p. 6-11	[454]	R6-107	p. 6-29
[351]	R5-267	p. 5-122	[403]	R6-40	p. 6-11	[455]	R6-108	p. 6-29
[352]	R5-268	p. 5-123	[404]	R6-41	p. 6-11	[456]	R6-109	p. 6-29
[353]	R5-269	p. 5-124	[405]	R6-42	p. 6-11	[457]	R6-110	p. 6-29
[354]	O5-270	p. 5-124	[406]	R6-43	p. 6-11	[458]	R6-111	p. 6-29
[355]	R5-271	p. 5-125	[407]	R6-44	p. 6-11	[459]	R6-112	p. 6-30
[356]	O5-272	p. 5-125	[408]	R6-45	p. 6-11	[460]	R6-113	p. 6-30
[357v2]	R5-273	p. 5-129	[409]	R6-46	p. 6-11	[461v2]	R6-114	p. 6-30
[358]	R5-279	p. 5-132	[410]	R6-47	p. 6-12	[462]	R6-115	p. 6-30

[463]	...	R6-116	...	p. 6-30	[515]	...	R6-193	...	p. 6-49	[567v2]	...	R6-250	...	p. 6-59
[464]	...	R6-117	...	p. 6-30	[516]	...	R6-194	...	p. 6-49	[568v2]	...	R6-252	...	p. 6-61
[465]	...	R6-123	...	p. 6-31	[517]	...	R6-195	...	p. 6-49	[569v3]	...	O6-253	...	p. 6-61
[466]	...	R6-124	...	p. 6-31	[518]	...	R6-196	...	p. 6-49	[570v2]	...	R6-255	...	p. 6-62
[467v2]	...	R6-125	...	p. 6-32	[519v3]	...	R6-197	...	p. 6-49	[571v2]	...	R6-257	...	p. 6-62
[468]	...	R6-126	...	p. 6-32	[520]	...	Rx-x	...	p. A-2	[572v2]	...	O6-258	...	p. 6-62
[469]	...	R6-132	...	p. 6-32	[521v2]	...	R6-198	...	p. 6-50	[573]	...	R6-259	...	p. 6-62
[470]	...	R6-133	...	p. 6-33	[522]	...	Rx-x	...	p. A-2	[574v3]	...	R6-261	...	p. 6-63
[471]	...	R6-134	...	p. 6-33	[523]	...	R6-199	...	p. 6-50	[575]	...	Rx-x	...	p. A-2
[472]	...	R6-135	...	p. 6-34	[524]	...	R6-200	...	p. 6-50	[576v3]	...	R6-263	...	p. 6-63
[473]	...	O6-136	...	p. 6-34	[525]	...	R6-202	...	p. 6-50	[577]	...	Rx-x	...	p. A-2
[474]	...	Rx-x	...	p. A-2	[526]	...	R6-203	...	p. 6-51	[578v2]	...	R6-264	...	p. 6-63
[475]	...	Rx-x	...	p. A-2	[527]	...	R6-204	...	p. 6-51	[579v2]	...	R6-266	...	p. 6-63
[476]	...	R6-137	...	p. 6-34	[528v5]	...	R6-205	...	p. 6-51	[580v4]	...	R6-267	...	p. 6-65
[477]	...	O6-138	...	p. 6-34	[529]	...	R6-206	...	p. 6-51	[581v2]	...	R6-269	...	p. 6-65
[478v2]	...	R6-139	...	p. 6-34	[530]	...	Rx-x	...	p. A-2	[582v2]	...	R6-270	...	p. 6-65
[479]	...	R6-140	...	p. 6-34	[531v2]	...	R6-207	...	p. 6-51	[583v2]	...	R6-271	...	p. 6-65
[480v2]	...	R6-141	...	p. 6-34	[532]	...	Rx-x	...	p. A-2	[584v2]	...	R6-275	...	p. 6-66
[481]	...	R6-142	...	p. 6-34	[533]	...	R6-208	...	p. 6-51	[585v2]	...	R6-283	...	p. 6-67
[482]	...	O6-143	...	p. 6-34	[534]	...	R6-209	...	p. 6-52	[586v2]	...	R6-276	...	p. 6-66
[483]	...	Rx-x	...	p. A-2	[535]	...	R6-210	...	p. 6-52	[587v2]	...	R6-285	...	p. 6-67
[484]	...	Rx-x	...	p. A-2	[536]	...	R6-212	...	p. 6-52	[588v2]	...	R6-286	...	p. 6-67
[485v2]	...	R6-144	...	p. 6-34	[537]	...	R6-213	...	p. 6-52	[589v2]	...	CR6-288	...	p. 6-68
[486v2]	...	O6-145	...	p. 6-35	[538]	...	R6-214	...	p. 6-52	[590]	...	Rx-x	...	p. A-2
[487]	...	Rx-x	...	p. A-2	[539v2]	...	R6-215	...	p. 6-53	[591v3]	...	CR6-287	...	p. 6-68
[488]	...	R6-146	...	p. 6-35	[540]	...	R6-217	...	p. 6-53	[592v2]	...	R6-294	...	p. 6-69
[489v2]	...	R6-147	...	p. 6-35	[541v2]	...	R6-218	...	p. 6-53	[593v2]	...	CR6-292	...	p. 6-68
[490]	...	Rx-x	...	p. A-2	[542]	...	R6-219	...	p. 6-54	[594]	...	Rx-x	...	p. A-2
[491v2]	...	R6-148	...	p. 6-37	[543]	...	Rx-x	...	p. A-2	[595v2]	...	R6-295	...	p. 6-69
[492v2]	...	O6-149	...	p. 6-37	[544v2]	...	CR6-220	...	p. 6-54	[596v2]	...	R6-298	...	p. 6-69
[493]	...	Rx-x	...	p. A-2	[545]	...	Rx-x	...	p. A-2	[597v2]	...	O6-299	...	p. 6-69
[494]	...	Rx-x	...	p. A-2	[546]	...	Rx-x	...	p. A-2	[598]	...	R6-300	...	p. 6-69
[495v2]	...	R6-150	...	p. 6-37	[547]	...	R6-223	...	p. 6-54	[599]	...	O6-301	...	p. 6-69
[496v2]	...	O6-151	...	p. 6-37	[548]	...	R6-224	...	p. 6-54	[600]	...	CRx-x	...	p. A-2
[497v2]	...	R6-152	...	p. 6-37	[549v2]	...	R6-227	...	p. 6-55	[601]	...	Rx-x	...	p. A-2
[498]	...	R6-153	...	p. 6-37	[550v2]	...	R6-228	...	p. 6-55	[602]	...	Rx-x	...	p. A-2
[499]	...	R6-154	...	p. 6-37	[551v2]	...	R6-229	...	p. 6-55	[603]	...	Rx-x	...	p. A-2
[500]	...	Rx-x	...	p. A-2	[552v3]	...	R6-231	...	p. 6-56	[604]	...	Rx-x	...	p. A-2
[501]	...	R6-155	...	p. 6-38	[553v3]	...	R6-232	...	p. 6-56	[605]	...	Rx-x	...	p. A-3
[502]	...	O6-156	...	p. 6-38	[554]	...	R6-233	...	p. 6-56	[606]	...	Ox-x	...	p. A-3
[503]	...	Rx-x	...	p. A-2	[555]	...	R6-234	...	p. 6-57	[607]	...	Rx-x	...	p. A-3
[504]	...	Rx-x	...	p. A-2	[556v2]	...	R6-236	...	p. 6-58	[608]	...	Ox-x	...	p. A-3
[505v2]	...	R6-157	...	p. 6-38	[557v3]	...	O6-237	...	p. 6-58	[609v2]	...	R6-305	...	p. 6-71
[506v2]	...	O6-158	...	p. 6-38	[558v2]	...	R6-239	...	p. 6-58	[610v4]	...	R6-302	...	p. 6-71
[507]	...	Rx-x	...	p. A-2	[559v2]	...	R6-241	...	p. 6-58	[611v2]	...	R6-306	...	p. 6-71
[508]	...	R6-159	...	p. 6-38	[560v2]	...	O6-242	...	p. 6-59	[612]	...	R6-303	...	p. 6-71
[509]	...	R6-160	...	p. 6-38	[561]	...	R6-243	...	p. 6-59	[613]	...	Rx-x	...	p. A-3
[510]	...	Rx-x	...	p. A-2	[562v3]	...	R6-245	...	p. 6-59	[614]	...	Rx-x	...	p. A-3
[511]	...	Rx-x	...	p. A-2	[563]	...	Rx-x	...	p. A-2	[615]	...	Rx-x	...	p. A-3
[512v2]	...	R6-190	...	p. 6-48	[564v3]	...	R6-247	...	p. 6-59	[616]	...	Rx-x	...	p. A-3
[513v2]	...	R6-191	...	p. 6-49	[565]	...	Rx-x	...	p. A-2	[617v2]	...	R6-308	...	p. 6-72
[514]	...	R6-192	...	p. 6-49	[566v2]	...	R6-248	...	p. 6-59	[618v2]	...	R6-309	...	p. 6-72

[619] . . . Rx-x . . . p. A-3	[671] . . . Rx-x . . . p. A-3	[723v3] . . O6-418 . . p. 6-142
[620] . . . R6-318 . . p. 6-92	[672] . . . R6-366 . . p. 6-119	[724v3] . . CR6-419 . p. 6-142
[621] . . . R6-310 . . p. 6-87	[673] . . . R6-367 . . p. 6-119	[725v2] . . R6-161 . . . p. 6-43
[622v2] . . R6-311 . . . p. 6-87	[674] . . . R6-368 . . p. 6-122	[726] . . . Rx-x p. A-3
[623v2] . . R6-319 . . . p. 6-93	[675] . . . R6-369 . . p. 6-123	[727] . . . Rx-x p. A-3
[624] . . . Rx-x p. A-3	[676] . . . R6-373 . . p. 6-123	[728] . . . R6-433 . . . p. 6-144
[625] . . . R6-315 . . . p. 6-92	[677] . . . R6-374 . . p. 6-123	[729] . . . CR6-435 . p. 6-144
[626v3] . . R6-312 . . . p. 6-88	[678] . . . R6-375 . . p. 6-123	[730] . . . O6-434 . . . p. 6-144
[627v2] . . R6-316 . . . p. 6-92	[679] . . . R6-376 . . p. 6-123	[731v4] . . R6-448 . . . p. 6-146
[628v2] . . R6-317 . . . p. 6-92	[680] . . . R6-377 . . p. 6-123	[732v3] . . R6-453 . . . p. 6-147
[629v3] . . R6-313 . . . p. 6-89	[681] . . . R6-378 . . p. 6-123	[733] . . . Rx-x p. A-3
[630] . . . Rx-x p. A-3	[682] . . . R6-381 . . p. 6-125	[734] . . . Ox-x p. A-3
[631] . . . Rx-x p. A-3	[683v4] . R6-379 . . . p. 6-124	[735v4] . . R6-451 . . . p. 6-146
[632] . . . R6-320 . . . p. 6-93	[684] . . . Rx-x p. A-3	[736v3] . . R6-452 . . . p. 6-146
[633] . . . R6-323 . . . p. 6-94	[685v2] . CR6-382 . p. 6-135	[737] . . . R6-457 . . . p. 6-148
[634v2] . . R6-324 . . . p. 6-100	[686] . . . CR6-385 . p. 6-135	[738] . . . R6-460 . . . p. 6-148
[635v3] . . R6-329 . . . p. 6-101	[687v2] . R6-386 . . . p. 6-135	[739] . . . Rx-x p. A-3
[636] . . . R6-326 . . . p. 6-100	[688v2] . R6-387 . . . p. 6-136	[740v2] . . R6-423 . . . p. 6-143
[637v2] . . R6-327 . . . p. 6-100	[689v2] . R6-389 . . . p. 6-136	[741] . . . R6-424 . . . p. 6-143
[638v3] . . R6-331 . . . p. 6-101	[690] . . . R6-391 . . . p. 6-136	[742] . . . Rx-x p. A-3
[639v2] . . R6-325 . . . p. 6-100	[691v4] . R6-392 . . . p. 6-136	[743] . . . R6-425 . . . p. 6-143
[640v3] . . R6-337 . . . p. 6-103	[692] . . . Rx-x p. A-3	[744] . . . R6-427 . . . p. 6-144
[641] . . . R6-338 . . . p. 6-104	[693] . . . O6-394 . . . p. 6-137	[745] . . . R6-426 . . . p. 6-144
[642] . . . R6-339 . . . p. 6-104	[694v4] . R6-395 . . . p. 6-137	[746] . . . R6-428 . . . p. 6-144
[643v5] . . O6-341 . . . p. 6-108	[695v2] . R6-321 . . . p. 6-94	[747] . . . R6-429 . . . p. 6-144
[644] . . . Rx-x p. A-3	[696] . . . Ox-x p. A-3	[748v2] . . R6-430 . . . p. 6-144
[645v3] . . R6-348 . . . p. 6-110	[697] . . . Rx-x p. A-3	[749] . . . O6-464 . . . p. 6-151
[646] . . . Ox-x p. A-3	[698] . . . Rx-x p. A-3	[750] . . . R6-465 . . . p. 6-151
[647] . . . Ox-x p. A-3	[699] . . . Ox-x p. A-3	[751v2] . . R6-466 . . . p. 6-151
[648v2] . . Rx-x p. A-3	[700] . . . Rx-x p. A-3	[752] . . . Rx-x p. A-3
[649] . . . Rx-x p. A-3	[701] . . . Rx-x p. A-3	[753] . . . Rx-x p. A-3
[650] . . . Ox-x p. A-3	[702] . . . Rx-x p. A-3	[754] . . . Rx-x p. A-3
[651] . . . Rx-x p. A-3	[703v3] . O6-322 . . . p. 6-94	[755v2] . . R6-462 . . . p. 6-149
[652] . . . Rx-x p. A-3	[704] . . . Rx-x p. A-3	[756] . . . R6-467 . . . p. 6-151
[653] . . . Rx-x p. A-3	[705] . . . CRx-x p. A-3	[757v2] . . O6-468 . . . p. 6-152
[654] . . . Rx-x p. A-3	[706] . . . R6-400 . . . p. 6-139	[758] . . . Rx-x p. A-3
[655] . . . Rx-x p. A-3	[707] . . . O6-401 . . . p. 6-139	[759] . . . R6-470 . . . p. 6-153
[656v2] . . R6-349 . . . p. 6-111	[708] . . . Ox-x p. A-3	[760] . . . R6-471 . . . p. 6-154
[657] . . . R6-350 . . . p. 6-111	[709] . . . O6-402 . . . p. 6-139	[761] . . . O7-3 p. 7-2
[658] . . . CR6-351 . p. 6-111	[710] . . . R6-403 . . . p. 6-139	[762v2] . . R7-4 p. 7-2
[659v2] . . R6-352 . . . p. 6-111	[711] . . . R6-404 . . . p. 6-139	[763] . . . R7-5 p. 7-2
[660v2] . . R6-353 . . . p. 6-111	[712] . . . R6-405 . . . p. 6-140	[764] . . . R7-6 p. 7-3
[661] . . . R6-355 . . . p. 6-115	[713] . . . R6-406 . . . p. 6-140	[765] . . . R7-7 p. 7-3
[662] . . . R6-356 . . . p. 6-115	[714] . . . R6-407 . . . p. 6-140	[766] . . . R7-8 p. 7-3
[663] . . . R6-357 . . . p. 6-115	[715] . . . CR6-408 . p. 6-140	[767] . . . R7-9 p. 7-3
[664v2] . . R6-358 . . . p. 6-115	[716] . . . Rx-x p. A-3	[768] . . . R7-10 p. 7-3
[665] . . . Ox-x p. A-3	[717] . . . R6-412 . . . p. 6-140	[769] . . . Rx-x p. A-3
[666] . . . Rx-x p. A-3	[718] . . . R6-413 . . . p. 6-141	[770] . . . Rx-x p. A-4
[667] . . . R6-359 . . . p. 6-115	[719] . . . R6-414 . . . p. 6-141	[771] . . . R7-11 p. 7-3
[668] . . . R6-360 . . . p. 6-115	[720] . . . O6-415 . . . p. 6-141	[772] . . . R7-12 p. 7-3
[669] . . . R6-361 . . . p. 6-118	[721] . . . R6-416 . . . p. 6-141	[773] . . . R7-13 p. 7-4
[670] . . . R6-362 . . . p. 6-118	[722v3] . O6-417 . . . p. 6-142	[774] . . . R7-14 p. 7-4

[775] . . . Ox-xp. A-4	[827] . . . R8-59. . . p. 8-19	[879] . . . R8-112 . . p. 8-35
[776v2] . R8-1.p. 8-6	[828] . . . R8-60. . . p. 8-19	[880v2] . . R8-113 . . p. 8-35
[777v2] . R8-2.p. 8-6	[829] . . . R8-61. . . p. 8-19	[881v2] . . R8-114 . . p. 8-35
[778v2] . R8-3.p. 8-6	[830] . . . R8-62. . . p. 8-19	[882] . . . R8-117 . . p. 8-36
[779v2] . O8-4.p. 8-7	[831] . . . R8-63. . . p. 8-20	[883v2] . . R8-119 . . p. 8-36
[780v2] . R8-5.p. 8-7	[832] . . . R8-64. . . p. 8-20	[884] . . . R8-120 . . p. 8-36
[781v2] . R8-6.p. 8-7	[833] . . . R8-65. . . p. 8-20	[885] . . . R8-121 . . p. 8-36
[782v2] . R8-7.p. 8-7	[834] . . . R8-66. . . p. 8-20	[886] . . . R8-122 . . p. 8-36
[783] . . . R8-8.p. 8-8	[835] . . . O8-67 . . p. 8-20	[887] . . . R8-124 . . p. 8-39
[784v3] . R8-9.p. 8-8	[836] . . . R8-68. . . p. 8-20	[888] . . . R8-125 . . p. 8-39
[785] . . . R8-10p. 8-9	[837] . . . R8-69. . . p. 8-20	[889] . . . R8-126 . . p. 8-39
[786] . . . R8-11p. 8-10	[838] . . . R8-70. . . p. 8-20	[890] . . . R8-127 . . p. 8-41
[787v3] . R8-12p. 8-10	[839] . . . R8-72. . . p. 8-21	[891v2] . . R8-128 . . p. 8-41
[788v2] . R8-15p. 8-10	[840] . . . Rx-xp. A-4	[892] . . . CRx-x. . . .p. A-4
[789] . . . Rx-xp. A-4	[841] . . . R8-73. . . p. 8-21	[893] . . . R8-130 . . p. 8-44
[790v3] . R8-16p. 8-11	[842] . . . R8-74. . . p. 8-21	[894] . . . R8-132 . . p. 8-44
[791v2] . R8-20p. 8-12	[843] . . . R8-75. . . p. 8-21	[895] . . . R8-133 . . p. 8-45
[792v2] . CRx-x. . . .p. A-6	[844v2] . . R8-76. . . p. 8-21	[896v2] . . R8-134 . . p. 8-45
[793v4] . CR8-17 . . p. 8-11	[845v2] . . R8-77. . . p. 8-21	[897] . . . R8-135 . . p. 8-45
[794v2] . R8-22p. 8-12	[846] . . . R8-78. . . p. 8-21	[898] . . . R8-136 . . p. 8-45
[795v3] . R8-23p. 8-12	[847] . . . R8-79. . . p. 8-22	[899] . . . R8-137 . . p. 8-45
[796v2] . R8-25p. 8-13	[848] . . . R8-80. . . p. 8-22	[900v2] . . R8-138 . . p. 8-45
[797v2] . R8-29p. 8-13	[849] . . . R8-81. . . p. 8-22	[901] . . . R8-139 . . p. 8-45
[798] . . . R8-30p. 8-13	[850v3] . . R8-82. . . p. 8-22	[902] . . . R8-140 . . p. 8-45
[799v2] . R8-31p. 8-14	[851] . . . R8-83. . . p. 8-23	[903] . . . R3-69 . . . p. 3-64
[800v2] . R8-32p. 8-14	[852] . . . R8-84. . . p. 8-23	[904] . . . R3-76 . . . p. 3-66
[801v2] . R8-33p. 8-14	[853] . . . R8-85. . . p. 8-23	[905] . . . R3-83 . . . p. 3-68
[802v2] . R8-34p. 8-14	[854] . . . R8-87. . . p. 8-23	[906] . . . R3-90 . . . p. 3-72
[803v2] . R8-35p. 8-14	[855] . . . R8-88. . . p. 8-25	[907] . . . R3-91 . . . p. 3-72
[804] . . . R8-36p. 8-15	[856] . . . R8-89. . . p. 8-25	[908] . . . O3-118 . . p. 3-87
[805] . . . R8-37p. 8-15	[857] . . . R8-90. . . p. 8-25	[909] . . . O3-131 . . p. 3-93
[806] . . . R8-38p. 8-15	[858] . . . R8-91. . . p. 8-26	[910] . . . CR4-25 . . p. 4-64
[807v2] . R8-39p. 8-15	[859] . . . R8-92. . . p. 8-26	[911v3] . . R5-4.p. 5-2
[808] . . . R8-40p. 8-15	[860] . . . R8-93. . . p. 8-26	[912] . . . R5-5.p. 5-2
[809v2] . R8-41p. 8-15	[861] . . . R8-94. . . p. 8-26	[913] . . . CR5-65 . . p. 5-32
[810] . . . R8-42p. 8-16	[862] . . . R8-95. . . p. 8-27	[914] . . . R5-70 . . . p. 5-34
[811] . . . R8-43p. 8-16	[863] . . . R8-96. . . p. 8-27	[915] . . . CR5-71 . . p. 5-34
[812] . . . R8-44p. 8-16	[864] . . . R8-97. . . p. 8-27	[916] . . . R5-72 . . . p. 5-34
[813v4] . R8-45p. 8-16	[865] . . . R8-98. . . p. 8-27	[917] . . . O5-102 . . p. 5-47
[814v2] . R8-47p. 8-16	[866] . . . R8-99. . . p. 8-27	[918v2] . . R5-114 . . p. 5-58
[815] . . . Rx-xp. A-4	[867] . . . R8-100 . . p. 8-27	[919v2] . . R5-127 . . p. 5-64
[816] . . . R8-48p. 8-16	[868] . . . R8-101 . . p. 8-27	[920v2] . . R5-128 . . p. 5-64
[817] . . . R8-49p. 8-17	[869v2] . . R8-102 . . p. 8-33	[921v2] . . R5-139 . . p. 5-68
[818] . . . CR8-50 . . p. 8-17	[870] . . . O8-103 . . p. 8-33	[922v2] . . R5-156 . . p. 5-79
[819v2] . R8-51p. 8-17	[871] . . . O8-104 . . p. 8-33	[923v2] . . R5-141 . . p. 5-69
[820] . . . O8-52p. 8-17	[872] . . . O8-105 . . p. 8-33	[924v2] . . R5-142 . . p. 5-69
[821] . . . O8-53p. 8-17	[873] . . . R8-106 . . p. 8-33	[925v2] . . R5-143 . . p. 5-69
[822v2] . R8-54p. 8-17	[874] . . . O8-107 . . p. 8-33	[926] . . . R5-147 . . p. 5-70
[823] . . . R8-55p. 8-19	[875] . . . R8-108 . . p. 8-33	[927] . . . Rx-x.p. A-4
[824v2] . R8-56p. 8-19	[876] . . . R8-109 . . p. 8-34	[928v3] . . O5-165 . . p. 5-81
[825] . . . R8-57p. 8-19	[877] . . . R8-110 . . p. 8-34	[929] . . . Ox-x.p. A-4
[826] . . . R8-58p. 8-19	[878v2] . . R8-111 . . p. 8-35	[930v2] . . R5-164 . . p. 5-81

[931]	...	R5-258	..	p. 5-117	[983]	...	R6-293	..	p. 6-69	[1035]	...	Ox-x	...	p. A-6
[932]	...	R5-274	..	p. 5-130	[984]	...	CR6-296	..	p. 6-69	[1036]	...	O8-21	...	p. 8-12
[933]	...	R5-275	..	p. 5-130	[985]	...	CR6-297	..	p. 6-69	[1037]	...	R8-71	...	p. 8-20
[934]	...	R5-276	..	p. 5-130	[986]	...	R6-304	..	p. 6-71	[1038]	...	R8-86	...	p. 8-23
[935]	...	R5-277	..	p. 5-130	[987]	...	R6-307	..	p. 6-71	[1039]	...	O8-115	...	p. 8-35
[936]	...	O5-278	..	p. 5-130	[988]	...	R6-314	..	p. 6-91	[1040v2]	...	R8-116	...	p. 8-35
[937v2]	...	O5-286	..	p. 5-137	[989]	...	R6-328	..	p. 6-100	[1041]	...	R8-118	...	p. 8-36
[938]	...	R6-3	...	p. 6-2	[990]	...	R6-330	..	p. 6-101	[1042]	...	R8-123	...	p. 8-38
[939]	...	R6-4	...	p. 6-4	[991v2]	...	R6-333	..	p. 6-102	[1043]	...	R8-129	...	p. 8-41
[940]	...	O6-55	...	p. 6-17	[992v2]	...	R6-336	..	p. 6-103	[1044v2]	...	R3-103	...	p. 3-79
[941]	...	R6-70	...	p. 6-20	[993]	...	Rx-x	...	p. A-4	[1045]	...	O3-104	...	p. 3-80
[942v2]	...	CR6-71	..	p. 6-20	[994v3]	...	O6-342	..	p. 6-109	[1046]	...	R5-67	...	p. 5-33
[943v2]	...	CR6-72	..	p. 6-21	[995v2]	...	R6-380	..	p. 6-125	[1047]	...	R5-68	...	p. 5-33
[944v2]	...	R6-74	...	p. 6-21	[996]	...	R6-398	..	p. 6-138	[1048v2]	...	CR5-117	..	p. 5-61
[945]	...	R6-79	...	p. 6-24	[997v3]	...	R6-162	..	p. 6-43	[1049]	...	CR5-118	..	p. 5-62
[946]	...	R6-80	...	p. 6-24	[998]	...	Ox-x	...	p. A-4	[1050]	...	R5-119	...	p. 5-62
[947]	...	R6-81	...	p. 6-24	[999v2]	...	CR6-443	..	p. 6-145	[1051]	...	R5-129	...	p. 5-65
[948]	...	R6-82	...	p. 6-24	[1000v3]	...	R6-444	..	p. 6-145	[1052v2]	...	CR5-183	..	p. 5-87
[949]	...	R6-84	...	p. 6-24	[1001v3]	...	O6-171	..	p. 6-44	[1053v2]	...	R5-185	...	p. 5-87
[950]	...	R6-87	...	p. 6-25	[1002v3]	...	CR6-449	..	p. 6-146	[1054]	...	CRx-x	...	p. A-5
[951]	...	R6-216	...	p. 6-53	[1003]	...	Rx-x	...	p. A-6	[1055]	...	Rx-x	...	p. A-5
[952]	...	CR6-221	..	p. 6-54	[1004v2]	...	R6-450	..	p. 6-146	[1056]	...	Rx-x	...	p. A-5
[953]	...	O6-222	..	p. 6-54	[1005v3]	...	R6-454	..	p. 6-147	[1057]	...	CRx-x	...	p. A-5
[954]	...	R6-225	...	p. 6-54	[1006v3]	...	R6-455	..	p. 6-147	[1058]	...	Rx-x	...	p. A-5
[955]	...	R6-226	...	p. 6-55	[1007v3]	...	CR6-456	..	p. 6-147	[1059]	...	Rx-x	...	p. A-6
[956]	...	O6-230	..	p. 6-56	[1008]	...	R6-458	..	p. 6-148	[1060]	...	Rx-x	...	p. A-6
[957v2]	...	O6-235	..	p. 6-58	[1009]	...	CR6-459	..	p. 6-148	[1061]	...	CR5-237	..	p. 5-103
[958]	...	R6-238	...	p. 6-58	[1010]	...	O6-469	..	p. 6-152	[1062]	...	R5-238	...	p. 5-103
[959v2]	...	O6-240	..	p. 6-58	[1011v2]	...	CR8-24	..	p. 8-12	[1063]	...	R5-239	...	p. 5-103
[960]	...	R6-244	...	p. 6-59	[1012]	...	R8-131	...	p. 8-44	[1064v2]	...	R6-127	...	p. 6-32
[961]	...	R6-246	...	p. 6-59	[1013]	...	R5-8	...	p. 5-11	[1065v2]	...	R6-128	...	p. 6-32
[962v3]	...	R6-249	...	p. 6-59	[1014]	...	Rx-x	...	p. A-5	[1066v2]	...	R6-129	...	p. 6-32
[963v2]	...	O6-251	..	p. 6-61	[1015]	...	R5-171	..	p. 5-82	[1067v2]	...	R6-130	...	p. 6-32
[964]	...	R6-254	...	p. 6-61	[1016]	...	CR5-172	..	p. 5-83	[1068v2]	...	R6-131	...	p. 6-32
[965v2]	...	O6-256	..	p. 6-62	[1017v2]	...	R5-173	..	p. 5-83	[1069v2]	...	R6-164	...	p. 6-43
[966]	...	R6-260	...	p. 6-62	[1018v2]	...	CR5-179	..	p. 5-86	[1070v2]	...	R6-166	...	p. 6-44
[967]	...	R6-262	...	p. 6-63	[1019]	...	R5-180	..	p. 5-86	[1071]	...	R6-167	...	p. 6-44
[968v3]	...	R6-265	...	p. 6-63	[1020]	...	R5-181	..	p. 5-86	[1072]	...	R6-169	...	p. 6-44
[969]	...	R6-268	...	p. 6-65	[1021]	...	R5-182	..	p. 5-86	[1073]	...	R6-170	...	p. 6-44
[970]	...	CR6-272	..	p. 6-66	[1022]	...	R5-212	..	p. 5-95	[1074]	...	O6-168	...	p. 6-44
[971]	...	R6-273	...	p. 6-66	[1023v2]	...	R5-219	..	p. 5-97	[1075]	...	R6-172	...	p. 6-45
[972]	...	R6-274	...	p. 6-66	[1024v2]	...	R5-236	..	p. 5-102	[1076]	...	O6-173	...	p. 6-45
[973]	...	R6-277	...	p. 6-66	[1025]	...	R5-264	..	p. 5-121	[1077]	...	R6-174	...	p. 6-45
[974]	...	R6-278	...	p. 6-66	[1026]	...	R5-265	..	p. 5-121	[1078]	...	R6-175	...	p. 6-45
[975]	...	R6-279	...	p. 6-67	[1027]	...	R5-280	..	p. 5-132	[1079v2]	...	R6-201	...	p. 6-50
[976]	...	CR6-280	..	p. 6-67	[1028]	...	R5-282	..	p. 5-134	[1080v2]	...	R6-211	...	p. 6-52
[977]	...	R6-281	...	p. 6-67	[1029]	...	R5-283	..	p. 5-135	[1081]	...	R6-332	...	p. 6-101
[978]	...	R6-282	...	p. 6-67	[1030]	...	R5-284	..	p. 5-135	[1082]	...	R6-344	...	p. 6-109
[979]	...	R6-284	...	p. 6-67	[1031v2]	...	R5-287	..	p. 5-137	[1083]	...	O6-345	...	p. 6-109
[980]	...	R6-289	...	p. 6-68	[1032v2]	...	O5-288	..	p. 5-138	[1084]	...	O6-347	...	p. 6-109
[981]	...	R6-290	...	p. 6-68	[1033v2]	...	O6-354	..	p. 6-112	[1085]	...	CR6-363	..	p. 6-118
[982]	...	R6-291	...	p. 6-68	[1034v3]	...	R6-396	..	p. 6-137	[1086]	...	R6-364	...	p. 6-119

[1087] . . . R6-365 . . . p. 6-119	[1139] . . . R3-11. . . . p. 3-11	[1191] . . . R6-118 . . . p. 6-30
[1088] . . . CR6-370 . . . p. 6-123	[1140] . . . R3-12. . . . p. 3-11	[1192] . . . R6-119 . . . p. 6-30
[1089] . . . R6-371 . . . p. 6-123	[1141] . . . R3-16. . . . p. 3-28	[1193] . . . R6-120 . . . p. 6-31
[1090] . . . R6-372 . . . p. 6-123	[1142] . . . R3-17. . . . p. 3-32	[1194] . . . R6-121 . . . p. 6-31
[1091] . . . CR6-383 . . . p. 6-135	[1143] . . . R3-18. . . . p. 3-32	[1195] . . . R6-122 . . . p. 6-31
[1092] . . . O6-384 . . . p. 6-135	[1144] . . . R3-20. . . . p. 3-35	[1196] . . . CR6-163. . . p. 6-43
[1093] . . . R6-388 . . . p. 6-136	[1145] . . . O3-22 p. 3-35	[1197] . . . CR6-165. . . p. 6-44
[1094] . . . O6-390 . . . p. 6-136	[1146] . . . R3-25. . . . p. 3-36	[1198] . . . CR6-176. . . p. 6-45
[1095v2] . . . R6-393 . . . p. 6-136	[1147] . . . R3-32. . . . p. 3-41	[1199] . . . R6-177 . . . p. 6-45
[1096] . . . O6-420 . . . p. 6-142	[1148] . . . R3-37. . . . p. 3-44	[1200] . . . R6-178 . . . p. 6-46
[1097] . . . O6-421 . . . p. 6-142	[1149] . . . R3-43. . . . p. 3-51	[1201] . . . CR6-179. . . p. 6-46
[1098] . . . O6-422 . . . p. 6-142	[1150] . . . R3-44. . . . p. 3-52	[1202] . . . R6-180 . . . p. 6-46
[1099v2] . . . R8-46 p. 8-16	[1151] . . . CR3-47. . . . p. 3-57	[1203] . . . R6-181 . . . p. 6-46
[1100] . . . R2-1. p. 2-3	[1152] . . . R3-67. . . . p. 3-64	[1204] . . . O6-182 . . . p. 6-46
[1101v2] . . . R3-10 p. 3-9	[1153] . . . R3-84. . . . p. 3-70	[1205] . . . R6-183 . . . p. 6-46
[1102] . . . R3-24 p. 3-36	[1154] . . . R3-85. . . . p. 3-70	[1206] . . . R6-184 . . . p. 6-46
[1103v2] . . . CR3-27 p. 3-38	[1155] . . . R3-86. . . . p. 3-71	[1207] . . . O6-185 . . . p. 6-47
[1104v2] . . . R3-28 p. 3-38	[1156] . . . R3-105. . . . p. 3-80	[1208] . . . R6-186 . . . p. 6-47
[1105] . . . Rx-x p. A-4	[1157] . . . R4-3 p. 4-8	[1209] . . . O6-187 . . . p. 6-47
[1106] . . . R4-9. p. 4-13	[1158] . . . R4-6 p. 4-9	[1210] . . . R6-188 . . . p. 6-47
[1107v2] . . . R4-12 p. 4-14	[1159] . . . R4-7 p. 4-9	[1211] . . . R6-189 . . . p. 6-47
[1108v2] . . . R4-16 p. 4-16	[1160] . . . R5-22. p. 5-17	[1212] . . . CR6-334. . . p. 6-102
[1109v2] . . . R5-134 p. 5-67	[1161] . . . O5-23 p. 5-17	[1213] . . . R6-335 . . . p. 6-102
[1110] . . . O5-136 p. 5-68	[1162] . . . R5-24. p. 5-18	[1214] . . . R6-343 . . . p. 6-109
[1111v2] . . . R5-148 p. 5-70	[1163] . . . R5-25. p. 5-18	[1215] . . . R6-346 . . . p. 6-109
[1112] . . . Ox-x p. A-6	[1164] . . . R5-26. p. 5-18	[1216] . . . O6-431 . . . p. 6-144
[1113] . . . O5-177 p. 5-85	[1165] . . . O5-27 p. 5-18	[1217] . . . CR6-432. . . p. 6-144
[1114] . . . O5-178 p. 5-85	[1166] . . . CR5-28. p. 5-19	[1218] . . . CR6-436. . . p. 6-145
[1115] . . . CR5-184 p. 5-87	[1167] . . . R5-29. p. 5-19	[1219] . . . O6-437 . . . p. 6-145
[1116] . . . O5-210 p. 5-94	[1168] . . . CR5-30. p. 5-19	[1220] . . . CR6-438. . . p. 6-145
[1117] . . . R5-211 p. 5-95	[1169] . . . R5-31. p. 5-19	[1221] . . . CR6-439. . . p. 6-145
[1118] . . . O5-216 p. 5-97	[1170] . . . R5-32. p. 5-19	[1222] . . . CR6-440. . . p. 6-145
[1119] . . . O5-249 p. 5-106	[1171] . . . R5-33. p. 5-20	[1223] . . . CR6-441. . . p. 6-145
[1120] . . . R5-250 p. 5-106	[1172] . . . R5-34. p. 5-22	[1224] . . . O6-442 . . . p. 6-145
[1121] . . . Ox-x p. A-6	[1173] . . . R5-35. p. 5-24	[1225] . . . O6-445 . . . p. 6-146
[1122] . . . O6-25 p. 6-9	[1174] . . . R5-36. p. 5-24	[1226] . . . CR6-446. . . p. 6-146
[1123] . . . O6-53 p. 6-12	[1175] . . . O5-37 p. 5-24	[1227] . . . R6-447 . . . p. 6-146
[1124] . . . R6-56 p. 6-17	[1176] . . . R5-38. p. 5-24	[1228] . . . R6-461 . . . p. 6-148
[1125] . . . R6-98 p. 6-27	[1177] . . . CR5-120 p. 5-63	[1229] . . . CR6-463. . . p. 6-150
[1126] . . . CR6-340 p. 6-108	[1178] . . . R5-121 p. 5-63	[1230] . . . CR8-13 . . . p. 8-10
[1127v2] . . . R6-397 p. 6-137	[1179] . . . R5-122 p. 5-63	[1231] . . . R8-14 p. 8-10
[1128v2] . . . R6-399 p. 6-138	[1180] . . . R5-146 p. 5-70	[1232] . . . CR8-26 . . . p. 8-13
[1129] . . . R6-409 p. 6-140	[1181] . . . O5-157. p. 5-79	[1233] . . . CR8-27 . . . p. 8-13
[1130] . . . R6-410 p. 6-140	[1182] . . . R5-158 p. 5-80	[1234] . . . R8-28 p. 8-13
[1131] . . . R6-411 p. 6-140	[1183] . . . O5-159 p. 5-80	
[1132] . . . R7-1. p. 7-1	[1184] . . . O5-170 p. 5-82	
[1133] . . . R7-2. p. 7-1	[1185] . . . R5-174 p. 5-83	
[1134v2] . . . O7-15 p. 7-5	[1186] . . . R5-209 p. 5-94	
[1135v2] . . . R7-16 p. 7-6	[1187] . . . O5-228 p. 5-100	
[1136] . . . Rx-x p. A-7	[1188] . . . R5-229 p. 5-101	
[1137] . . . R8-18 p. 8-11	[1189] . . . O6-75 p. 6-22	
[1138] . . . R8-19 p. 8-12	[1190] . . . O6-76 p. 6-22	



TELCORDIA ENTERPRISE LICENSE AGREEMENT AND LIMITED WARRANTY

For Technical Documents: Generic Requirements (GRs), Special Reports (SRs), Technical References (TRs), Technical Advisories (TAs), Family of Requirements (FRs), Family of Documents (FDs), Framework Advisories (FAs), Science Technologies (STs), Message Driven Program (MDPs), Information Publications (IPs), Audio Visuals (AVs) and Telcordia Practices (BRs)

IMPORTANT! READ CAREFULLY

USE OF THIS PRODUCT INDICATES THAT YOU (LICENSEE OR USER) HAVE READ AND ACCEPT THE TERMS OF THIS AGREEMENT. IF YOU DO NOT AGREE WITH THE TERMS OF THIS AGREEMENT, PROMPTLY RETURN THE UNUSED PRODUCT WITH ALL SEALS INTACT TO THE ADDRESS LISTED BELOW FOR A TELCORDIA CREDIT.

1. LICENSE GRANT

Telcordia grants to customer ("Licensee") a non-exclusive, non-transferable, limited license to use this Licensed Product by employees of Licensee ("Users") for internal business purposes only. All intellectual property rights, title and interest in all Licensed Products furnished to Licensee remain in Telcordia. This License does not preclude the execution of additional license agreements with Licensee for the Licensed Product(s).

Telcordia has exclusive rights to all Licensed Products which are protected by United States and international copyright laws.

2. LICENSEE'S USE:

a) Licensee may place the Licensed Products on a Local Area Network, Wide Area Network, server, internal web site, or other electronic computing platform shared or accessible to employees or affiliates of Licensee. Licensee may make paper and electronic copies of Licensed Products as determined by Licensee to be necessary for Licensee's internal purposes; provided all copies, in whole or in part, of the Licensed Products shall bear the same Telcordia copyright and disclaimer notices legend as appear on the Licensed Products originally furnished to Licensee by Telcordia.

b) Subject to the preceding paragraph, Licensee may reproduce and distribute Licensed Products to "Affiliates" defined as (i) the parent entity (corporation or partnership) which directly or indirectly owns the majority of the outstanding shares or interests of Licensee, (ii) a sibling entity (corporation or partnership) the majority of whose outstanding shares or interests are owned by its parent entity, or (iii) a subsidiary entity (corporation or partnership) the majority of whose outstanding shares or interests are owned by Licensee, provided, however, that such entity shall continue to remain an Affiliate hereunder only as long as the applicable ownership interest as described above exists.

Licensee may sublicense the rights granted in this section to an Affiliate, provided Licensee shall remain responsible for any breach by such Affiliate. Licensee shall ensure that such Affiliate as assignee agrees to be bound by the rights, obligations and limitations set forth herein, and such Affiliate shall be responsible for any breach by such Affiliate and Licensee shall ensure that Telcordia shall have the right of direct enforcement of such obligations against such Affiliate. If a direct enforcement claim is denied, for any reason, it is agreed that Licensor may assert such claim against Licensee.

c) Licensee may copy portions of Licensed Products to create specifications and related documentation (the "Licensee Documentation").

d) Licensee may, in marketing a product or related services (collectively, "Licensee Product"), (i) make reference to the Licensed Product utilized in the development of Licensee Product; provided that Licensee shall make no statement, representation or warranty on behalf of Telcordia including but not limited to a certification by Telcordia of a product's or related service's compliance with the Licensed Product, unless otherwise agreed to by the parties in writing; or (ii) distribute the

Licensee Documentation to a third party prior to sale of the Licensee Product.

e) Licensee may refer to and/or incorporate portions of such Licensed Products in the Licensee Documentation for the Licensee Product and copy the Licensee Documentation for distribution in conjunction with the sale of the Licensee Product to any third party so long as the original Telcordia and copyright legends, as applicable, are acknowledged on the specifications and/or documentation.

f) Licensee must treat the Licensed Product(s) like any other copyrighted material.

g) Except as otherwise stated, it is understood that the foregoing license does not include the right to make copies of the Licensed Products for sale to third parties or to create derivative works for sale.

USER MAY NOT:

a) Copy the Licensed Product, except as provided above;

b) Make copies of the Licensed Product or portions thereof as are permitted above for internal purposes that contain provisions that conflict or differ in content from comparable provisions of the Licensed Product, unless such differences are identified specifically, and it is made clear in such copies that the results are not part of the Licensed Product;

c) Transfer the Licensed Product to another party, except as provided above;

d) Licensee may not make the Licensed Product available, in whole or in part for the purposes of external distribution to third parties other than Affiliates.

e) Grant sublicenses, leases, or other rights to the Licensed Product or rent the Licensed Products to others, except as provided above; or

f) Make telecommunications data transmissions of the Licensed Product to the public or any third party.

g) Data, in whole or in part, may not be extracted from the Licensed Product(s) for use in any derivative Licensee product or used to verify and subsequently modify data in any Licensee product which is sold, licensed or otherwise provided to third parties unless Licensee has executed a separately negotiated Special License Agreement with Telcordia, except as provided above.

3. AUDITS

Upon reasonable written notice to Licensee, Telcordia shall have the right to review Licensee's compliance with the terms and conditions of this License Agreement ("Agreement"). If such review reveals a violation of the requirements set forth herein, in addition to any other remedies it may have, Telcordia may terminate this Agreement in accordance with the Termination section of this Agreement.

4. FEES AND PAYMENTS

All fees and charges due hereunder shall be paid in full within thirty (30) days of the date of the invoice. Overdue payments are subject to a late payment charge, calculated and compounded monthly, and calculated at an annual rate of either (1) one percent (1%) over the prime rate available in New York City, as published in *The Wall Street Journal* on the first Monday (or the next bank business day) following the payment due date; or (2) 12 percent (12%), whichever shall be

higher. If the amount of the late payment charge exceeds the maximum permitted by law, the charge will be reduced to that maximum amount.

Licensee shall pay or reimburse Telcordia for all sales or use taxes, duties, or levies imposed by any authority, government or government agency (other than those levied on the net income of Telcordia) in connection with this Agreement. If Telcordia is required to collect a tax to be paid by Licensee, Licensee shall pay this tax on demand. If Licensee fails to pay these taxes, duties or levies, Licensee shall pay all reasonable expenses incurred by Telcordia, including reasonable attorney's fees, to collect such taxes, duties or levies.

Telcordia shall provide Licensee with one (1) Copy of the Licensed Product. Any additional copies in cd or paper media will be provided to Licensee at a cost of \$75.00 per copy. Please contact our Customer Call Center noted below.

5. LIMITED WARRANTY

Telcordia warrants that the media on which the Licensed Product is provided is free from defects in materials and workmanship for 90 days. Licensee's sole remedy for breach of this warranty is Telcordia's Product Replacement Plan described below. This warranty applies only to the original Licensee.

6. DISCLAIMER OF WARRANTIES

EXCEPT AS SET FORTH ABOVE, THE LICENSED PRODUCT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, EVEN IF TELCORDIA HAS BEEN MADE AWARE OF SUCH PURPOSE, OR ANY WARRANTY AGAINST INFRINGEMENT OF PATENTS OR OTHER INTELLECTUAL PROPERTY RIGHTS. LICENSEE ASSUMES RESPONSIBILITY FOR THE SELECTION OF THE LICENSED PRODUCT TO ACHIEVE ITS INTENDED RESULTS, AND FOR THE USE AND RESULTS OBTAINED FROM THE LICENSED PRODUCT.

7. LIMITATION OF LIABILITY

THE ENTIRE LIABILITY OF TELCORDIA, AND LICENSEE'S EXCLUSIVE REMEDY, IS THE REPLACEMENT OF ANY LICENSED PRODUCT WHICH DOES NOT MEET THE TELCORDIA LIMITED WARRANTY AND IS RETURNED TO TELCORDIA WITHIN 90 DAYS.

IN NO EVENT WILL TELCORDIA BE LIABLE TO LICENSEE FOR ANY DAMAGES, INCLUDING DIRECT DAMAGES, LOST PROFITS, OR OTHER INDIRECT, SPECIAL, INCIDENTAL, EXEMPLARY OR CONSEQUENTIAL DAMAGES ARISING OUT OF THIS AGREEMENT, EVEN IF TELCORDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE WARRANTY GIVES LICENSEE SPECIFIC LEGAL RIGHTS, AND LICENSEE MAY ALSO HAVE OTHER RIGHTS WHICH VARY FROM STATE TO STATE. SOME STATES DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THE ABOVE LIMITATION MAY NOT APPLY TO LICENSEE.

8. THIRD PARTY PRODUCTS AND INFORMATION WARRANTY

Telcordia does not warrant Third Party products or information which Telcordia may use to prepare the Licensed Product. Third Party products or information may be warranted by Third Parties as expressly provided in the documentation accompanying the Third Party product or information, if any. Licensee's exclusive remedy under any Third Party warranty is as provided in the Third Party documentation accompanying the Third Party product or information, if any.

9. LICENSED PRODUCT REPLACEMENT PLAN

During the first 30 days after Licensee licenses the Telcordia Licensed Product, Telcordia will replace at no charge any Licensed Product which is returned to Telcordia because its media is defective in materials or workmanship. Returns for replacement of a defective Licensed Product should be sent postpaid to Telcordia using the Return Policy procedures stated below.

10. RETURN POLICY

Licensed Product(s) may be returned within 30 days of receipt for Telcordia credit only. Returned Licensed Products must be in their original packaging with all seals intact. Returns not found to be defective in materials or workmanship will be subject to a 10%

restocking fee. Licensed Products that have been delivered electronically (downloaded from the Superstore) are **not** eligible for credits, refunds or returns, even if duplicative with Licensed Products that are the subject of prior or contemporaneous orders. Licensee assumes all responsibility for managing its inventory of Licensed Product(s).

11. TERMINATION

If Licensee or its User breaches one or more of its obligations under this Agreement, Telcordia may elect at any time, in addition to any other remedy, to terminate the license and rights granted. Prior to the termination, Telcordia must give Licensee two (2) months written notice specifying the breach. Telcordia may terminate the license and rights granted if Licensee does not remedy all breaches specified in the written notice within the two (2) month notice period. Upon termination of the license and rights granted, Licensee shall destroy or return all Licensed Product(s) and Documentation, including all copies, and certify in writing to Telcordia the destruction or return.

12. PUBLICITY

Notwithstanding anything herein to the contrary, each party is prohibited from using in advertising, publicity, promotion, marketing, or other similar activity, any name, trade name, trademark, or other designation including any abbreviation, contraction or simulation of the other without the prior, express, written permission of the other.

13. GENERAL

Reexport. Licensee acknowledges that any commodities and/or technical data provided under this Agreement is subject to the Export Administration Regulations ("the EAR") administered by the U.S. Commerce Department and that any export or re-export thereof must be in compliance with the EAR. Licensee agrees that it shall not export or reexport, directly or indirectly, either during the term of this Agreement or after its expiration, any commodities and/or technical data (or direct products thereof) provided under this Agreement in any form to destinations in Country Groups D:1 or E:2, as specified in Supplement No. 1 to Part 740 of the EAR, and as modified from time to time by the U.S. Department of Commerce, or to recipients or destinations that are otherwise controlled or embargoed under U.S. law.

Foreign Tax Payment. For a Licensee which is not a United States corporation, Telcordia will not accept remittance of less than the full amount billed to Licensee as full payment unless:

- a) Licensee withholds that amount to satisfy tax withholding requirements imposed by the country (other than the United States) in which Licensee resides or in which Licensee has accepted delivery of the Licensed Product; and
- b) Licensee furnishes a receipt issued by the withholding tax jurisdiction and certifying deposit of the withheld amount into its treasury or other tax depository to Telcordia's sole credit, or a certification on Licensee's stationery that Licensee has deposited the withheld amount into its tax jurisdiction's treasury or other tax depository to Telcordia's sole credit.

Further, to ensure the orderly processing of Telcordia tax returns, Licensee shall provide to Telcordia a summary of all amounts withheld during the year no later than ten business days after December 31 of each year.

Governing Law. This Agreement is a contract between Telcordia and the Licensee of the Licensed Product. This contract is to be interpreted in the federal and state courts of New Jersey, in accordance with the laws of the State of New Jersey without regard to its conflict of laws principles, and the parties consent to the jurisdiction of such courts for this purpose.

Entire Agreement. Licensee further agree that this is the complete and exclusive statement of the Agreement between Licensee and Telcordia and supersedes any proposal or prior Agreement, oral or written, or any other communication between us relating to the subject matter of this Agreement.

All questions about this Agreement should be directed to:

Telcordia Technologies, Inc.
Customer Service Center
One Telcordia Drive, RRC 1B180
Piscataway, NJ 08854

Phone: 1.866.672.6997 (USA)
+1.732.699.6700 (Worldwide)

END OF TERMS AND CONDITIONS

Agreed:

Company: _____

—

Name: _____

Signature: _____

Date: _____

Revised 3/06