PRENTICE HALL SERIES IN COMPUTER NETWORKING AND DISTRIBUTED SYSTEMS

NEXT GENERATION OPTICAL NETWORKS

The Convergence of IP Intelligence and Optical Technologies



Peter Tomsu • Christian Schmutzer

Exhibit 1016
IPR2023-00581
U.S. Patent 8,886,772

NEXT GENERATION OPTICAL NETWORKS

The Convergence of IP Intelligence and Optical Technology

ISBN 0-13-028226-X



PRENTICE HALL SERIES IN COMPUTER NETWORKING AND DISTRIBUTED SYSTEMS

NEXT GENERATION OPTICAL NETWORKS

The Convergence of IP Intelligence and Optical Technology

Peter Tomsu • Christian Schmutzer



Prentice Hall PTR
Upper Saddle River, NJ 07458
www.phptr.com

Library of Congress Catalog Card Number: 2001033944

Editorial/production supervision: *Donna Cullen-Dolce*
Acqusition Editor: *Mary Franz*
Editorial Assistant: *Noreen Regina*
Marketing Manager: *Dan DePasquale*
Manufacturing Manager: *Alexis R. Heydt*
Cover Design Director: *Jerry Votta*
Interior Design: *Gail Cocker-Bogusz*



© 2002 Prentice Hall PTR
Prentice-Hall, Inc.
Upper Saddle River, NJ 07458

All rights reserved. No part of this book may be
reproduced, in any form or by any means, without
permission in writing from the publisher

The publisher offers discounts on this book when ordered in bulk quantities.
For more information, contact
Corporate Sales Department,
Prentice Hall PTR
One Lake Street
Upper Saddle River, NJ 07458
Phone: 800-382-3419; FAX: 201-236-714
E-mail (Internet): corpsales@prenhall.com

Printed in the United States of America

10 9 8 7 6 5 4 3 2

ISBN 0-13-028226-X

Prentice-Hall International (UK) Limited, *London*
Prentice-Hall of Australia Pty. Limited, *Sydney*
Prentice-Hall Canada Inc., *Toronto*
Prentice-Hall Hispanoamericana, S.A., *Mexico*
Prentice-Hall of India Private Limited, *New Delhi*
Prentice-Hall of Japan, Inc., *Tokyo*
Prentice-Hall (Singapore) Pte. Ltd., *Singapore*
Editora Prentice-Hall do Brasil, Ltda., *Rio de Janeiro*

To our family and friends

Contents

Acknowledgments xix

Preface xxi

Chapter 1

Introduction to Carrier Network Architectures 1

IP as Unifying Protocol Infrastructure 1

The Traditional Carrier Network Architecture 2

Achieving Long-Haul Connectivity with Optical Transmission 4

Delivering Multiplexed Services 4

Synchronous Transmission Standards 5

Building a Complex SONET/SDH Carrier Network 7

Multiplexing and Framing Standards 9

Providing Carrier Class Network Availability 13

Adding Versatility with Asynchronous Transfer Mode 14

ATM Reference Model 15

Introducing Statistical Multiplexing 16

Transport Multiple Services across a Single Network 17

AAL1 18

AAL2 19

AAL3/4 20

AAL5 20

ATM Services 22

Permanent Virtual Connection 23
Switched Virtual Connection 23
ATM Virtual Connections 23
ATM Switching Operation 24
ATM Addressing 24
Subnetwork Model of Addressing 25
NSAP-Format ATM Addresses 25
ATM Address Common Components 25
ATM Quality of Service 26
ATM Signaling and Connection Establishment Overview 27
Signaling Standards 28
ATM Connection Establishment Procedure 28
Connection Request Routing and Negotiation 29
Providing Internet Services at the IP Layer 29
Introducing the IP Layer on Top of the ATM Layer 29
Using Multiprotocol Label Switching for Delivering IP Services 31
Scalability Issues of Standard IP Routing Protocols 31
The Concept of Using Labels as Forwarding Information 32
MPLS-Based Applications 36
Next-Generation Carrier Networks 38
Eliminating Intermediate Layers in the Backbone 39
IP over ATM over SONET/SDH over WDM (1) 40
IP over ATM over WDM (2) 41
IP over SONET/SDH over WDM (3) 41
IP over WDM (4) 41
Handling Fiber Capacity Exhaust 41
Adding Intelligence to the Optical Layer 42
Summary 44
Recommended Reading 44

Chapter 2

Optical Networking Standardization 47

Standardization Activities and Targets 47
Activities of the T1 Committee and Its T1X1 Subcommittee 48
Optical Networking Standardization Framework (G.871) 48
OTN Architecture (G.872) 50
Layered Structure 50
OCH Layer Network 51
Optical Multiplex Section (OMS) Layer Network 52
Optical Transmission Section (OTS) Layer Network 52
Interlayer Adaptation 53
Client Signal/OCH Adaptation 53
OCH/OMS Adaptation 53
OMS/OTS Adaptation 53
Activities of the Institute of Electrical and Electronics Engineering (IEEE) 54
10-Gbps Ethernet 54
Resilient Packet Rings 55
Activities of the Optical Internetworking Forum (OIF) 57
OIF Mission 57
Standard Physical Layer Interface 59
IP-Centric Control and Signaling for Optical Paths 60
Optical User-to-Network Interface (O-UNI) 60
Activities of the IETF 64
Summary 64
Recommended Reading 65

chapter 3

Optical Networking Technology Fundamentals 67

Optical Transmission Technologies 67

Wavelength Division Multiplexing (WDM) 68

Attenuation 70

Dispersion 71

Nonlinear Effects 73

Optical Signal-to-Noise Ratio (OSNR) 74

Optical Transmission Systems 74

Optical Amplifiers 74

Erbium-Doped Fiber Amplifiers (EDFAs) 75

Praseodynamium-Doped Fiber Amplifiers (PDFAs) 77

Amplifier Control 78

DWDM Systems 79

Wavelength Routers 83

Electrical Wavelength Routers 85

Hybrid Wavelength Routers 86

Optical Wavelength Routers 87

Data Transmission Technologies 90

Packet over SONET/SDH (POS) 91

PPP Encapsulation 92

Framing 94

Interface Format 95

Transmission Rate 96

Payload Scrambling 97

PPP over SONET/SDH at OC-192c/STM-64c 98

Parallel Implementation 98

Word-Oriented POS Framing 98

Automatic Protection Switching (APS) 101

APS Channel Protocol 101
Cisco's APS Implementation 103
APS Switchover Process 105
APS Network Layer Interaction 105
Dynamic Packet Transport (DPT) 105
Encapsulation 107
IP Data Packets 109
ATM Data Packets 110
Control Packets 110
Usage Packets 112
SRP Physical Layer Implementation 112
Framing 112
Interface Format 113
Transmission Rate 113
Packet Handling Procedures 113
Topology Discovery 113
Packet Processing 115
Multicasting 116
Packet Priority 117
SRP Fairness Algorithm 118
Intelligent Protection Switching 120
IPS Messages 121
IPS Message Handling State Machine 124
Node Pass-Through Mode 126
Network Management 126
Clocking and Synchronization 126
Data Transmission Technology Comparison 127
Encapsulation Overhead 127
Framing Overhead 128
MPLS Traffic Engineering (MPLS-TE) 133

MPLS-TE Architecture 133
Why MPLS-TE 133
MPLS-TE Components 135
Defining the Term *Traffic Trunk* 137
MPLS-TE Attributes 138
Trunk Attributes 138
Bandwidth 138
Resource Class Affinity 138
Path Selection Policy 138
Priority/Preemption 139
Adaptability 139
Resilience 139
Resource Attributes 140
Available Bandwidth 140
Resource Class 140
Administrative Weight 140
Information Distribution 140
Bandwidth Accounting 141
OSPF Extensions 142
Router Address TLV 145
Link TLV 145
IS-IS Extensions 146
Router ID TLV 147
Extended IP Reachability TLV 147
Extended IS Reachability TLV 147
Path Computation and Selection 148
LSP Tunnel Setup 149
LSP Setup Procedure 149
LSP Setup Request 150
Admission Control 151
Explicit Routing 151
Label Allocation and LSP Establishment 151
RSVP Extensions 152
IP Routing using MPLS-TE 153

MPLS-TE Protection 157
Path Protection 157
Restoration Time 159
Optimizing Path Protection 159
Link/Node Protection 160
Network Survivability Principles 161
Defining Survivability 162
Survivability Concepts 163
Protection 164
Restoration 164
Protection Techniques 165
Protection Types 165
1 + 1 Protection 165
1:1 Protection 165
1:N Protection 165
Protection Switching Characteristics 166
Protection in Ring Networks 167
Two-Fiber UPSR 167
Two-Fiber Bidirectional Line-Switched Rings (2-Fiber BLSRs) 169
Four-Fiber Bidirectional Line-Switched Rings (4-Fiber BLSR) 171
Protection Ring Comparison 172
MPLS Restoration 173
Prenegotiated versus Dynamic Protection 173
End-to-End versus Local Restoration 174
Protection Switching Options 174
1 + 1 Protection 174
1:N and N:M Protection 174
Network Survivability Design 175
Survivability Mechanism Categories 175
Dedicated Protection 175
Shared Protection 176

Restoration 176
Restoration Time 176
Survivability Mechanism Comparison 178
Multilayer Survivability 180
Survivability Design Trends for Optical Networks 180
Eliminating SONET/SDH and ATM 180
Putting Intelligence into the Optical Core 181
Summary 182
Recommended Reading 183
Optical Transmission Technologies 183
Optical Transmission Systems 183
Data Transmission Technologies 184
Packet over SONET/SDH 184
Dynamic Packet Transport (DPT) 184
MPLS Traffic Engineering (MPLS-TE) 184
Network Survivability Principles 185

Chapter 4

Existing and Future Optical Control Planes 187

Static IP Optical Overlay Control Plane 189
Static IP Optical Overlay Model 189
IP Service Infrastructure 189
IP/Optical Adaptation 191
Two-Layer Architecture Implementation 192
Multiwavelength Transmission 194
Bandwidth Management 197
Static Optical Control Plane 198
IP Service Infrastructure Implementation 198
Small Point-to-Point Backbones or Backbone Interconnections 199

Large-Scale Hierarchical Backbones 200
POP Aggregation 204
Switched POP Aggregation 204
Point-to-Point POP Aggregation 206
DPT Fiber Ring POP Aggregation 207
Optical Transport Infrastructure Implementation 208
DWDM System Building Blocks 208
Transponders 210
WDM Network Design 211
Topology 211
Channel Spacing 212
Maximum Line Rate 213
Optical Power Budget 213
Restoration 215
Optical Protection 216
Optical Line Protection 216
Optical Channel Protection 218
Optical Multiplex Section Protection 219
Service Layer Restoration 221
Automatic Protection Switching 221
Intelligent Protection Switching 222
IP Restoration 223
Application-Based Restoration 223
Layer 3 Route Convergence 224
Load balancing 224
MPLS-TE Fast Reroute 225
Dynamic IP Optical Overlay Control Plane 226
Wavelength Routing Overlay Model 227
Architecture and Elements 228
Wavelength Routing Control Plane 231
Lightpath Provisioning 234
Lightpath Attributes 236
Centralized Lightpath Routing 236

Distributed Lightpath Routing 237
Optical UNI Signaling 239
Wavelength Conversion 240
Restoration 240
Protection Strategy 241
Prenegotiated Protection 241
On-Demand Protection 243
QoS Levels 243
Advanced Protection Concepts 243
Local Protection 243
Hierarchical Protection 243
Integrated IP Optical Peer Control Plane 244
Applicability Considerations 245
Advantages 245
Comparing WRs and LSRs 245
Comparing LSPs and Lightpaths 246
MPLmS Model 247
Architecture and Elements 248
MPLmS Control Plane 249
Lightpath Provisioning 252
Optical Routing 252
Topology Discovery 253
Optical Link State Advertisement 254
Optical LSA Objects 254
Constrained-Based Routing 256
Lightpath (Optical LSP) Signaling 257
Concept of Nested LSPs 257
Optical LSP Setup 258
RSVP Setup Procedure 258
Optical LSP Setup in an OTN without Wavelength Conversion 26
LSP Aggregation 261

Using MPLmS in the Overlay Model 262
Restoration 263
Summary 263
Recommended Reading 264
Static IP Optical Overlay Control Plane 264
Dynamic IP Optical Overlay Control Plane 265
Integrated IP Optical Peer Control Plane 265

Chapter 5

Optical Networking Applications and Case Examples 267

Optical End-to-End Networking Design Trends 267
Service POP 268
Metro Solutions 272
IP and Optical Metro Evolution 273
Core Solutions 276
IP and Optical Core Evolution 277
Conclusions 278
Case Example A: Next-Generation Storage Networks 279
Application Requirements 283
Common Solutions 283
Case Example B: Next-Generation Internet Service Provider 285
Application Requirements 285
Common Solutions 286
Case Example C: Next-Generation Carrier 293
Application Requirements 294
Common Solutions 294
Summary 300
Recommended Reading 301

Glossary 303

Notes 311

Index 317

1

Introduction to Carrier Network Architectures

In the first chapter, we give an overview about existing and possible future Carrier Network Architectures. We will explain the traditional architectural elements and describe the basic concepts and requirements of these architectures in order to put each of the different elements and its functionality into context. We do not go deep into the details of all used technologies, but rather point out the most important basics to provide the reader with just the necessary level of knowledge to understand the more detailed discussions of next-generation optical networks throughout the rest of the book.

IP as Unifying Protocol Infrastructure

The explosive growth of Internet/intranet traffic is making its mark on the existing transport infrastructure. An unprecedented shift has occurred in traffic content, pattern, and behavior. It has transformed the design of multiservice networks and created a commercial demand for Internet Protocol (IP) networks that operate in excess of 1 Gigabit per second. A change in usage patterns from connection-oriented, fixed configured services to dynamic, connectionless IP services is currently underway. According to several studies, telephone company revenue will grow significantly, with data services—particularly IP—accounting for most of this increase. For public carriers, IP is critical for future revenue growth.

Exponential growth in IP traffic volumes is associated with continuing growth of Internet and enterprise intranet network usage, rapid expansion of enhanced IP services based on Voice over IP (VoIP) and multicasting capabilities, and cost-effective, high-bandwidth residential connectivity via Digital Subscriber Line (DSL), cable modem, and wireless technologies.

From the networking perspective, IP is the only protocol that runs over and all transport technologies. Even end-to-end solutions based on technologies such as Asynchronous Transfer Mode (ATM) rely on systems and applications to map IP traffic onto backbone circuits. This forms a common and standardized interface between services and the technologies that are used to deliver the services. Such a standardized enables operators to adapt quickly to rapidly changing markets, the introduction of new technologies, and increasing competition.

An intelligent network is one that is able to separate services from technologies. Such an intelligent network will recognize individual users and applications, authenticate them as valid users of the network, enable them to select services, and deliver the appropriate level of performance to the applications.

Considering that the service providers' business of today is a mix of different services, such as dial access and Frame Relay, ATM has to be integrated. One way is to form an IP + ATM network infrastructure. An IP+ATM network infrastructure combines the application layer visibility of IP and the traffic management capabilities of ATM, all on a single platform. This enables service providers to provision services such as dial access and Frame Relay, as well as next-generation VoIP and integrated access services from a unified network infrastructure without compromising Quality of Service (QoS).

The Traditional Carrier Network Architecture

Service providers have been using a mix of various kinds of networking technologies for building up their national or international carrier networks (see Figure 1–1). In doing so, they had to cope with several constraints and challenges. Each networking technology introduced by the service provider handled—as many networks, still handles—the issues of one or more of these challenges

ic volumes is associated with continued strong
e intranet network usage, rapid emergence of
oice over IP (VoIP) and multicasting capabili-
dwidth residential connectivity via Digital Sub-
n, and wireless technologies.

tive, IP is the only protocol that runs over any
Even end-to-end solutions based on backbone
ous Transfer Mode (ATM) rely on external sys-
P traffic onto backbone circuits. Therefore, IP
ed interface between services and the transport
liver the services. Such a standardized interface
ly to rapidly changing markets, the introduction
sing competition.

that is able to separate services from technology
ill recognize individual users and application
s of the network, enable them to select service
l of performance to the applications.

providers' business of today is a mix of differ
d Frame Relay, ATM has to be integrated. (
network infrastructure. An IP+ATM netw
olication layer visibility of IP and the traffic n
all on a single platform. This enables service
h as dial access and Frame Relay, as well as
d access services from a unified network arc
uality of Service (QoS).

etwork Architecture

sing a mix of various kinds of networking t
ional or international carrier networks (se
cope with several constraints and challeng
duced by the service provider handled-
—the issues of one or more of these challe

Exponential growth in IP traffic volumes is associated with continued strong growth of Internet and enterprise intranet network usage, rapid emergence of enhanced IP services based on Voice over IP (VoIP) and multicasting capabilities, and cost-effective, high-bandwidth residential connectivity via Digital Subscriber Line (DSL), cable modem, and wireless technologies.

From the networking perspective, IP is the only protocol that runs over any and all transport technologies. Even end-to-end solutions based on backbone technologies such as Asynchronous Transfer Mode (ATM) rely on external systems and applications to map IP traffic onto backbone circuits. Therefore, IP forms a common and standardized interface between services and the transport technologies that are used to deliver the services. Such a standardized interface enables operators to adapt quickly to rapidly changing markets, the introduction of new technologies, and increasing competition.

An intelligent network is one that is able to separate services from technology. Such an intelligent network will recognize individual users and applications, authenticate them as valid users of the network, enable them to select services, and deliver the appropriate level of performance to the applications.

Considering that the service providers' business of today is a mix of different services, such as dial access and Frame Relay, ATM has to be integrated. One way is to form an IP + ATM network infrastructure. An IP+ATM network infrastructure combines the application layer visibility of IP and the traffic management capabilities of ATM, all on a single platform. This enables service providers to provision services such as dial access and Frame Relay, as well as next-generation VoIP and integrated access services from a unified network architecture without compromising Quality of Service (QoS).

The Traditional Carrier Network Architecture

Service providers have been using a mix of various kinds of networking technologies for building up their national or international carrier networks (see Figure 1–1). In doing so, they had to cope with several constraints and challenges. Each networking technology introduced by the service provider handled—and many networks, still handles—the issues of one or more of these challenges.

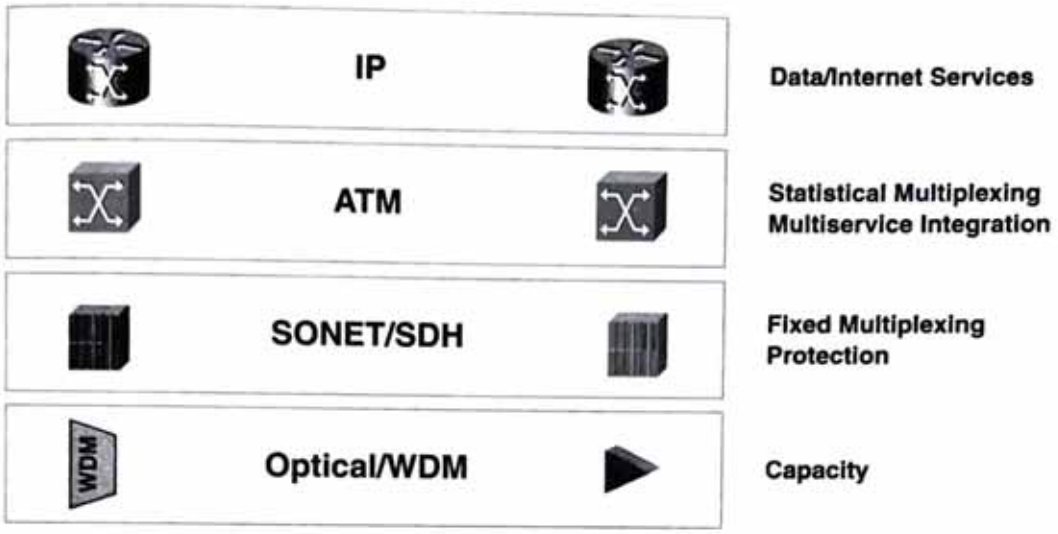


Figure 1–1 The traditional network architecture consists of multiple layers

As a consequence, a traditional service provider network architecture is built of multiple layers. The optical/Wavelength Division Multiplexing (WDM) layer forms the physical transport medium providing sheer bandwidth; in the past, this layer did not have too much routing intelligence. To allocate bandwidth in a proper way, the Synchronous Optical Network (SONET)/Synchronous Digital Hierarchy (SDH) layer is used in many traditional networks. It offers mechanisms for efficient bandwidth utilization plus intelligent protection mechanisms but does not allow intelligent routing. The ATM layer above gives additional possibilities for statistical multiplexing while allowing multiservice integration at the same time. This basically enhances efficient utilization of the layers below (the SONET/SDH and the optical/WDM layer). ATM also defines routing mechanisms intended to optimize traffic delivery throughout the network in terms of the different ATM service offerings.

If service providers do not offer any IP services, the three-layer infrastructure described so far is more than sufficient. Because IP-based applications enforced IP service offerings, several mechanisms evolved integrating IP with the infrastructure described so far, ranging from pure overlay methods to very advanced integration models such as Multiprotocol Label Switching (MPLS).

This layered network architecture grew and changed as the services provided and sold by the providers changed and networking technology evolved. To understand this dramatic change in the architecture of today's networks fully,

the functions and features of each of these layers will now be discussed in this chapter in more detail.

Achieving Long-Haul Connectivity with Optical Transmission

Optical fiber as a transport medium offers a tremendous amount of capacity. We will look into optical transmission details in greater detail in the following chapters, but for the moment, we will try to keep it simple.

Because of the fibers' transmission capabilities, a modulated optical signal, which is usually inserted by a laser, can be transmitted several kilometers before it has to be recovered. With technologies available today (lasers as well as fiber), the signal has to be recovered every 40 to 80 km. The typical bit rates achievable with advanced modulation techniques today allow data to be sent at 10 Gbps per wavelength (optical frequency) and can be expected to go even higher as technology evolves. If multiple frequencies are used in parallel over one fiber (which generally is called ***WDM***), the transmission capacity per fiber can be extended today into the Tb range.

By deploying these technologies, service providers can easily build optical networks with tremendous transmission capacity, ranging from several Gbps to Tbps, forming the basis for the current and next-generation backbone carrier networks.

Delivering Multiplexed Services

In the past, service providers typically delivered telephony services to their customers. For this, they used nonsynchronous hierarchies called ***Plesiochronous Digital Hierarchies*** (PDHs) to carry the low-bit-rate signals representing either the customer's voice or data connections. PDH networks have several limitations. To begin with, the highest multiplexing rate is limited to 139.264 Mbps. In addition, the transport efficiency is low because there is a high amount of fill capacity required for accommodating delays between the nonsynchronous clocks of the network equipment. Furthermore, interoperability between different vendors is difficult because only multiplexing—not the transmission—is defined by the standards.

Synchronous Transmission Standards

SONET and SDH overcome the drawbacks of PDH networks and enable service providers also to provision high-bit-rate connections above 155 Mbps. These connections are typically required for today's increasing data transport needs. SONET and SDH allowed service providers to build a Time Division Multiplexing (TDM) network on top of their physical fiber plant. Although SONET is the North American synchronous TDM standard, SDH is the TDM standard commonly used in Europe and Japan. Because SDH can be seen as a global standard (and SONET being a subset of SDH), interoperability at certain levels is ensured, as will be outlined in more detail later. In addition to the summary in this chapter, the Nortel whitepaper "SONET 101" [NORT-1], the Bellcore standard GR-253 [BELL-1], and the International Telecommunication Union (ITU) T recommendation G.707 [ITU-3] can be consulted for all the details on SONET/SDH.

Both SONET and SDH define a digital multiplexing hierarchy and should ensure compatibility of equipment and implement **synchronous networking**. The basic functionality is that client signals of different service types, such as E0, E1, DS0, T1, ATM, and others, are mapped into appropriate payloads that are then multiplexed into synchronous optical signals.

Both SONET and SDH accommodate nonsynchronous TDM hierarchies. SONET includes the North American hierarchy, which is based on the DS1 signal, combining 24 DS0s (56-Kbps channels) into one 1.54-Mb stream. SDH integrates the European hierarchy, which is based on the E1 signal, combining 32 E0s (64-Kbps channels) into one E1 at a speed of 2.048 Mbps, as can be seen from Table 1–1.

Table 1–1 The North American and European TDM Hierarchies

NORTH AMERICAN RATE			EUROPEAN RATE		
SIGNAL	SPEED	CHANNELS	SIGNAL	SPEED	CHANNELS
DS0	64 Kbps	1 DS-0	E0	64 Kbps	1E-0
DS1	1.54 Mbps	24 DS-0s	E1	2.048 Mbps	32 E-0s
DS2	6.3 Mbps	96 DS-0s	E2	8.45 Mbps	128 E-0s
DS3	44.8 Mbps	28 DS-1s	E3	34 Mbps	16 E-1s
Not defined			E4	140 Mbps	64 E-1s

SONET, which was developed first, specifies a basic transmission rate of 51.84 Mbps, called the *synchronous transport signal 1* (STS-1). The equivalent optical signal is called ***OC-1*** (optical carrier 1). To ensure that both SONET and SDH match into a common multiplexing hierarchy, SDH defines its base level, the synchronous transport module 1 (STM-1), at 155.52 Mbps, which is three times the SONET base level. The optical line rates of both hierarchies are shown in Table 1–2.

Table 1–2 Interface Rates of the SONET/SDH Multiplexing Hierarchy

SONET	BIT RATE	SDH
STS-1 / OC-1	51.84 Mbps	-
STS-3 / OC-3	155.52 Mbps	STM-1
STS-12 / OC-12	622.08 Mbps	STM-4
STS-24 / OC-24	1244.16 Mbps	-
STS-48 / OC-48	2488.32 Mbps	STM-16
STS-192 / OC-192	9953.28 Mbps	STM-64

Building a Complex SONET/SDH Carrier Network

A typical SONET/SDH network basically consists of four different network elements:

1. Add/Drop Multiplexer (ADM)
2. Terminal Multiplexer (TM)
3. Digital Cross-Connect (DXC)
4. Regenerator

All of these elements are interconnected using the service provider's fiber plant and form a typical SONET/SDH network, as shown in Figure 1–2.

ADMs used for ring networks and TMs used in linear topologies may be interconnected directly with dark fiber. If the distance between two multiplexers exceeds approximately 40 km, a regenerator must be placed in between the multiplexers. The regenerator ensures proper transmission through regenerating the optical signal, which has been degraded during optical transmission across the fiber.

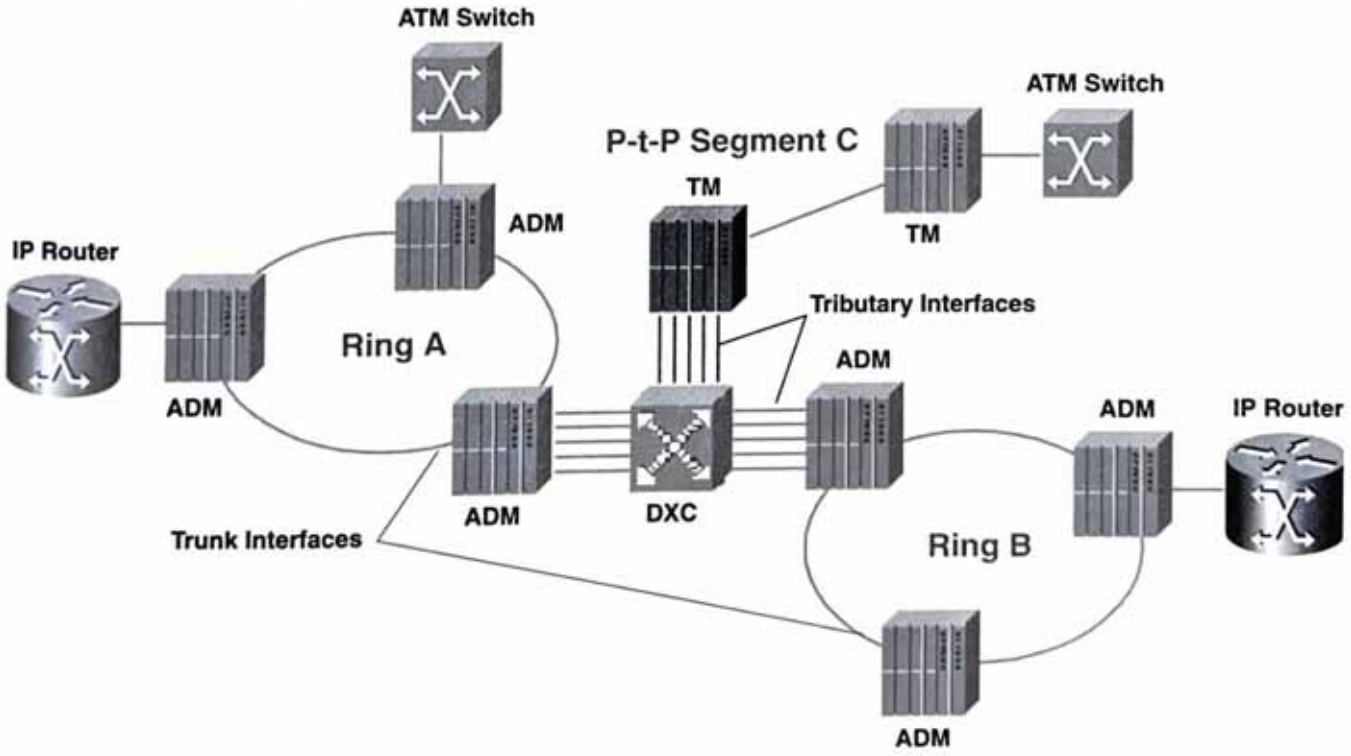


Figure 1–2 A typical SONET/SDH network consists of multiple multiplexers forming rings or linear links, which are interconnected with digital cross-connects

Multiplexers are equipped with two types of interfaces, trunk and tributary interfaces. ***Trunk interfaces*** are used to interconnect multiplexers. Trunk interfaces can range from OC-3/STM-1 to OC-192/STM-64, as specified in the current SONET/SDH recommendations. ***Tributary interfaces*** are used to attach client equipment, such as IP routers, ATM switches, or telephone switches, to the multiplexers. The range of available tributary interfaces for a multiplexer typically starts at the hierarchy level of the multiplexer's trunk interface and goes down to DS-0 or E1 interfaces.

To be able to switch a connection from a ring network segment to another or onto a point-to-point segment, DXCs are used. For example, in Figure 1-2, a DXC is placed between ring A, ring B, and the linear segment C. One multiplexer from each segment is attached to the DXC, using multiple tributary interfaces. Depending on what level of the TDM hierarchy connections have to be cross-connected, a certain type of DXC and the appropriate tributaries have to be used. If connections at DS1/E1 level are to be cross-connected, E1/DS1 tributaries must be used to attach the multiplexers to the DXC, and a DXC capable of switching at DS1/E1 level must be selected.

A network reference model is specified by the SONET/SDH standards to define a proper structure for all of the mechanisms in this complex network scenario. Three layers are defined (Figure 1–3). Each network element may be part of only one or more of these layers, according to its tasks and its requirements.

The lowest layer, called ***section layer*** (within SONET) or ***regenerator section layer*** (within SDH), incorporates optical transmission and signal regeneration. All network elements are part of this layer. Moreover, regenerators are only part of this layer, because they simply regenerate the optical signals along very long-distance interconnections.

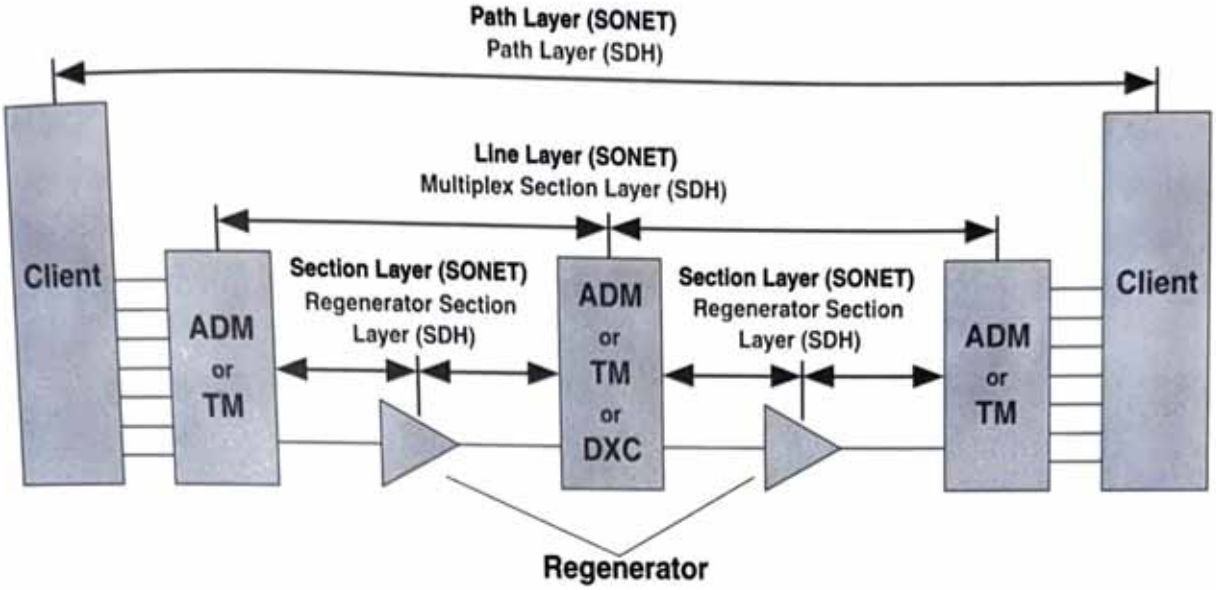


Figure 1–3 SONET/SDH network reference model introduces three functional layers

The second layer is called the ***line layer*** (SONET) or ***multiplex section layer*** (SDH). At this layer, several low-speed signals are multiplexed to and demultiplexed out of the high-bit-rate signal of the trunk interfaces. Multiplexers and DXCs are part of this layer.

The third layer is responsible for end-to-end connection delivery and is called the ***path layer*** (SONET and SDH). The client equipment terminating the endpoints of connections is part of this layer.

Certain control information is required at each layer for signaling, performance monitoring, or protection switching. This control information is carried in the ***overhead*** of each layer. Thus, a multiplexer is adding some path overhead (POH) to the signal coming from the tributary interface. It multiplexes multiple signals together and adds the section overhead (SONET) or multiplex section overhead (SDH). Finally, before sending it out to the next network element, it adds the line overhead (LOH) within SONET or the regenerator section overhead (RSOH) within SDH.

Multiplexing and Framing Standards

A wide variety of signals can be transported across a SONET/SDH infrastructure. A signal inserted or extracted through a tributary interface is mapped into an appropriate container. This container is called ***virtual tributary*** (VT) within

SONET and ***virtual container*** (VC) within SDH. In addition to the sig itself, the container also includes the POH information, ensuring end-to-e connectivity.

How signals are multiplexed and mapped is specified in the SONET/SD recommendations. The mapping and multiplexing scheme for the most impo tant signal rates is outlined in Figure 1–4. Please keep in mind that this is not complete picture of both the SONET and SDH multiplexing hierarchy.

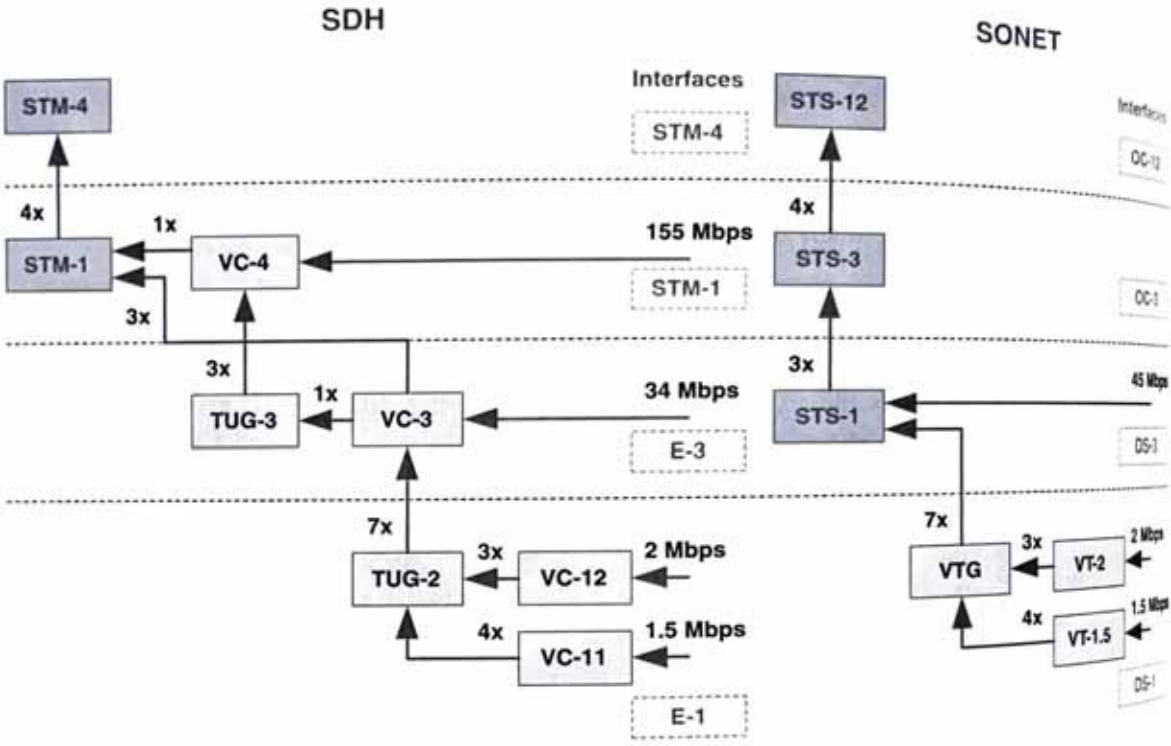


Figure 1–4 SONET and SDH multiplexing hierarchy

Within SONET, four VT1.5s and three VT2s are combined together into VT group (VTG) at the DS1 level. Seven VTGs are then combined and mapped into the basic building block the STS-1, which resides at the DS3 level. Because the first optical interface specified (the OC-1 interface) is not commonly used in practice, the first optical SONET interface to be considered is the OC-3 interface running at 155 Mbps. Thus, after mapping the VTGs into the STS-1, three STS-1s are mapped into an OC-3 interface.

Within SDH, four VC-11s and three VC12s are combined together into transport unit group (TUG-2) at the E1 level. Seven TUG-2s are then com

bined together and mapped into a VC-3 at the E3 level. As opposed to SONET, SDH also allows a high-bit-rate signal at 34 Mpbs to be mapped into a VC-3. The first specified optical interface within SDH is the STM-1 interface running at 155 Mbps (equivalent to the OC-3 interface of SONET). A VC-3 itself can be mapped in two different ways into an STM-1 interface. It can be mapped into a TUG-3, and three TUG-3s can be combined into a VC-4, which is then mapped into an STM-1 interface. The second possibility is to map three VC-3s directly into an STM-1 interface.

The basic building block of SONET is the STS-1, and the basic building block of SDH is the STM-1 frame. A high-level logical picture of the frame structure of both the STS-1 and STM-1 is shown in Figure 1–5. Keep in mind that the STS-1 frame is used for a 51.84-Mbps interface, and the STM-1 frame is used for a 155-Mpbs interface. Thus, the STM-1 frame has three times the columns of the STS-1 frame. From a functional point of view, the structures of the two frames are most likely the same.

The first columns of the frame are used for carrying the overhead of the first two SONET/SDH layers, the section overhead (SOH) and LOH, as defined in SONET, or the RSOH and the multiplex section overhead (MSOH), as defined in SDH. The remaining columns are used for the payload carrying the multiplexed containers. As already mentioned, the overhead of the third layer (the POH) is carried within each container.

The ***payload pointers***, which are part of the LOH (SONET) or MSOH (SDH), are the key function in both synchronous hierarchies. The VTs and VCs possibly may not be synchronous with the STS/STM-Frame. Ideally, SONET/SDH networks are synchronous, but in practice, there are always small clocking differences, and these have to be accommodated. As a consequence, the VTs and VCs are mapped into the STS/STM-frame at slightly varying boundaries. The payload pointers then are used to maintain a fixed relationship between the STS/STM-frame boundaries and the position of the containers mapped into the payload.

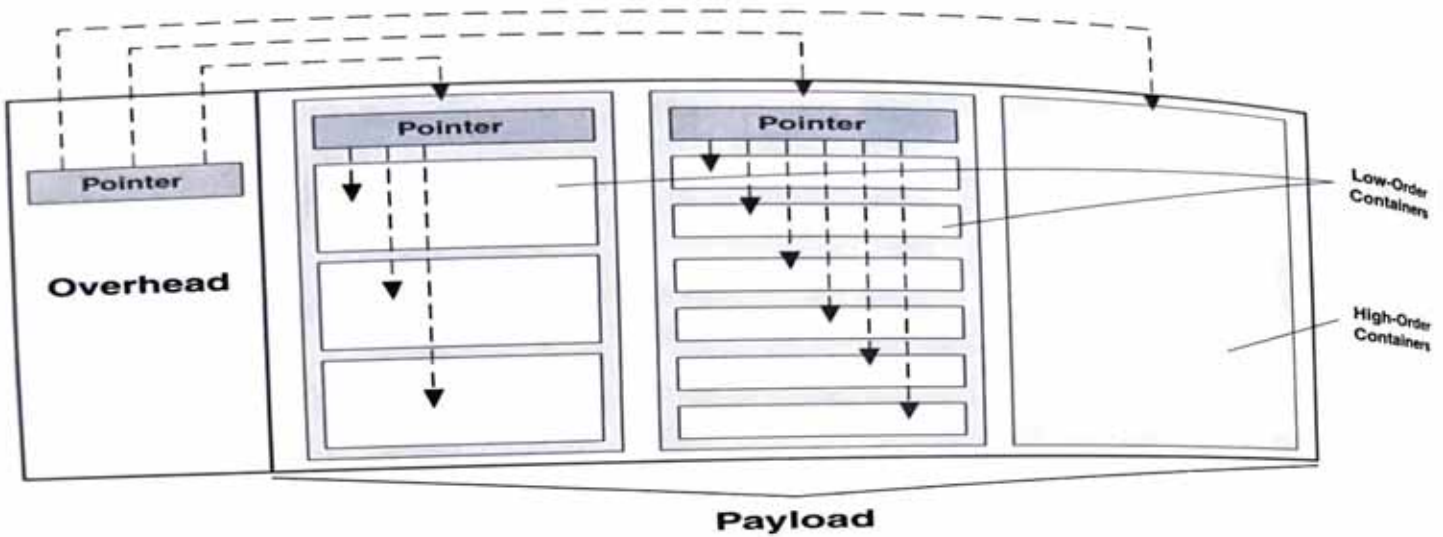


Figure 1–5 The SONET STS and SDH STM frame consist of an overhead and a payload portion

The pointers of the SOH/MSOH are referring to the position of the high-order containers in the frame. A high-order container can either carry a signal directly (a VC-3 may directly carry a 34-Mbps signal) or may carry low-order containers and have its own payload pointers, which point to the position of the low-order containers.

The second most important information carried in the SOH/MSOH is the K1/K2 overhead bytes that are used for protection switching purposes. Using the information carried in the K1/K2 bytes, a protection-signaling protocol is implemented and used by the SONET/SDH network elements to monitor the network status, detect failures, and react appropriately to restore failed traffic. Details on protection signaling and the protection architectures used in SONET/SDH networks are covered in Chapter 3, "Optical Networking Technology Fundamentals."

Much more information is carried in the LOH/RSOH and SOH/MSOH of the STS/STM-frame but the details are outside of the scope of this section. Additional information can be found in the SONET recommendation GR-253 [BELL-1] and the SDH recommendation G.707 [ITU-3].

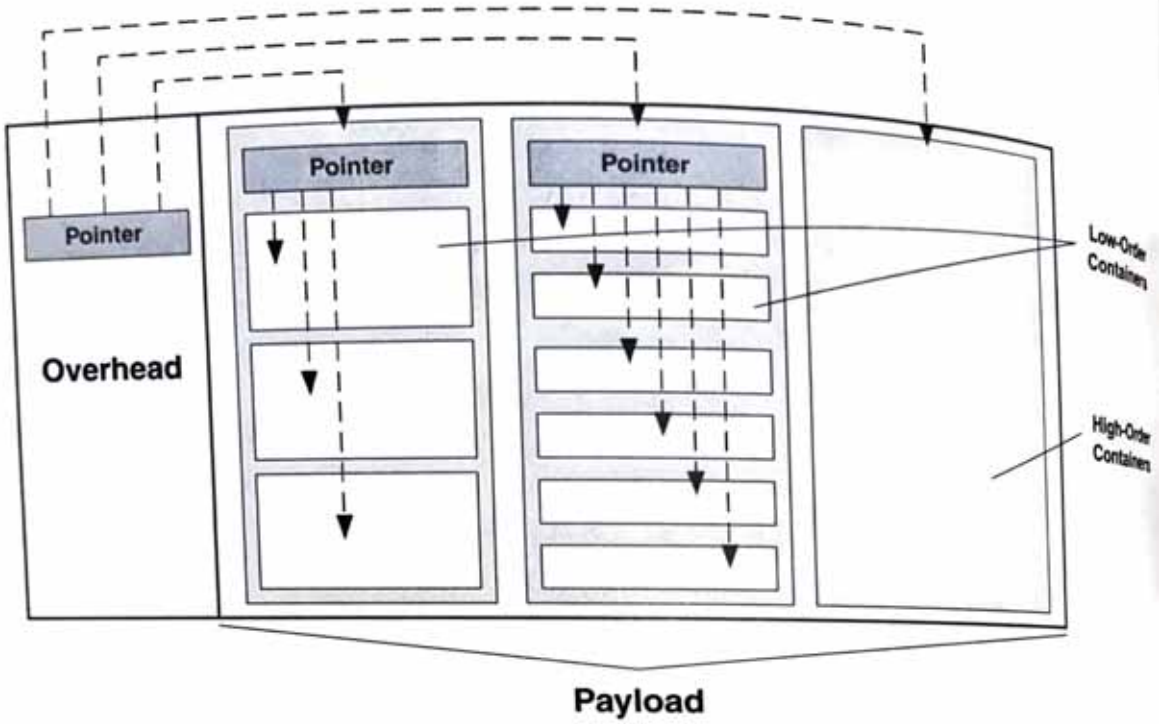


Figure 1–5 The SONET STS and SDH STM frame consist of an overhead and a payload portion

The pointers of the SOH/MSOH are referring to the position of the high-order containers in the frame. A high-order container can either carry a signal directly (a VC-3 may directly carry a 34-Mbps signal) or may carry low-order containers and have its own payload pointers, which point to the position of the low-order containers.

The second most important information carried in the SOH/MSOH is the K1/K2 overhead bytes that are used for protection switching purposes. Using the information carried in the K1/K2 bytes, a protection-signaling protocol is implemented and used by the SONET/SDH network elements to monitor the network status, detect failures, and react appropriately to restore failed traffic. Details on protection signaling and the protection architectures used in SONET/SDH networks are covered in Chapter 3, "Optical Networking Technology Fundamentals."

Much more information is carried in the LOH/RSOH and SOH/MSOH of the STS/STM-frame but the details are outside of the scope of this section. Additional information can be found in the SONET recommendation GR-253 [BELL-1] and the SDH recommendation G.707 [ITU-3].

Providing Carrier Class Network Availability

Network availability, in particular, becomes a key design criteria in the case of providing telephony services, because these services usually require a nearly 100% uptime. The network must be designed to handle fiber cuts, as well as equipment failures. SONET/SDH has been specifically designed to fulfill that requirement.

In a SONET/SDH network, there is always some capacity reserved for protection purposes. Figure 1–6 illustrates that, in a typical ring architecture, there is always a working fiber (ring) and a protection fiber (ring). Usually in SONET/SDH, the protection path is not used under normal transport conditions; only in the case of a failure is the capacity of the protection path used to reroute and restore traffic affected by the failure. This rerouting is done solely based on the SONET/SDH network availability information and does not take into account any higher layer forwarding and reachability requests, i.e., ATM or IP reachability and availability information.

In the SONET/SDH frame overhead performance monitoring and protection information is exchanged. This information is used to implement a signaling protocol for protection switching, which is then used by the SONET/SDH network elements to communicate with each other and to perform required protection switching.

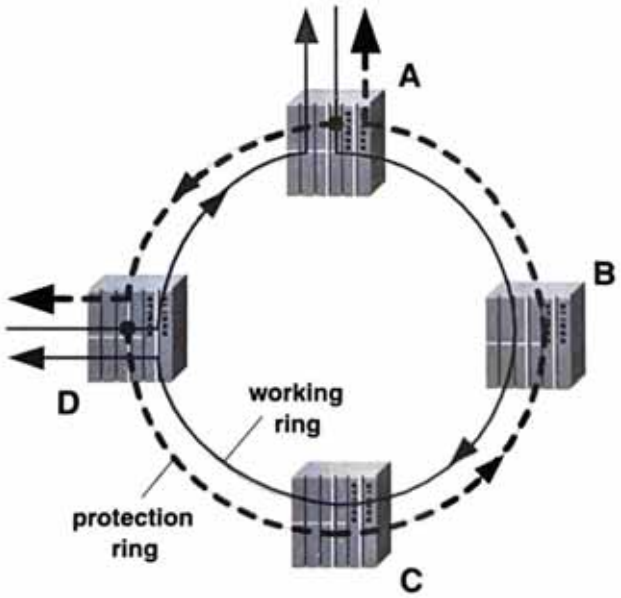


Figure 1–6 SONET/SDH provides comprehensive protection options

With the protection switching mechanisms implemented in the SONET/SDH layer, failures typically can be restored within 60 ms; thus, the upper layers will usually not realize that a failure has occurred. It is important to understand that, whenever another technology or combination of technologies is going to replace SONET/SDH, all the protection mechanisms must be guaranteed by this new implementation. Otherwise, no service provider would replace an existing, reliable part of the network. This will be emphasized in the following chapters when we deal with replacement strategies in detail.

Adding Versatility with Asynchronous Transfer Mode

SONET/SDH perfectly fulfills the needs of service providers that are mainly delivering telephony services. Old-world TDM telephony is connection-oriented and requires constant bandwidth (usually 64 Kbps per voice channel) even if no data has to be delivered (no speech has to be transferred via the connection). This is the reason why multiplexing and container allocation in SONET/SDH networks is optimized for voice. Static and fixed bandwidth is allocated permanently as soon as the desired connection is provisioned.

In the 1980s, standardization bodies driven by telcos and vendors developed a versatile layer on top of SONET/SDH to get a more flexible allocation of bandwidth, as well as more advanced service offerings. The new technology—ATM—works with synchronous slots called ***cells***, but these cells do not have to be sent only periodically, as is the case with timeslots in TDM networks. Instead, there are also defined services sending cells on demand, for example whenever new collected voice samples of a voice connection have to be transferred. This means that, during silence, no cells are sent and, thus, no bandwidth is blocked by idle cells.

Furthermore, numerous studies have shown that data services would play a more dominant part in the service portfolio of service providers. With this in mind and the fact that data traffic is more dynamic than static, it becomes obvious that static provisioning of timeslots or containers was no longer the most efficient and preferred method for carrier networks. As a consequence another important goal of ATM was to integrate voice and data traffic in a more efficient way than SONET/SDH.

An ATM network basically consists of a mesh of ATM switches interconnected by point-to-point ATM links or interfaces. The transmission or switch

ing units of ATM are small, fixed-length data packets, so-called cells, with a fixed size of 53 bytes—thus the term ***cell switching***. Because the cells have a small and fixed size, ATM switches can forward cells from one interface to another very quickly. In fact, it was one of the key design goals at the time to get a forwarding plane capable of switching information quickly (at interface speeds up to several Gbps) between different interfaces of ATM switches. Interestingly enough, in today's most advanced packet switches (routers that forward variable-length packets), the internal structure is also very often based on fixed-length cells, although these cells are, in most cases, of different size than 53-byte ATM cells.

ATM allows carriers to transport traffic via end-to-end connections across ATM networks with a predictable and very small delay. This is the key enabler for delivering voice and data services with a certain QoS in terms of available bandwidth, delay, and jitter.

ATM switches are typically interconnected with optical interfaces to the SONET/SDH hierarchy. The physical specifications have been taken from the SONET/SDH recommendations to ensure interoperability and the capability to directly connect ATM switches to SONET/SDH network elements.

ATM Reference Model

The ATM architecture uses a logical model to describe the functionality it supports. ATM functionality corresponds to the physical layer and part of the data link layer of the Open Systems Interconnection (OSI) reference model.

The ATM reference model is composed of three planes that span all layers. The ***control plane*** is responsible for generating and managing signaling requests. The ***user plane*** is responsible for managing the transfer of data. Finally, there is the ***management plane,*** which consists of two components:

1. Layer management: This component manages layer-specific functions, such as detection of failures and protocol problems.
2. Plane management: This component manages and coordinates functions related to the complete system.

The ATM reference model is composed of the following ATM layers:

- The ***physical layer*** is analogous to the physical layer of the OSI ref... model. The ATM physical layer manages the medium-dependent ... mission.
- The ***ATM layer*** is tightly aligned with the ATM adaptation ... (AAL). The ATM layer is roughly analogous to the data link layer ... OSI reference model. The ATM layer is responsible for estab... connections and passing cells through the ATM network. To do ... uses information in the header of each ATM cell.
- The ***AAL***, which is strongly aligned with the ATM layer, is also ... analogous to the data link layer of the OSI model. The AAL is ... ble for isolating higher layer protocols from the details of the ATM ... cesses.
- Finally, the higher layers residing above the AAL accept user packe... data and hand it to the AAL. Figure 1–7 shows the ATM ref... model.

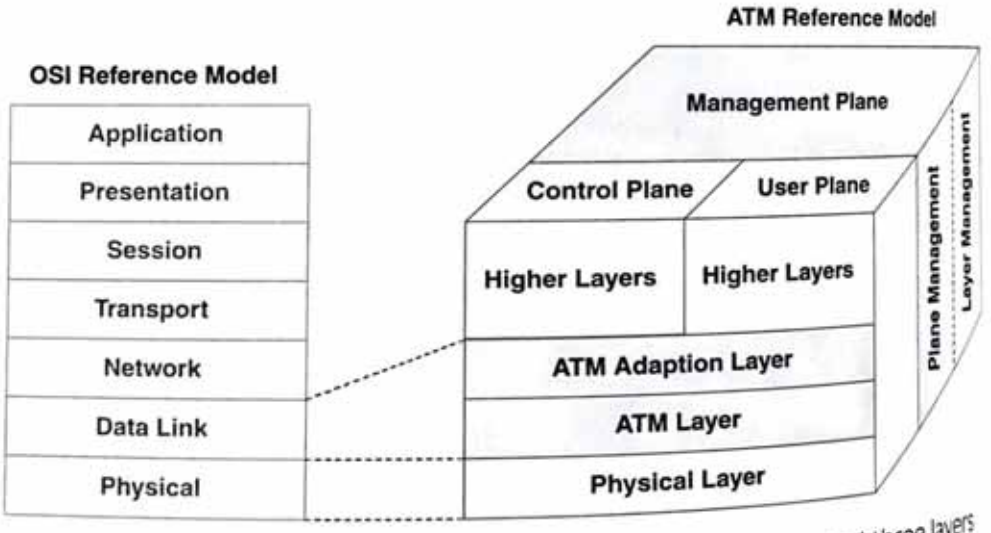


Figure 1–7 The ATM reference model includes three planes and three layers

Introducing Statistical Multiplexing

Instead of supporting pure TDM, the interfaces of ATM switches can ... payload of the SONET/SDH frame as one big container. This option is ... ***virtual concatenation*** within SONET/SDH. For example, a 622-Mbps ... interface matches the transmission rate of an STM-4 interface in the ...

world. The payload is used as one big container. That means that the four VC-4s are glued together; the big container is called a ***VC-4c.***

To be able to distinguish between a SONET/SDH interface in standard TDM mode or in concatenated mode, a "c" is appended to indicate concatenation. For example, an OC-3c is a 155-Mbps SONET interface that can be used with ATM switches. An OC-3 SONET interface cannot be used with ATM switches but is used instead for standard SONET applications. By using the entire interface bandwidth as one big transport unit, the ATM cells can be dynamically mapped into it. This method is also referred to as ***statistical multiplexing*** and is much more efficient for transporting dynamic and elastic data traffic than is fixed TDM.

On top of the SONET/SDH network, a meshed ATM network can now be set up (Figure 1–8). The ATM switches are interconnected by concatenated SONET/SDH interfaces, for example, STS-12c/STM-4c or others.

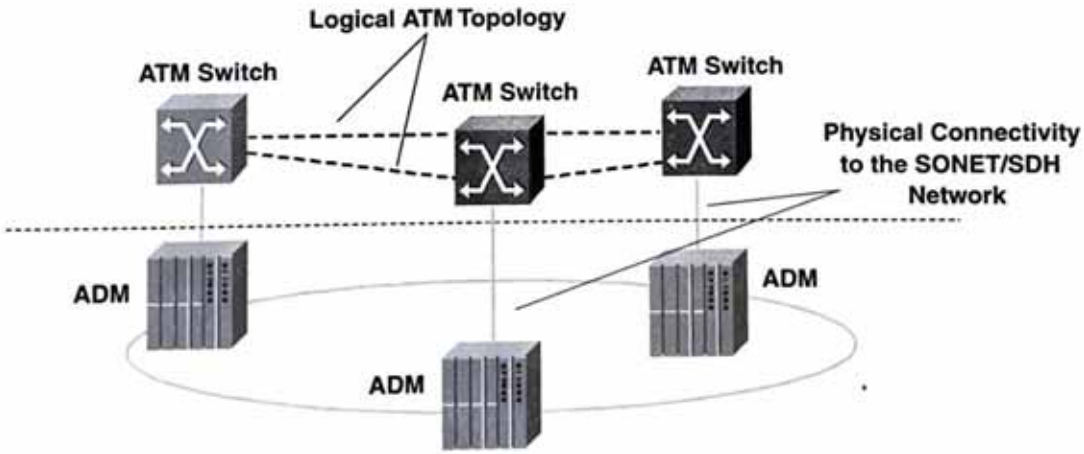


Figure 1–8 The ATM network is built on top of the SONET/SDH network to introduce statistical multiplexing

Transport Multiple Services across a Single Network

The ATM adaption layer (AAL) is responsible for transforming the client signal into ATM cells. To be able to transport different types of traffic, such as voice, mission-critical data, and Internet traffic, different types of AAL have been defined.

The AAL translates between larger service data units (SDUs; for example, video streams and data packets) of upper-layer processes and ATM cells. Specif-

ically, the AAL receives packets from upper-level protocols and breaks into the 48-byte segments that form the Payload field of the ATM cell.

Several AALs are specified to support a wide range of traffic requirements. Table 1–3 shows the AAL types and their characteristics.

Table 1–3 ATM Incorporates Four Different Types of ATM Adaptation Layers with Certain Characteristics

CHARACTERISTICS	AAL1	AAL2	AAL3/4	AAL5
Requires timing between source and destination	Yes	Yes	No	No
Bit rate	Constant	Variable	Available	Available
Connection mode	Connection-oriented	Connection-oriented	Connection-oriented or connectionless	Connection-oriented
Traffic types	Voice and circuit emulation	Packet voice and video	Data	Data

The AAL is again broken down into two sublayers, the ***convergence sublayer (CS)*** and the ***segmentation and reassembly sublayer (SAR)***. The CS adapts mation into multiples of octets. Padding can be added to achieve this The SAR segments information into 48-octet units at the source and bles them at the destination.

AAL1

AAL1, a connection-oriented service, is suitable for handling circuit applications, such as voice and video conferencing. Circuit emulation also allow the attachment of equipment currently using leased lines to backbone network. AAL1 requires timing synchronization between the and destination. For this reason, AAL1 depends on a media that supports ing, such as SONET. The AAL1 process prepares a cell for transmission following way, also shown in Figure 1–9.

Synchronous samples (for example, 1 byte of data at a sampling rate of 125 μseconds) are inserted in the Payload field. Sequence Number (SN) and Sequence Number Protection (SNP) fields are added to provide information that the receiving AAL1 uses to verify that it has received cells in the correct order. The remainder of the Payload field is filled with enough single bytes to equal 48 bytes.

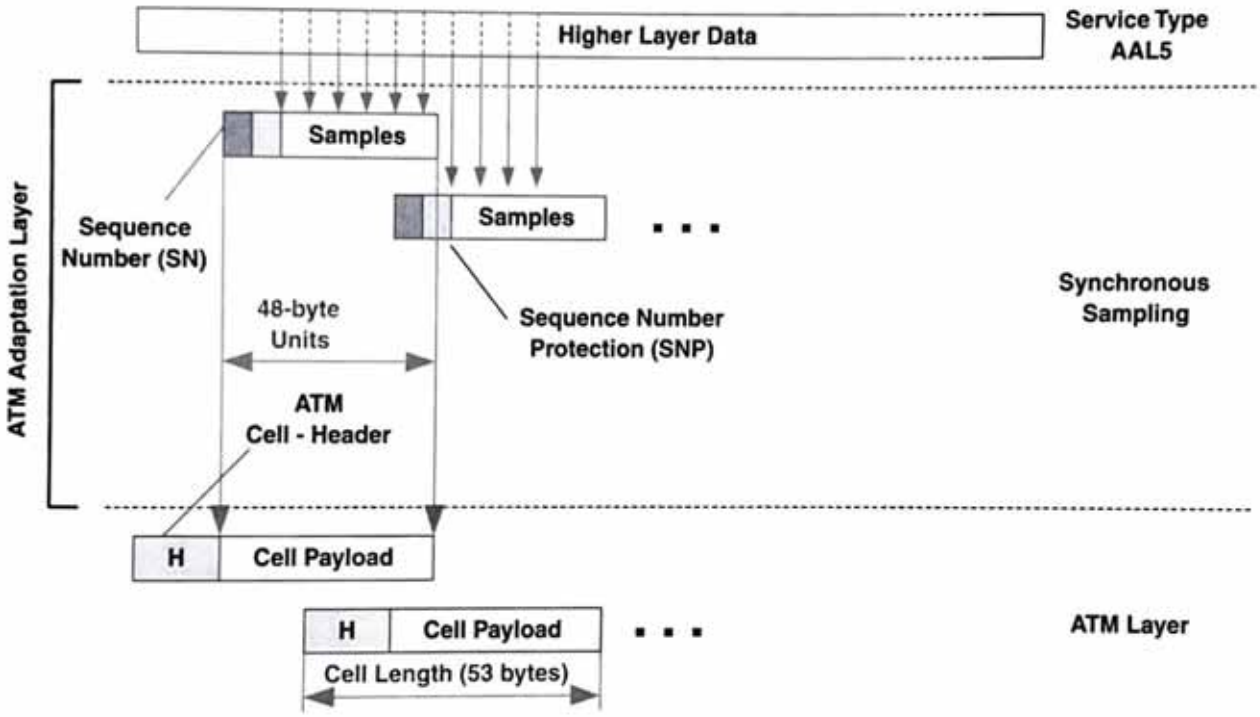


Figure 1–9 AAL1 maps synchronous samples into the 53-byte big ATM cells

AAL2

AAL2 is optimized for the transport of variable-bitrate, time-dependent traffic over ATM. As with AAL1, the SAR provides for a sequence number, but in addition, a given cell will contain a designator as to whether it is the beginning, middle, or end of a higher layer information frame.

Appended to the SAR SDU is a length indication (LI) and a cyclic redundancy check (CRC). The resultant SAR-protocol data unit (PDU) is delivered to the ATM layer, where it forms the payload (ATM-SDU) of the ATM cell (ATM-PDU). An AAL2 cell can also be partially filled; in this case, unused bytes are padded.

The most important use for AAL2 is transfer of video, although AAL5 got a lot more momentum for this application. With some vendors implementing

variable bit rate (VBR) over AAL1 for variable-rate voice, AAL2 may no longer be needed or used.

AAL3/4

AAL3/4 supports both connection-oriented and connectionless data. It was designed for network service providers and is closely aligned with Switched Multimegabit Data Service (SMDS). AAL3/4 will be used to transmit SMDS packets over an ATM network. AAL3/4 prepares a cell for transmission in the following way, also shown in Figure 1–10.

The CS creates a PDU by prepending a Beginning/End Tag header to the frame and appending a Length field as a trailer.

The SAR fragments the PDU and appends a header to it. The SAR sublayer also appends a CRC-10 trailer to each PDU fragment for error control. The completed SAR PDU becomes the Payload field of an ATM cell, to which the ATM layer appends the standard ATM header.

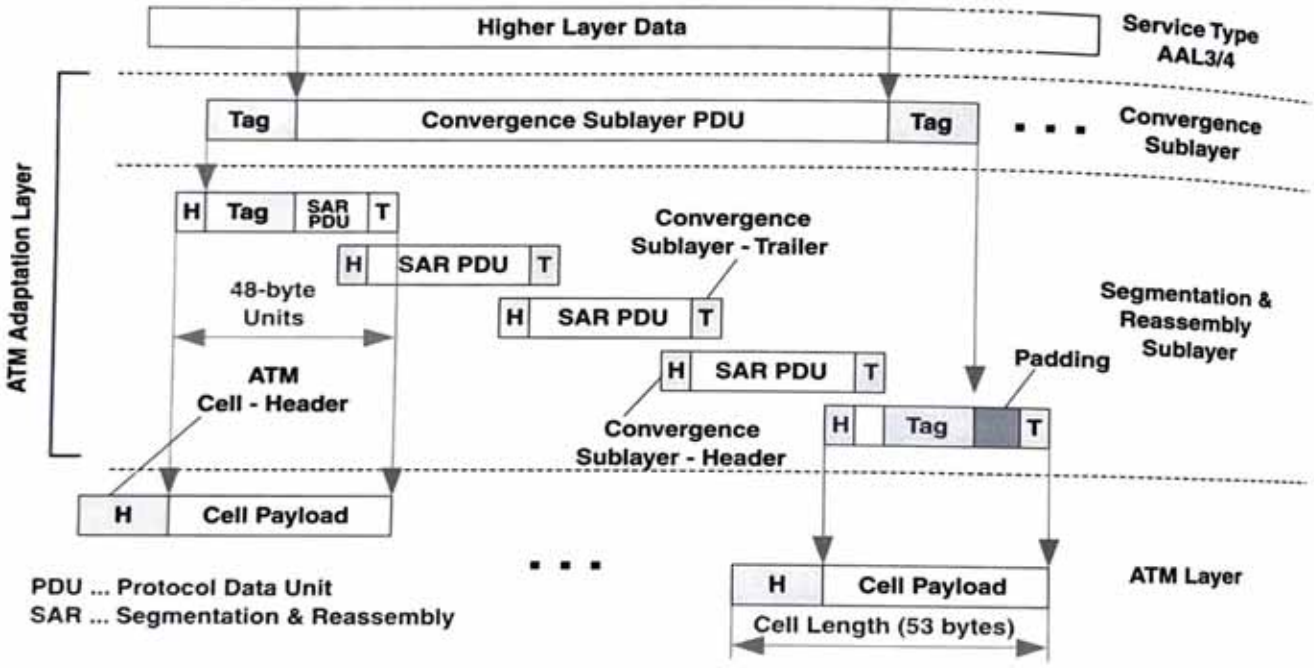
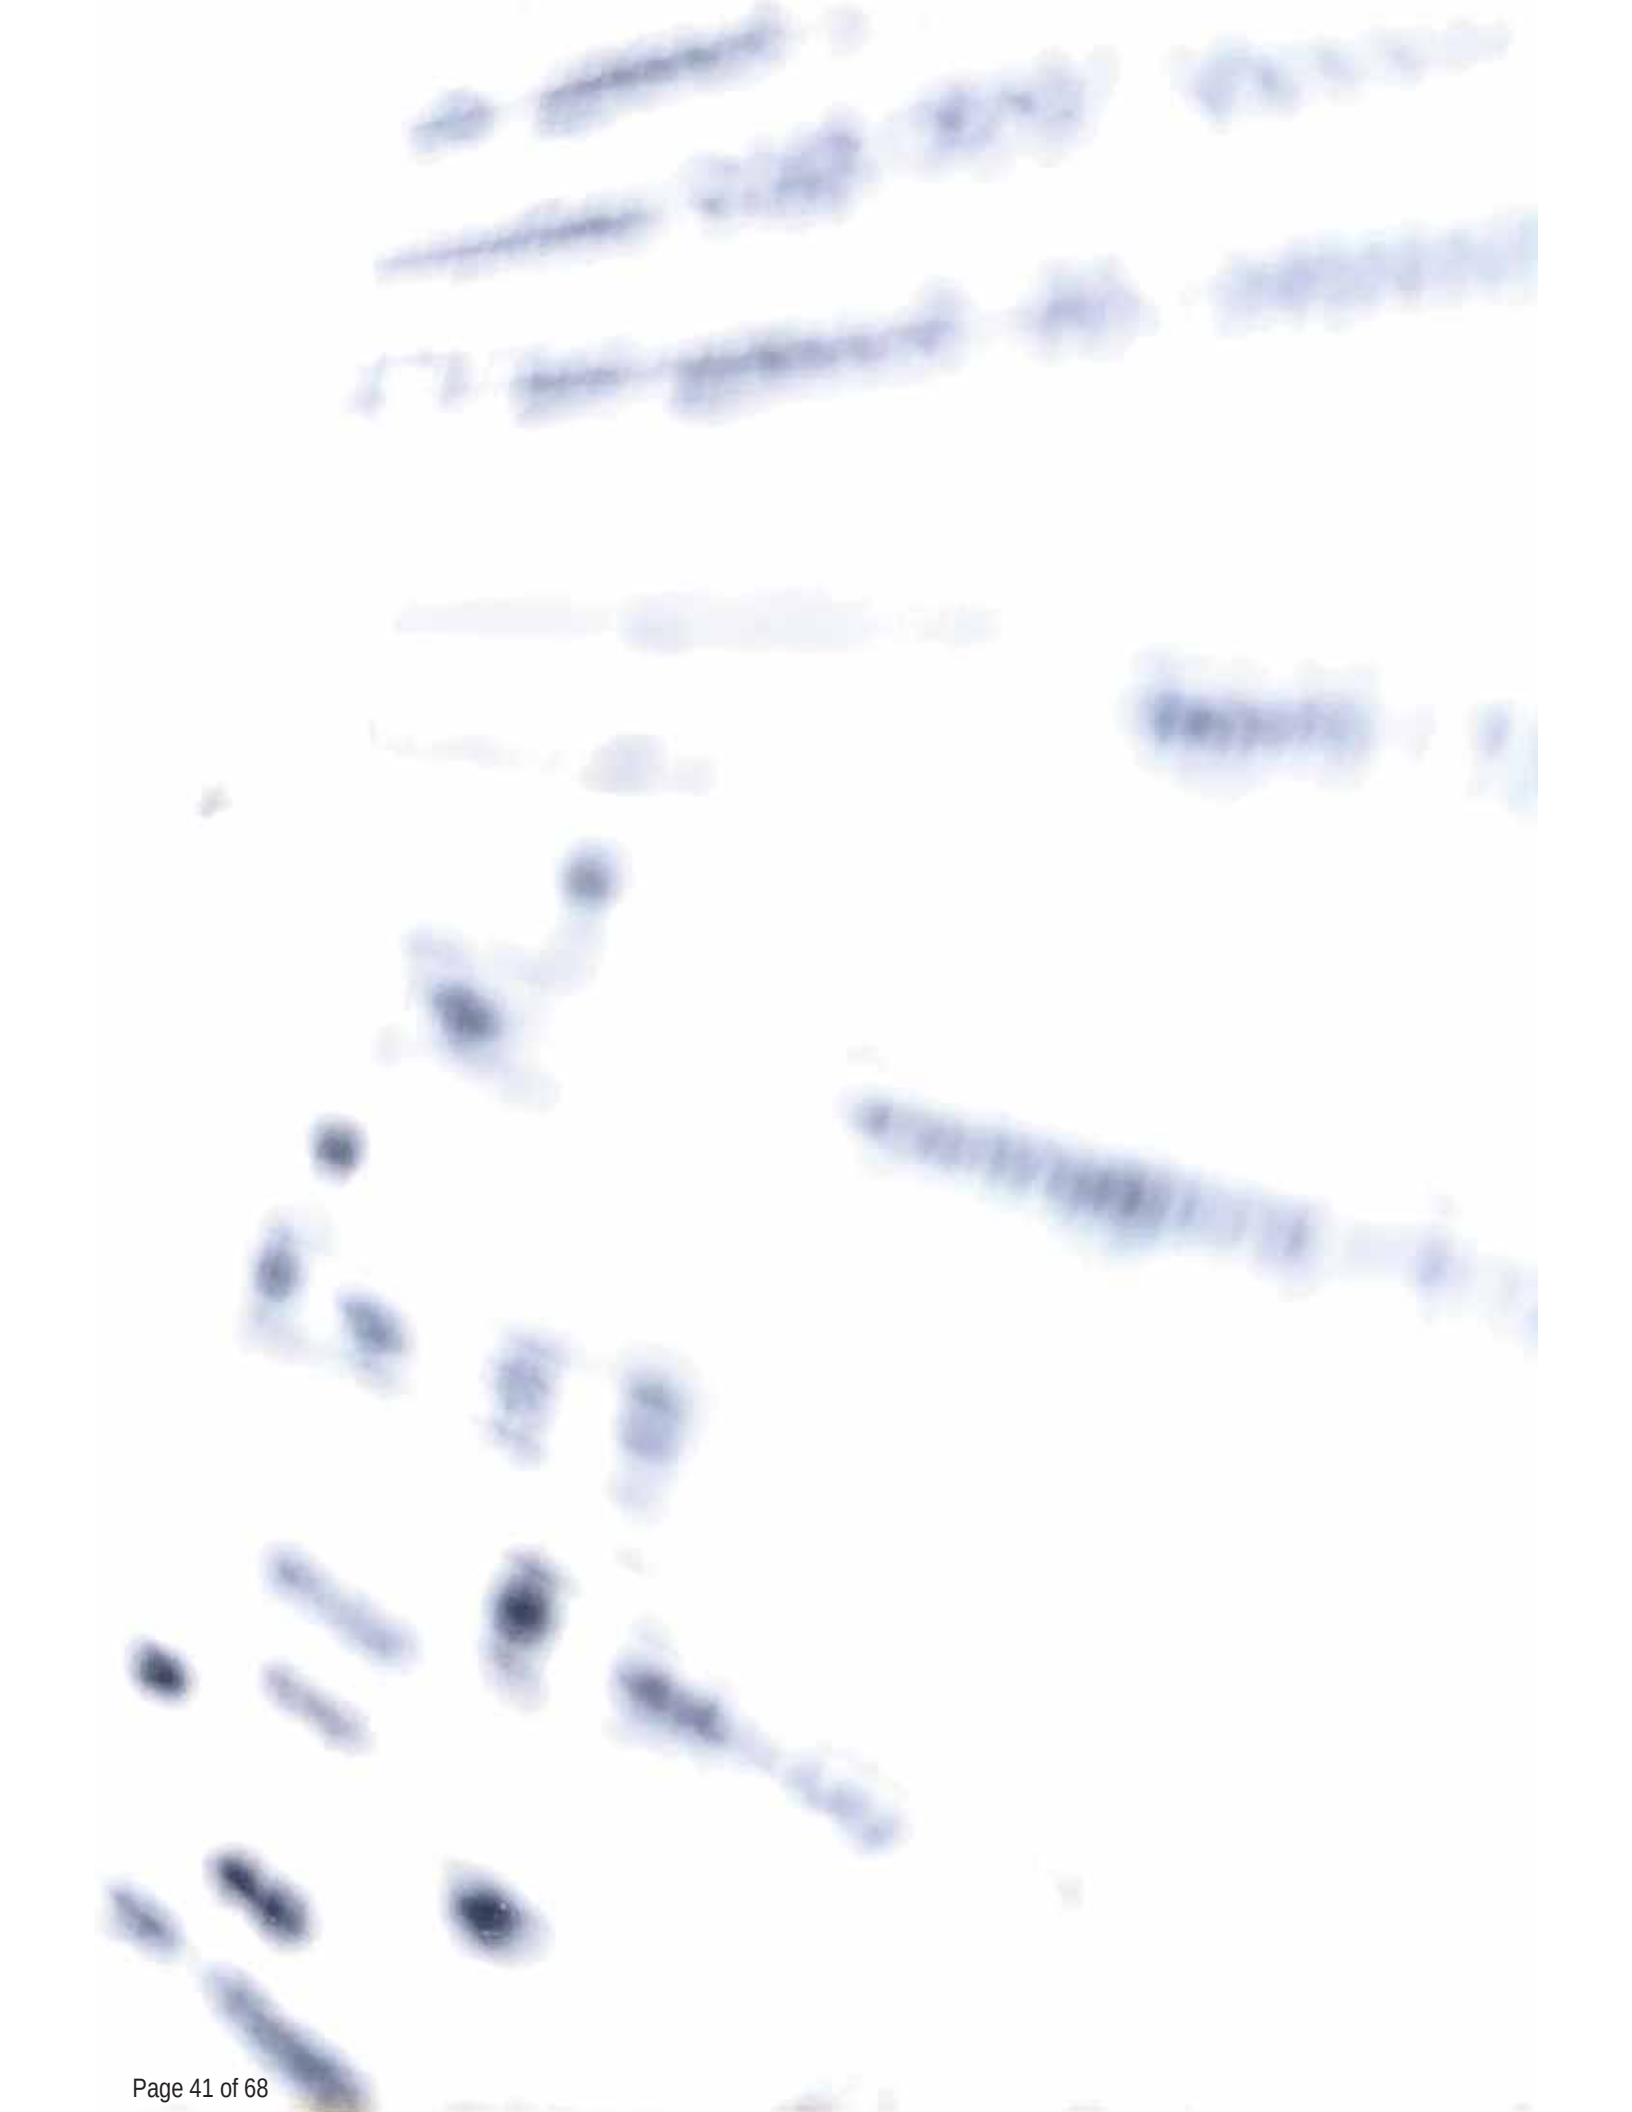


Figure 1–10 The AAL 3/4 incorporates two sublayers, the CS and the SAR

AAL5

AAL5 is the primary AAL for data and supports both connection-oriented and connectionless data. It is used to transfer most non-SMDS data, such as Classi-



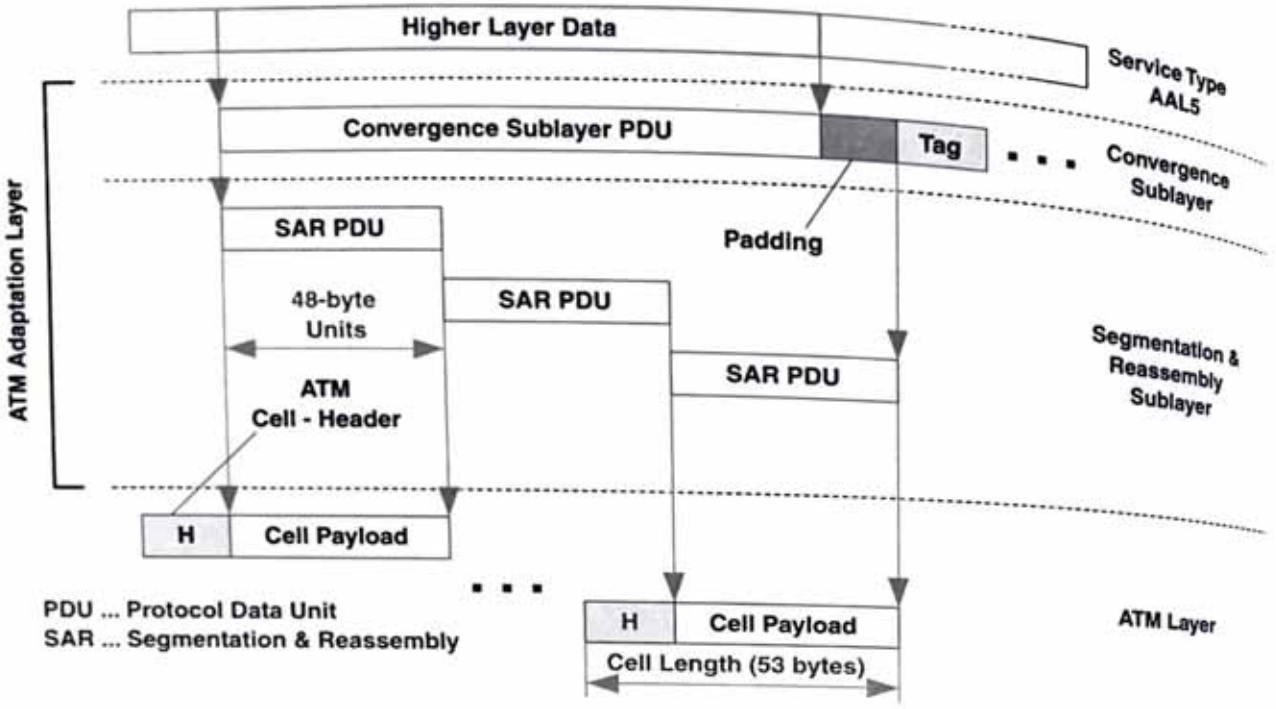


Figure 1–11 The AAL5 also incorporates two sublayers, the CS and the SAR

ATM Services

ATM is a connection-oriented networking technology. As we will see throughout the following description, connection-oriented networking differs significantly from connectionless networking. One can best explain the differences by thinking of normal life examples.

Connectionless networking means that one sends data into a network, and the network knows itself by analyzing certain parts of the data (in IP networks, for example, the destination IP network address) where to forward the data. This process is also generally known as ***routing***. Today's Internet is, to a large extent, built on that approach. For connectionless networking, the reader should be also aware that each data unit (in IP, generally called ***IP packet***) can take a different path through the network, whereas, as we will see in connection-oriented networking, all data always travels the same path.

Connection-oriented networking is very much the same as we know from the good old telephone service. If you want to place a call to a remote user, you pick up the phone (initiate connection setup), then dial the number (define the destination of the remote user where the call should be set up), and, finally, let the

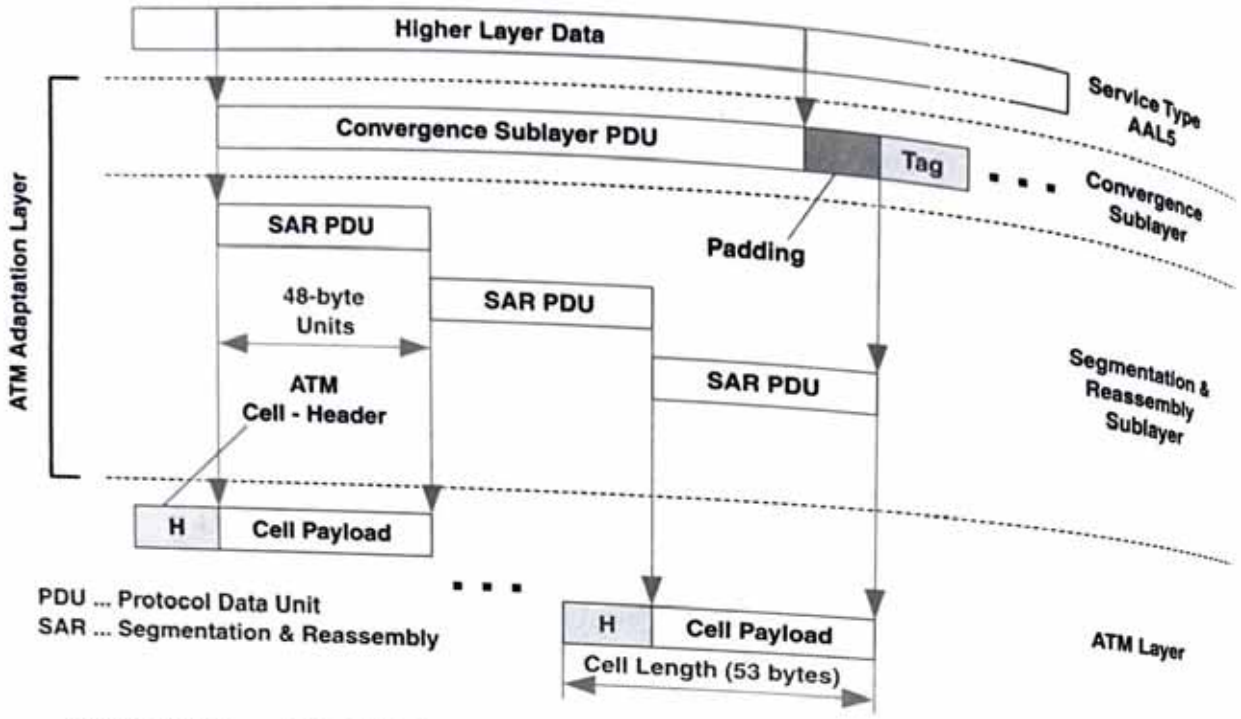


Figure 1–11 The AAL5 also incorporates two sublayers, the CS and the SAR

ATM Services

ATM is a connection-oriented networking technology. As we will see throughout the following description, connection-oriented networking differs significantly from connectionless networking. One can best explain the differences by thinking of normal life examples.

Connectionless networking means that one sends data into a network, and the network knows itself by analyzing certain parts of the data (in IP networks, for example, the destination IP network address) where to forward the data. This process is also generally known as ***routing***. Today's Internet is, to a large extent, built on that approach. For connectionless networking, the reader should be also aware that each data unit (in IP, generally called ***IP packet***) can take a different path through the network, whereas, as we will see in connection-oriented networking, all data always travels the same path.

Connection-oriented networking is very much the same as we know from the good old telephone service. If you want to place a call to a remote user, you pick up the phone (initiate connection setup), then dial the number (define the destination of the remote user where the call should be set up), and, finally, let the

telephone network establish the connection. After this successful setup, both parties can talk (exchange data in connection-oriented networking—often referred to as the ***data transfer phase***), and after finishing the conversation, both parties hang up (they tear down the connection). So we have three phases: connection setup, data transfer, and connection teardown.

End-to-end connectivity is achieved by provisioning virtual connections through the ATM network. Each ATM switch in the meshed network along the path of a virtual connection is configured to switch the virtual connection from the incoming interface to the right outgoing interface.

There are three types of ATM services:

1. Permanent virtual connection (PVC)
2. Switched virtual connection (SVC)
3. Connectionless service (which is similar to SMDS)

Permanent Virtual Connection

A PVC allows direct connectivity between sites. In this way, a PVC is similar to a leased line.

Advantages of PVCs are the guaranteed availability of a connection and the fact that no call setup procedures are required between switches. Disadvantages of PVCs include static connectivity and the manual administration required to set up.

Switched Virtual Connection

An SVC is created and released dynamically and remains in use only as long as data is being transferred. In this sense, it is similar to a telephone call. Dynamic call control requires a signaling protocol between the ATM endpoint and the ATM switch.

Advantages of SVC include connection flexibility and call setup that can be automatically handled by a networking device. Disadvantages include the extra time and overhead required to set up the connection.

ATM Virtual Connections

ATM networks are fundamentally connection-oriented. This means that a virtual channel (VC) needs to be set up across the ATM network prior to any data transfer. (A virtual channel is roughly equivalent to a virtual circuit.)

ATM connections are of two types:

1. Virtual paths, identified by virtual path identifiers (VPIs)
2. Virtual channels, identified by the combination of a VPI and a virtual channel identifier (VCI)

A virtual path is a bundle of virtual channels, all of which are switched transparently across the ATM network on the basis of the common VPI. All VCIs and VPIs, however, have only local significance across a particular link and are remapped, as appropriate, at each switch.

A transmission path is a bundle of VPs. Figure 1–12 shows how VCs concatenate to create VPs, which, in turn, concatenate to create a transmission path.

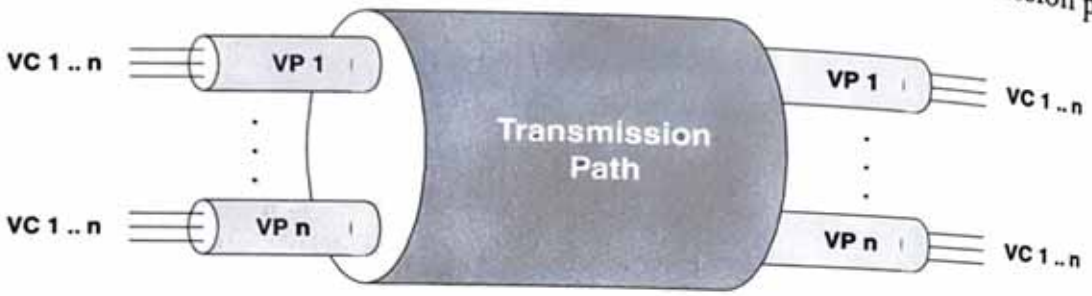


Figure 1–12 A typical ATM transmission path includes numerous virtual paths, which themselves includes multiple virtual circuits

ATM Switching Operation

The basic operation of an ATM switch is straightforward. A cell is received across a link on a known VCI or VPI value. The connection value is then looked up in a local translation table to determine the outgoing port (or ports) of the connection and the new VPI/VCI value of the connection on that link. Then the cell is transmitted on that outgoing link with the appropriate connection identifiers. Because all VCIs and VPIs have only local significance across a particular link, these values get remapped, as necessary, at each switch.

ATM Addressing

The ITU-T standardized upon the use of telephone-number-like E.164 addresses for public ATM (B-ISDN) networks.

The ATM Forum extended ATM addressing to include private networks. It decided on the "subnetwork" or "overlay" model of addressing, in which the ATM layer is responsible for mapping network layer addresses to ATM

1. Virtual paths, identified by virtual path identifiers (VPIs)
2. Virtual channels, identified by the combination of a VPI and a virtual channel identifier (VCI)

A virtual path is a bundle of virtual channels, all of which are switched transparently across the ATM network on the basis of the common VPI. All VCIs and VPIs, however, have only local significance across a particular link and are remapped, as appropriate, at each switch.

A transmission path is a bundle of VPs. Figure 1–12 shows how VCs concatenate to create VPs, which, in turn, concatenate to create a transmission path.



Figure 1–12 A typical ATM transmission path includes numerous virtual paths, which themselves includes multiple virtual circuits

ATM Switching Operation

The basic operation of an ATM switch is straightforward. A cell is received across a link on a known VCI or VPI value. The connection value is then looked up in a local translation table to determine the outgoing port (or ports) of the connection and the new VPI/VCI value of the connection on that link. Then the cell is transmitted on that outgoing link with the appropriate connection identifiers. Because all VCIs and VPIs have only local significance across a particular link, these values get remapped, as necessary, at each switch.

ATM Addressing

The ITU-T standardized upon the use of telephone-number-like E.164 addresses for public ATM (B-ISDN) networks.

The ATM Forum extended ATM addressing to include private networks. It decided on the "subnetwork" or "overlay" model of addressing, in which the ATM layer is responsible for mapping network layer addresses to ATM

addresses. This model is an alternative to using network layer protocol addresses (such as IP and IPX) and existing routing protocols (IGRP, RIP).

The ATM Forum defined an address format based on the structure of the OSI network service access point (NSAP) addresses.

Subnetwork Model of Addressing

The subnetwork model of addressing decouples the ATM layer from any existing higher layer protocols, such as IP or IPX. As such, it requires an entirely new addressing scheme and routing protocol.

All ATM systems need to be assigned an ATM address, in addition to any higher layer protocol addresses. This requires an ATM address resolution protocol (ATM_ARP) to map higher layer addresses to their corresponding ATM address.

NSAP-Format ATM Addresses

The 20-byte NSAP-format ATM addresses are designed for use within private ATM networks, whereas public networks typically use E.164 addresses that are formatted as defined by ITU-T. The ATM Forum did specify an NSAP encoding for E.164 addresses. This will be used for encoding E.164 addresses within private networks but can also be used by some private networks. Such private networks can base their own (NSAP format) addressing on the E.164 address of the public user-to-network interface (UNI) to which they are connected and take the address prefix from the E.164 number, identifying local nodes by the lower order bits.

ATM Address Common Components

All NSAP-format ATM address-es consist of three components:

1. ***Authority and format identifier (AFI)***: The AFI identifies the type and format of the initial domain identifier (IDI).
2. ***Initial domain identifier (IDI)***: The IDI identifies the address allocation and administrative authority.
3. ***Domain-specific part (DSP)***: The DSP contains actual routing information.

There are three formats of private ATM addressing that differ by the nature of the AFI and IDI:

1. ***NSAP-encoded E.164 format***: In this case, the IDI is an E.164 number.

2. ***DCC format***: In this case, the IDI is a data country code (DCC). These identify particular countries, as specified in ISO 3166. Such addresses are administered by the ISO National Member Body in each country.
3. ***ICD format***: In this case, the IDI is an international code designator (ICD). These are allocated by the ISO 6523 registration authority (the British Standards Institute). ICD codes identify particular international organizations.

The ATM Forum recommends that organizations or private network service providers use either the DCC or ICD formats to form their own numbering plan.

Figure 1–13 shows the three formats of ATM addresses for use in private networks.

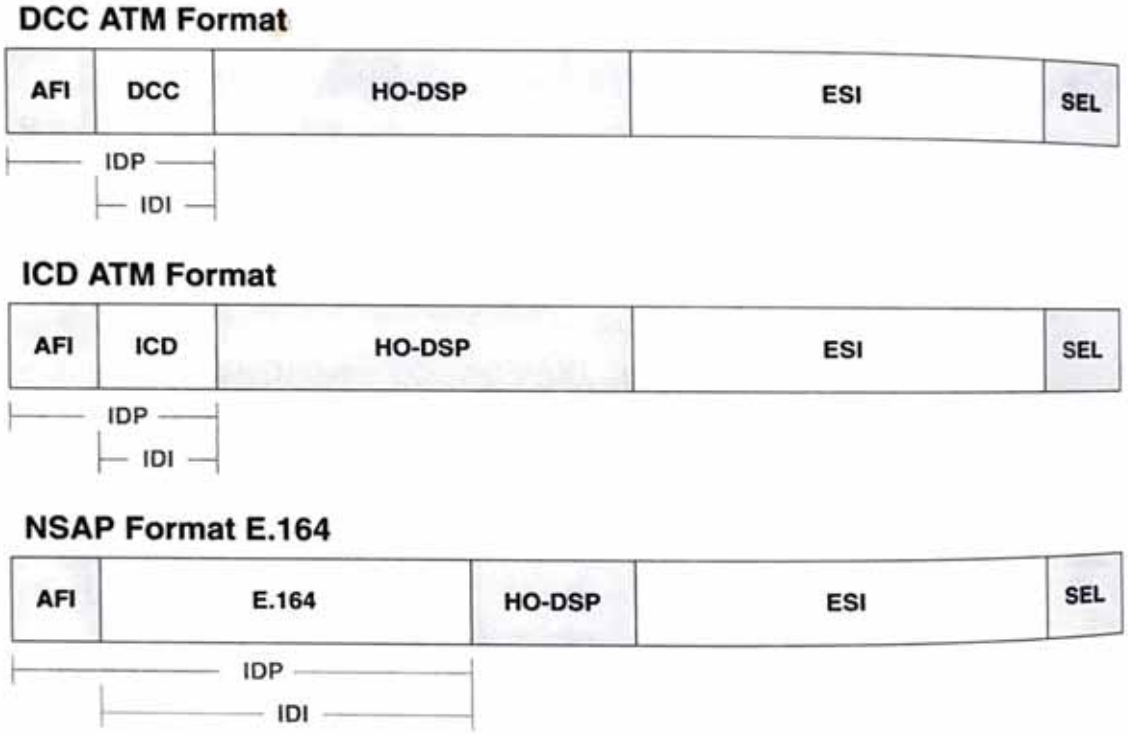


Figure 1–13 There are three address formats used in ATM networks: DCC, ICD, and NSAP addresses

ATM Quality of Service

ATM supports QoS guarantees comprising the following components: traffic contract, traffic shaping, and traffic policing.

When an ATM end-system connects to an ATM network, it is making a "contract" with the network based on QoS parameters. This ***traffic contract*** specifies an envelope that describes the intended data flow. This envelope specifies values for peak bandwidth, average sustained bandwidth, burst size, and more.

ATM devices are responsible for adhering to the contract by means of ***traffic shaping***. Traffic shaping is the use of queues to constrain data bursts, limit peak data rate, and smooth jitter so that traffic will fit within the promised envelope.

ATM switches can use ***traffic policing*** to enforce the contract. The switch can measure the actual traffic flow and compare it with the agreed-upon traffic envelope. If the switch finds that traffic is outside of the agreed-upon parameters, it can set the cell loss priority (CLP) bit of the offending cells. Setting the CLP bit makes the cell "discard eligible," which means that any switch handling the cell is allowed to drop the cell during periods of congestion.

ATM Signaling and Connection Establishment Overview

ATM signaling is a set of protocols used for call/connection establishment and clearing over ATM interfaces. The interfaces of interest to the ATM Forum are illustrated in Figure 1–14.

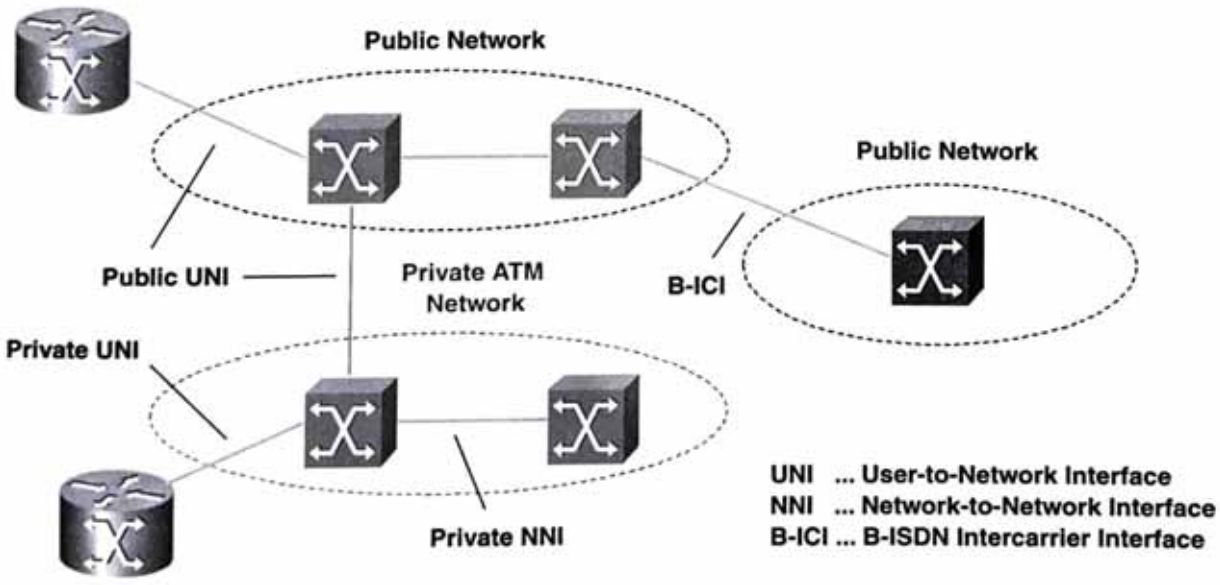


Figure 1–14 There are several types of interfaces specified by the ATM Forum

The public UNI is the user-to-network interface between an ATM user and a public ATM network. The private UNI is the user-to-network interface

between an ATM user and a private ATM network. B-ISDN intercarrier interface (B-ICI) is the network-to-network interface between two public networks or switching systems. Private-to-network node interface (NNI) is the network-to-network interface between two private networks or switching systems.

When an ATM device wants to establish a connection with another ATM device, it sends a signaling request packet to its directly connected ATM switch. This request contains the ATM address of the desired ATM endpoint, as well as any QoS parameters required for the connection.

ATM signaling protocols vary by the type of ATM link. ***UNI signaling*** is used between an ATM end-system and ATM switch across ATM UNI. ***NNI signaling*** is used across NNI links.

SIGNALING STANDARDS

The ATM Forum UNI 3.1 and UNI 4.0 specifications are the current standards for ATM UNI signaling. The UNI 3.1 specification is based on the Q.2931 public network signaling protocol developed by the ITU-T. UNI signaling requests are carried in a well-known default connection: VPI = 0, VPI = 5. For further reading, see the ATM Forum specification "ATM User to Network Interface Specification V3.1" [AF-2] or the ATM Forum Specification "UNI Signaling 4.0" [AF-3].

ATM CONNECTION ESTABLISHMENT PROCEDURE

ATM signaling uses the "one-pass" method of connection setup used in all modern telecommunication networks (such as the telephone network). The ATM connection setup procedure works as follows. The source end-system sends a connection signaling request. The connection request is propagated through the network and reaches the final destination. The destination either accepts or rejects the connection request and, in case of accept, the connections get set up through the network. Figure 1–15 highlights the "one-pass" method of ATM connection establishment.

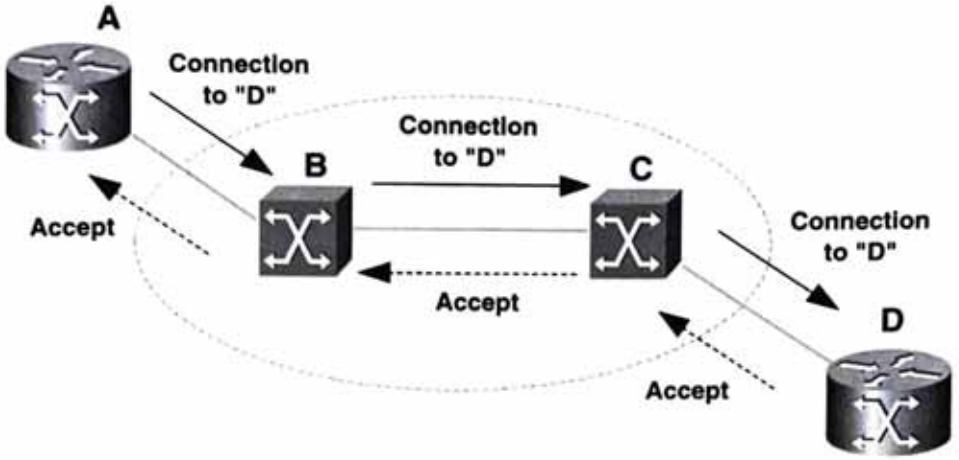


Figure 1–15 During the connection establishment, a request from the source end-system is propagated through the ATM network and either accepted or rejected by the destination end-system

Connection Request Routing and Negotiation

Routing of the connection request is governed by an ATM routing protocol. ATM routing protocols route connections based on destination and source addresses, and the traffic and QoS parameters requested by the source end-system. These routing protocols are usually either proprietary ones or the ATM Forum standard Private NNI routing protocol (PNNI).

The goal of the ATM Forum was to define NNI protocols for use within private ATM networks or, more specifically, within networks that use NSAP format addresses. Meanwhile, public networks that use E.164 numbers for addressing are interconnected using a different NNI protocol stack, based on the ITU-T B-ISDN user part (B-ISUP) ISUP signaling protocol. Negotiating a connection request that is rejected by the destination is limited because call routing is based on parameters of initial connection, and changing parameters might, in turn, affect the connection routing

Providing Internet Services at the IP Layer

Introducing the IP Layer on Top of the ATM Layer

For IP service delivery, the network layer and higher layers are required. Thus, IP routers are usually connected to ATM switches, which are interconnected across the ATM network by provisioning ATM virtual connections, as shown in Figure 1–16. These connections can be either static (where they are called *per-*

manent virtual connections, or PVCs) or dynamic (where they are called ***switched virtual connections***, or SVCs).

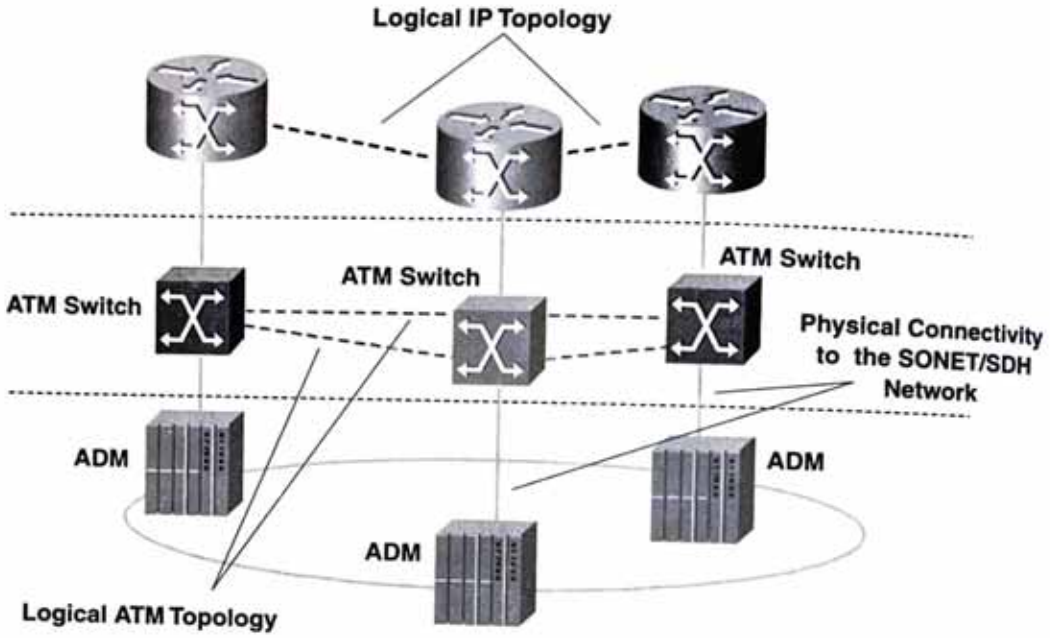


Figure 1–16 IP routers are connected to ATM switches, and virtual circuits are used to deploy the desired IP network topology

There are several methods used for running IP traffic over ATM networks. The first solutions used CLIP, defined by RFC 1577 [IETF-30].

In this case, the routers use an address resolution server—called the ***ARP server***—to find out the ATM addresses to connect to in order to reach an IP destination network. This approach is commonly known as the ***overlay model***, because the routers connected via an ATM cloud have no glue and, furthermore, cannot influence how the ATM network will set up a connection to reach the IP destination. There are several issues in terms of scalability and optimized routing throughout the ATM network, because CLIP works only in an IP subnet, and the ATM routing protocols have no information about the real needs of IP traffic delivery in terms of IP QoS, multicast, and others. Although CLIP was enhanced by additional protocols, such as Next Hop Resolution Protocol (NHRP) to overcome the limitation of suboptimal routing between subnets, the other issues (such as QoS mapping and IP-optimized connection setup) throughout the ATM network could not be solved.

Thus, the industry came up with other solutions by trying to combine the knowledge of IP destination reachability with the ATM backbones. These

approaches are commonly known as ***peer-to-peer solutions***, because they can be viewed as running IP and ATM routing protocols on the same layer, making them real peers and making the ATM switches IP aware, thus enabling them to set up IP-optimized routes during connection establishment. There were several approaches, mostly driven by single vendors, such as IP switching, or TAG switching. Finally, MPLS became the standard way to go. It is currently already very stable because it is defined by the Internet Engineering Task Force (IETF) in several Internet standards (RFC) and will be discussed in the following paragraphs of this section.

Using Multiprotocol Label Switching for Delivering IP Services

Scalability Issues of Standard IP Routing Protocols

Normal IP routing is based on the exchange of network reachability information via a routing protocol, such as Open Shortest Path First (OSPF) or others. The routers are examining the IP destination address contained in the IP header of each incoming packet and are using the network reachability information to determine how to forward the packet. This process, also called ***routing lookup***, is independently performed at each router-hop along the path of a packet.

As described before, the IP network is typically deployed on top of an ATM infrastructure. This layer-2 network is completely independent from the IP network. This fact leads to two major scalability issues.

First, to provide optimal routing in the IP network, any-to-any connectivity between all routers is desired. This leads to the demand for n*(n-1)/2 virtual connections provisioned in the ATM network. Hence, whenever a router is to be added, a virtual connection has to be provisioned to each other router.

Second, network failures or topology changes provoke massive routing protocol traffic. Each router has to send routing updates across each virtual connection it is connected to in order to inform its neighbor about the new IP network reachability situation.

From the routing design point of view, another issue can be seen. As an example, consider a typical Internet service provider (ISP) network in which there are multiple routers at the edge of the network having peering relationships with other ISP routers to exchange routing information and to achieve global IP connectivity. To be able to find the best path to any destination outside the ISP's network, the routers in the core of the ISP network must also know all the net-

work reachability information that the routers at the edge learn from the routers they are peering with. Hence, all core routers in the ISP network must maintain the entire Internet routing table, which requires a high amount of memory and leads to high CPU utilization. The solution to that issue is a forwarding technique where global network reachability is handled at the edge and packet forwarding "rules" are propagated to the core network.

THE CONCEPT OF USING LABELS AS FORWARDING INFORMATION

MPLS intends to address the outlined issues. MPLS has been standardized within IETF over the past few years. MPLS introduces a fundamental new approach for deploying IP networks. It separates the control mechanism from the forwarding mechanism and introduces the "label" used for packet forwarding.

MPLS can be deployed in router-only networks to address the routing protocol design issues or can also be deployed in ATM environments for integrating both the layer-2 and layer-3 infrastructure into an IP + ATM network. An MPLS network consists of label switch routers (LSRs) in the core of the network and edge label switch routers (Edge-LSRs) surrounding the LSRs, as shown in Figure 1–17.

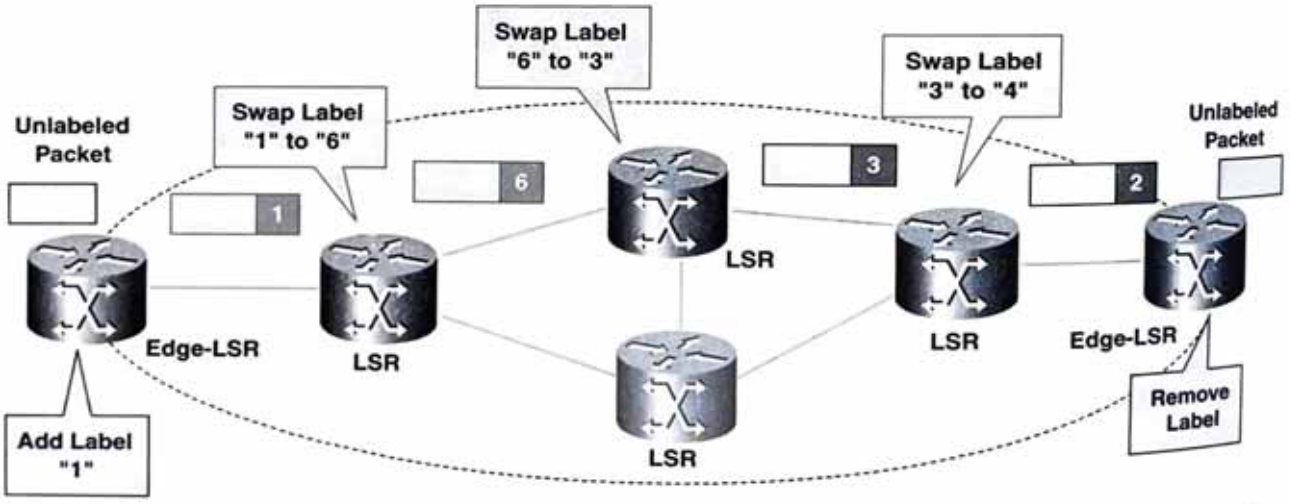


Figure 1–17 Edge-LSRs are imposing labels at the ingress side that are then used by the LSRs in the core to forward the traffic along the desired path and are removed again at the egress side by an Edge-LSR

Within the MPLS network, traffic is forwarded using labels. The Edge-LSRs at the ingress side of the MPLS cloud are responsible for assigning the label and

work reachability information that the routers at the edge learn from the routers they are peering with. Hence, all core routers in the ISP network must maintain the entire Internet routing table, which requires a high amount of memory and leads to high CPU utilization. The solution to that issue is a forwarding technique where global network reachability is handled at the edge and packet forwarding "rules" are propagated to the core network.

THE CONCEPT OF USING LABELS AS FORWARDING INFORMATION

MPLS intends to address the outlined issues. MPLS has been standardized within IETF over the past few years. MPLS introduces a fundamental new approach for deploying IP networks. It separates the control mechanism from the forwarding mechanism and introduces the "label" used for packet forwarding.

MPLS can be deployed in router-only networks to address the routing protocol design issues or can also be deployed in ATM environments for integrating both the layer-2 and layer-3 infrastructure into an IP + ATM network. An MPLS network consists of label switch routers (LSRs) in the core of the network and edge label switch routers (Edge-LSRs) surrounding the LSRs, as shown in Figure 1–17.

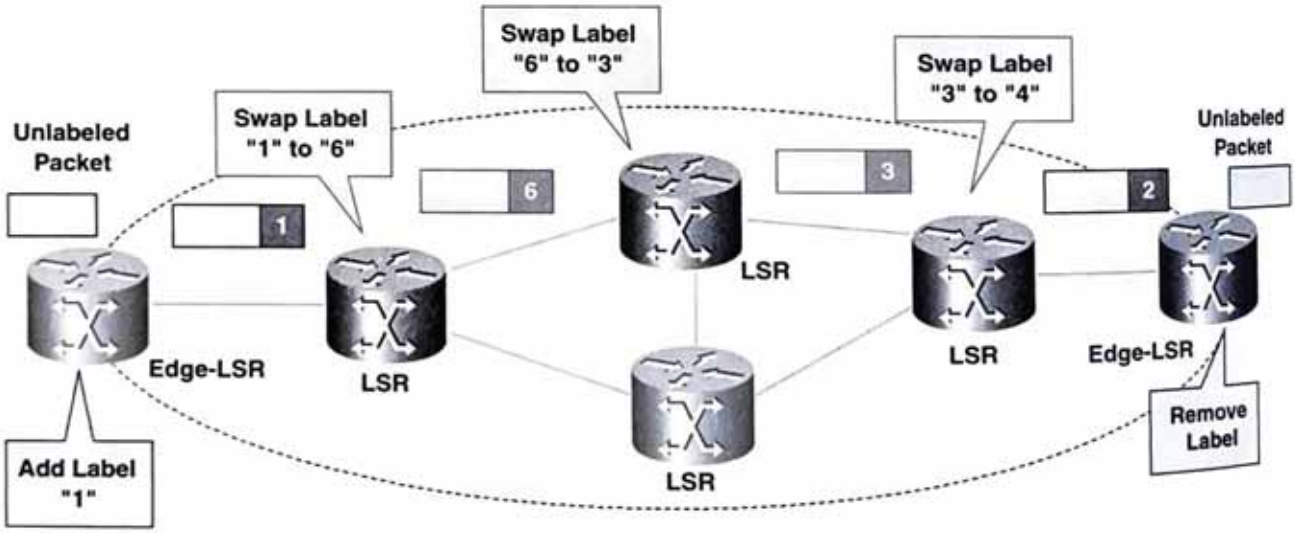


Figure 1–17 Edge-LSRs are imposing labels at the ingress side that are then used by the LSRs in the core to forward the traffic along the desired path and are removed again at the egress side by an Edge-LSR

Within the MPLS network, traffic is forwarded using labels. The Edge-LSRs at the ingress side of the MPLS cloud are responsible for assigning the label and

forwarding the packets on to the next-hop LSR along the path that the traffic follows through the MPLS cloud. All LSRs along the path use the label as an index into a table that holds the next hop information and a new label. The old label is exchanged with the new label from the table and the packet is forwarded to the next hop. Using this method implies that the label value has only local significance between two LSRs. At the egress side of the network, the label is removed, and traffic is forwarded, using normal IP-routing protocol mechanisms.

MPLS introduces the term ***forwarding equivalent class*** (FEC). An FEC is a group of packets that share the same attributes while they travel through the MPLS domain. These attributes can be the same destination IP network, the same quality of service class, the same Virtual Private Network (VPN), etc. An FEC can also be a combination of some of these attributes. A label is assigned to each FEC (Figure 1–18), so all packets belonging to the same FEC will get the same label assigned from the LSR.

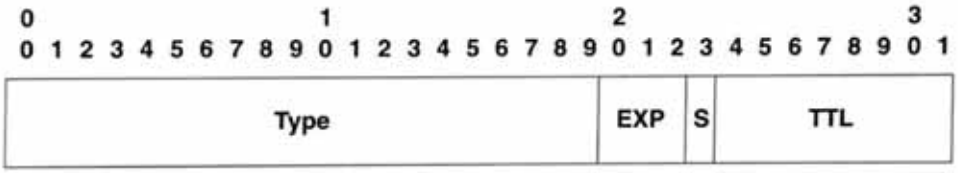


Figure 1–18 Each FEC gets its own locally unique label assigned, which is construed of four fields

The MPLS architecture allows that multiple labels forming a ***label stack*** are assigned to a packet. This might be for traffic engineering purposes, where one label is used to represent a tunnel and the other label to represent the FEC. Another possibility might be that one label is used to denote a certain VPN identifier, and another label is the FEC used for forwarding traffic across the provider network. To be able to distinguish whether label stacking is used, the ***bottom of stack*** (S) bit is part of each label. It is set to one for the last entry in the label stack (i.e., for the bottom of the stack) and zero for all other label stack entries.

Each IP packet has a ***time to live*** (TTL) field to prevent packets from looping endlessly in a network. The TTL field maintains a counter that gradually decrements down to zero, at which point the packet is discarded. A TTL 8-bit field is also included in the MPLS label, mainly as a troubleshooting mechanism.

A 3-bit field is reserved for ***experimental use***. These bits are already used by some applications to determine the traffic priority and are to be used for providing certain levels of QoS in the MPLS network.

The rest of the label is used for the ***label value*** itself. This 20-bit field carries the actual value of the label.

The control and forwarding mechanisms of each LSR and Edge-LSR are separated into a control and forwarding plane (Figure 1–19). Two information structures are part of the ***control plane***. The first one is the IP table routing maintained through running an Interior Gateway Protocol (IGP), such as OSPF or others.

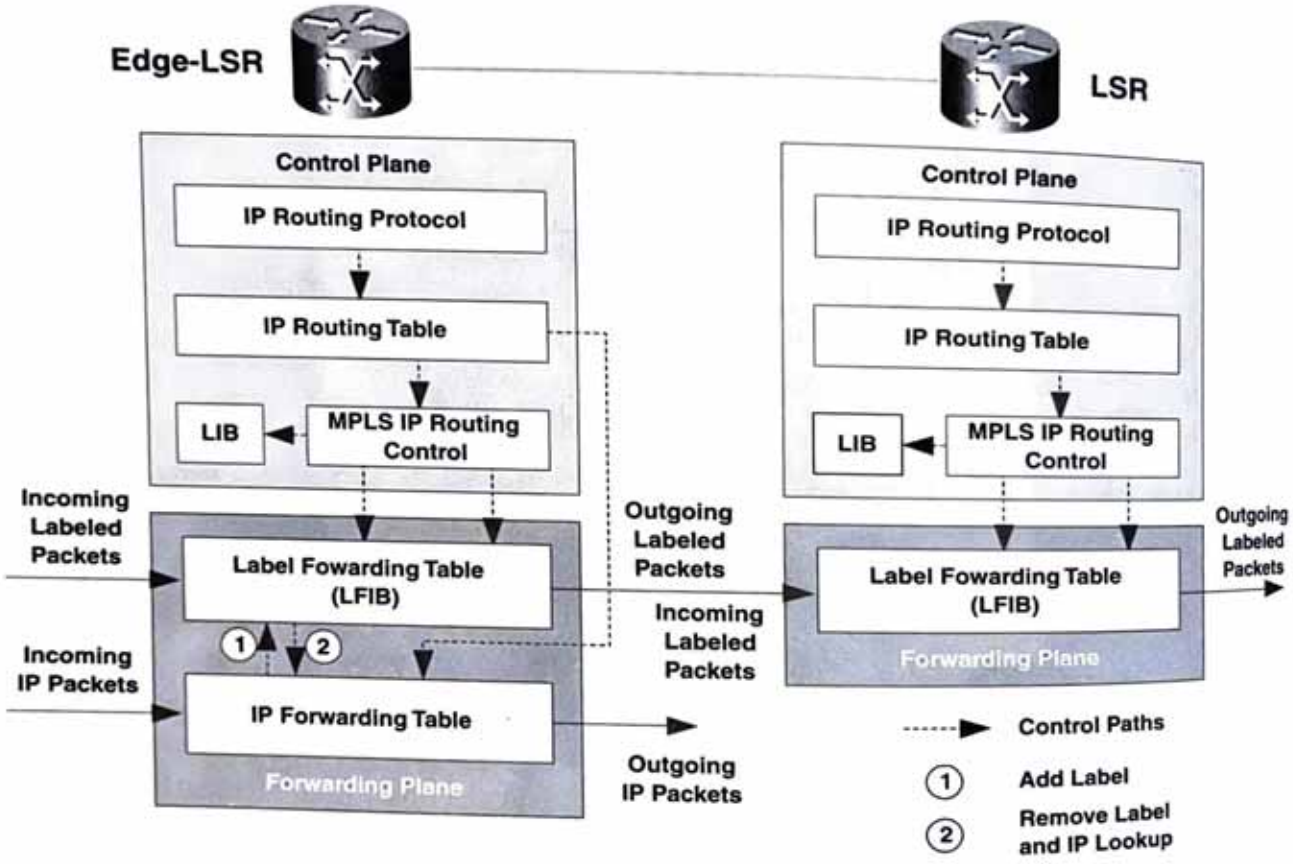


Figure 1–19 Each LSR (right side) or Edge-LSR (left side) has the control mechanism separated from the forwarding mechanism

The second one is a database called ***label information base*** (LIB). LSRs need to exchange the information on which label is associated with which FEC. There exist different protocols that perform this task. A protocol called *Label*

Distribution Protocol (LDP) is used to exchange the information on which prefix is associated with which label. To exchange the information on which VPN is associated with which label, the most common protocol is Multiprotocol BGP (MP-BGP). However, there is a variety of different protocols that enable the exchange of label-mapping information between LSRs.

Considering LDP as the protocol to be used, the label information is exchanged as follows. LDP uses ***downstream label allocation***. This means that an LSR is told by its downstream neighbor which label to use when forwarding a packet belonging to a certain FEC out of the desired interface (Figure 1–20). Basically, there might be two situations. If the downstream LSR explicitly asks for a label binding for a certain FEC, and the downstream LSR answers with a label-binding message, this is called ***downstream on demand***. If the upstream LSR did not ask for a specific label binding, and the downstream LSR sends out a label-binding message, this is called ***unsolicited downstream*** or simply ***downstream label distribution***. Binding for the incoming label values is performed locally.

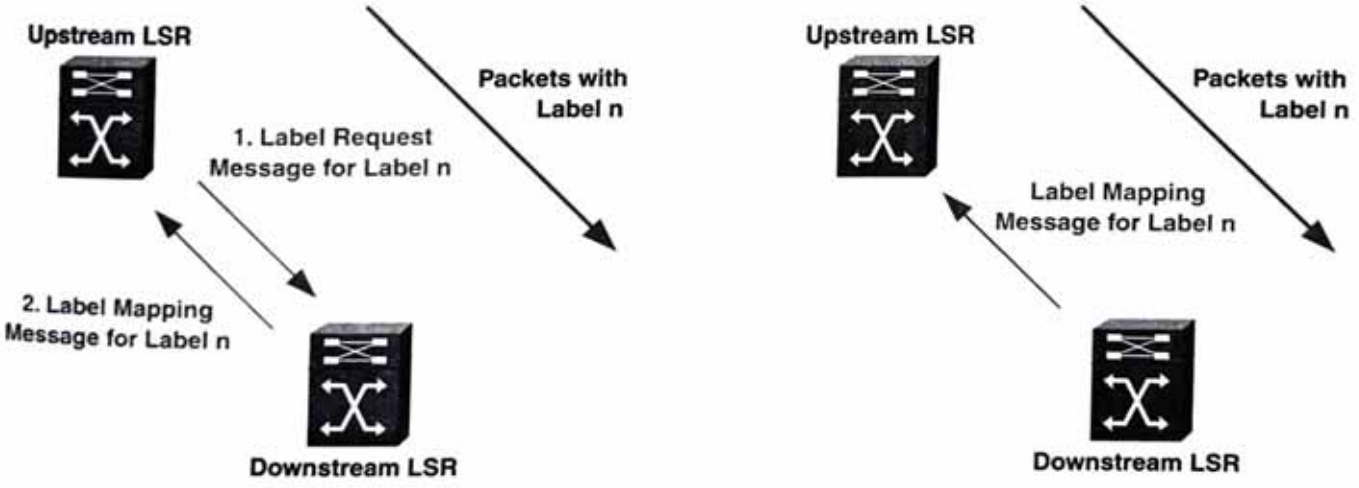


Figure 1–20 The LDP uses downstream label allocation, which might be requested by the upstream LSR (left half) or forced by the downstream LSR (right half)

The ***forwarding plane*** contains one information structure, which is a cache called ***label-forwarding information base*** (LFIB). The LFIB cache handles the actual packet-forwarding process. It contains information such as incoming label value, outgoing label value, prefix/FEC, outgoing interface (and encapsulation type of this interface), and the next hop address.

During the packet-forwarding process, this cache contains all required information to switch the packet or the cell through the LSR.

MPLS-BASED APPLICATIONS

In the early drafts of MPLS, we could see the main application of MPLS in integrating IP and ATM. In the later drafts, we see that the authors are focusing more on MPLS applications such as VPNs and traffic engineering, because these concepts can be used as real differentiators between legacy-routed networks and MPLS networks. Service providers are trying to create and sell value-added services to their customers that they can use to differentiate themselves from the competition and to take a step up in the value chain. The opportunity for these service providers to offer VPN services based on IP, such as BGP/MPLS VPNs, does make this technology very attractive for the market. It also explains a little bit of the "hype" for MPLS that we could see in the market for the last two years. In our view, there is no slowdown of this situation visible on the horizon.

The most-deployed MPLS application is the MPLS VPN. A typical MPLS VPN network is shown in Figure 1–21. The customer access routers, also called ***customer edge*** (CE) routers, are connected to Edge-LSRs, acting as provider edge (PE) routers. The PE routers are assigning two labels to each packet. One label is representing the VPN identifier, and the top label is used to forward the packet through the network (Figure 1–21). The LSRs in the core of the network are called ***provider*** (P) routers and are performing standard label switching using the top label. The PE router at the egress side of the network is removing both labels and is using the second label to determine to which CE (VPN) the packets are to be forwarded.

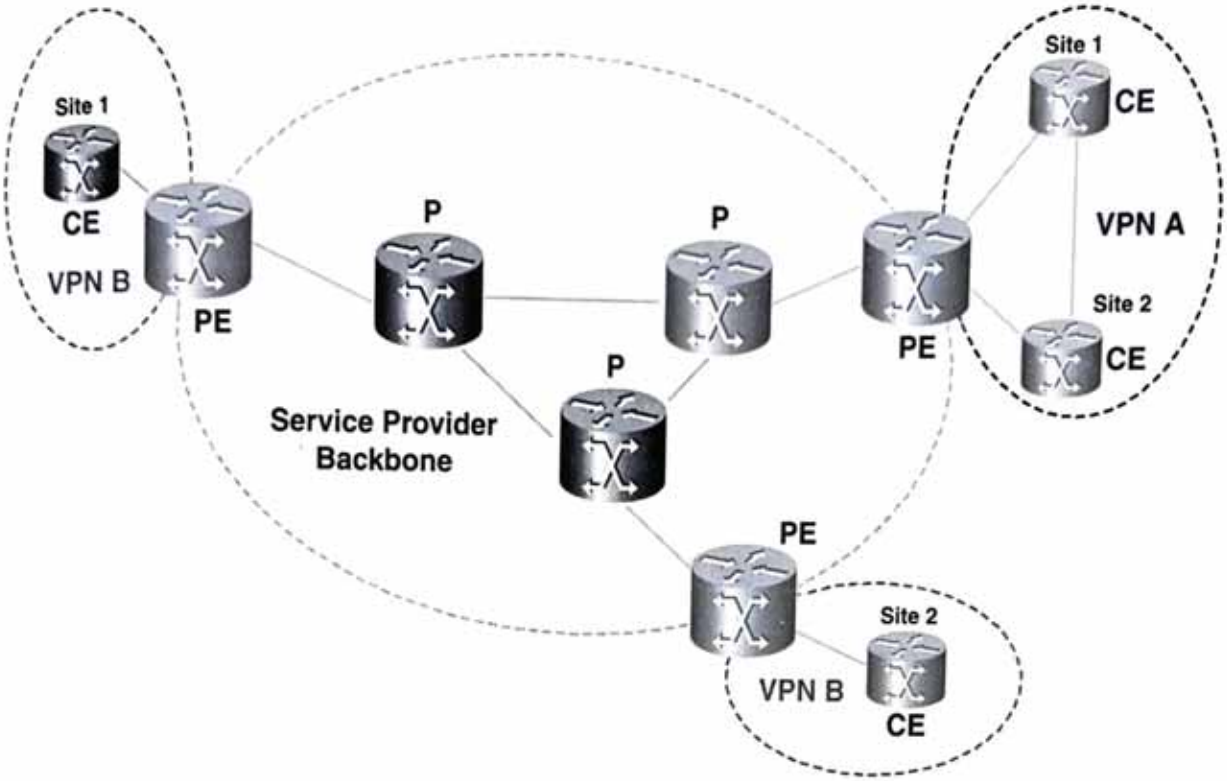


Figure 1–21 MPLS VPNs—The PE routers use the label stack to determine to which VPN the routing information and traffic belongs

The second application makes use of the fact that the control plane is completely separated from the forwarding plane. Standard routing protocols compute the optimum path from a source to a certain destination, considering a routing metric such as hop-count, cost, or link bandwidth. As a result, one single least-cost path is chosen. Although there might be an alternate path, only the one determined by the routing protocol is used to carry traffic. This leads to an inefficient utilization of network resources.

MPLS Traffic Engineering (MPLS-TE) introduces the term ***traffic trunk***, in which a group of traffic flows with the same requirements, such as resilience or traffic priority. MPLS-TE provides traffic-driven IGP route calculation functionality on a per-traffic-trunk basis. Other than with standard routing protocols, which are only topology driven, utilization of network resources and resilience attributes are analyzed and taken into account during path computation. As shown in Figure 1–22, multiple paths are possible from a source to a destination, and the best according to the current network situation is taken to ensure optimum network utilization.

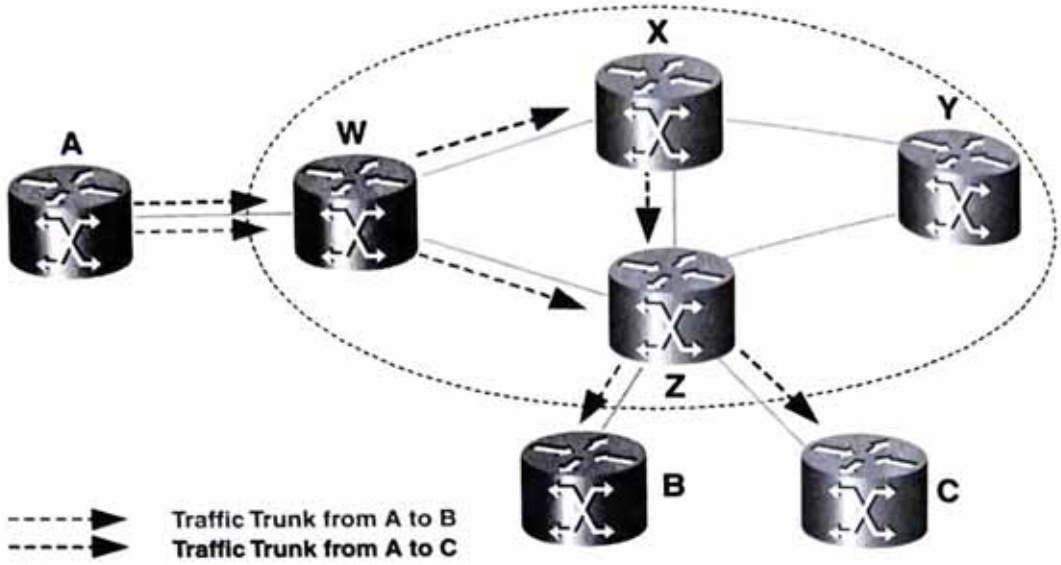


Figure 1–22 MPLS Traffic Engineering provides the ability to define and optimize the path to be taken by the traffic across the network

Next-Generation Carrier Networks

The explosive growth of Internet/intranet traffic is making its mark on the existing transport infrastructure. An unprecedented shift has occurred in traffic content, pattern, and behavior. It has transformed the design of multiservice networks and has created a commercial demand for IP networks that operate in excess of 1 Gigabit per second. A change in usage patterns from fixed, configured, connection-oriented services to dynamic, connectionless IP services is currently evolving. According to several studies, telecommunication revenue will grow significantly in the near future, with data services, particularly IP, accounting for most of this increase. For public carriers, therefore, IP is critical for future revenue growth. Figure 1–23 shows how fast IP data traffic has been growing during the last few years and is expected to grow similarly in the next years.

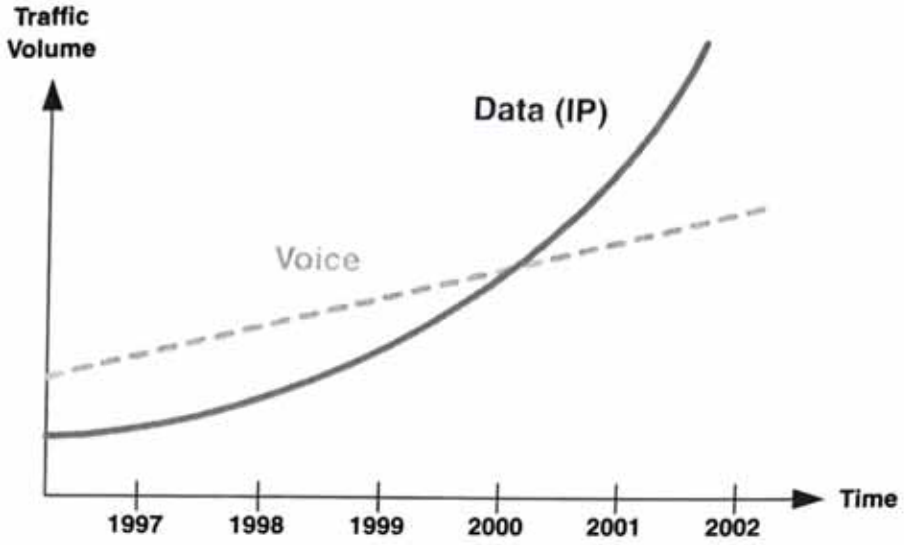


Figure 1–23 Data traffic has already overtaken voice traffic

As IP traffic volumes have exploded in recent years, bandwidth requirements for IP data have already reached limits that require the entire network architecture to be restructured in order to cope with the exponential growth.

IP backbone interconnections have already reached interface rates of OC-192/STM-64—thus, the highest level of SONET/SDH. Furthermore, the massive transport inefficiency of ATM, with its overhead up to 30%, overrides the advantages of the multiservice integration and QoS functionality.

As a consequence, the following trends and changes are forming the next-generation carrier networks.

Eliminating Intermediate Layers in the Backbone

With IP as the dominant traffic in the network, the layered network architecture is no longer appropriate. The goal is to minimize the transmission overhead to maximize useful transport bandwidth. Certainly, complexity in provisioning, operations, network planning, and engineering must also be reduced to minimize the service provider's operational expenses (OPEX) and maximize the profit.

As can be seen from Figure 1–24, the SONET/SDH and ATM layers are going to be eliminated, transforming the backbone network into a two-layer network at the final stage. Fault detection and resolution, and network restoration become critical issues. By reducing the intermediate SONET/SDH layer,

the comprehensive restoration mechanisms are removed and must be implemented within either the optical or IP layer. As shown in Figure 1–24, there are four network architectures possible. However, the question of the control plane for the optical network is important and not yet solved. Should the control be derived from the optical layer and run IP in an overlay fashion (notice the analogy to IP over ATM)? Should IP control the optical plane (peer model approach)? Should even other control planes deliver routing decisions (for example, content- or policy-derived control planes)? As stated earlier in this chapter, these issues will be discussed in greater detail in the Chapter 4, "Existing and Future Optical Control Planes."

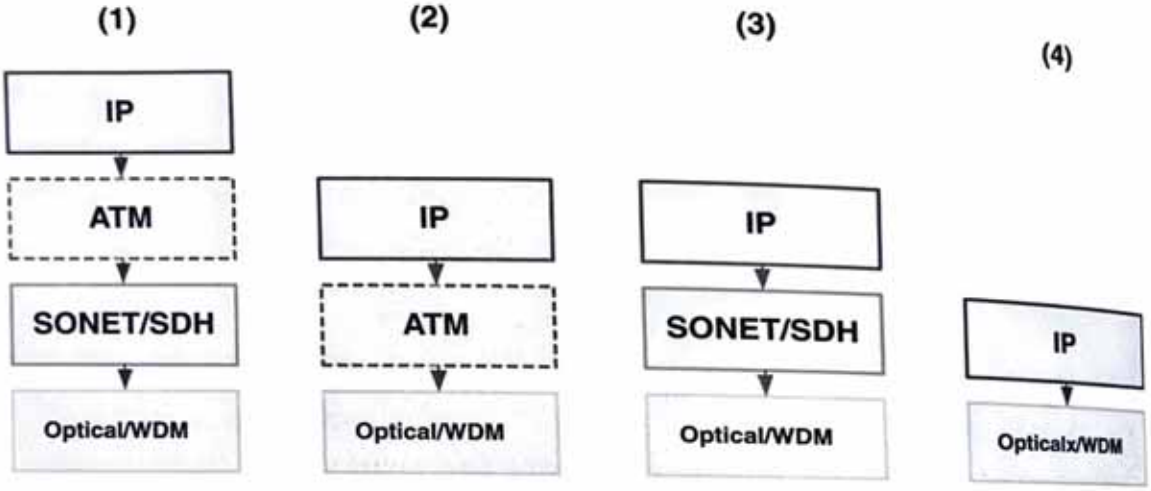


Figure 1–24 Transforming toward a two-layer network architecture

IP over ATM over SONET/SDH over WDM (1)

Traditional service providers offering TDM, ATM, and IP services use networks consisting of all four layers: IP, ATM, SONET/SDH, and optical/WDM. The transmission efficiency is very low because of the high ATM overhead, including AAL5 and ATM cell overhead. High resilience against fiber cuts is provided through the SONET/SDH infrastructure but with a high initial equipment cost and high operational costs, due to complex provisioning and network management. WDM might be used to increase the capacity of the physical fiber plant by transmitting multiple wavelengths over a single fiber.

IP over ATM over WDM (2)

A similar approach is to eliminate the SONET/SDH layer but to maintain the ATM layer. Service providers with a high amount of layer 2 services, such as leased lines and voice services, commonly use this approach and build a three-layer network, using the IP, ATM, and optical/WDM layer. The drawback of this architecture is, again, the inefficient transport of IP traffic because ATM is still there. On the other hand, all the other advantages of ATM, such as the strong QoS capabilities, remain.

IP over SONET/SDH over WDM (3)

Alternative datacentric service providers today commonly implement a three-layer network architecture using the IP, SDH, and optical/WDM layer. Packet over SONET/SDH (POS) technology is used for transporting IP data directly over SONET/SDH infrastructure to eliminate the inefficient ATM transmission layer. SONET/SDH infrastructure might already be installed because it is in widespread use (especially in Europe) and can be used to provide protection functionality. It can also be used to transport legacy voice traffic.

IP over WDM (4)

The fourth approach is to eliminate both the SONET/SDH and ATM layer, thus building only a two-layer network using the IP and optical/WDM layers. This approach is commonly known as an ***optical network*** and is mostly taken by service providers delivering data services and voiceover data service (VoIP).

The transport efficiency is optimized because the ATM layer is removed. The initial equipment is reduced because no SONET/SDH equipment is used. WDM technology may be used to increase the capacity of the fiber plant.

Handling Fiber Capacity Exhaust

To meet the exponential growth in bandwidth demand of today's provider networks, Wavelength Division Multiplexing (WDM) technology is used to increase fiber capacity. Multiple wavelengths are sent over a single fiber. Each wavelength channel is capable of carrying any type of client signal. These client signals can be either ATM or SONET/SDH traffic but will mostly be IP traffic.

WDM systems typically have 4 or 16 channels. As transmission techniques evolve, systems with more than 40 channels have been developed and are called ***Dense WDM*** (DWDM) systems. ITU-T has standardized 100-GHz and 50-GHz channel spacings to ensure vendor interoperability. Closed WDM systems use a proprietary method to transport traffic over the WDM trunks and to provide standardized interfaces, such as SONET/SDH, ATM, POS, or Gigabit Ethernet, using wavelength transponders. Open WDM systems provide standard optical interfaces according to the ITU-T DWDM grid specifications and with SONET/SDH framing at 2.5 or 10 Gbps. These systems can be used with any SONET/SDH equipment, ATM switch, or IP router equipped with ITU-T wavelength grid conform interfaces.

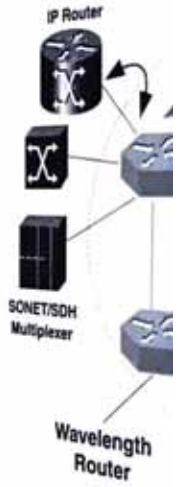
WDM systems can be used with point-to-point connections but also in ring applications. Service providers have the ability to build up multigigabit connections in their networks. By using WDM technology, networks become very scalable. Bandwidth demands can be easily met by adding some wavelength channels to the system, providing additional "virtual fibers."

Furthermore, service providers can offer "optical leased lines" by providing dedicated wavelengths to customers. This revolutionary new type of service delivers enhanced flexibility to customers because of the bit rate independence of the wavelength service. Customers can build up traditional TDM-based networks with SONET/SDH or new-world data networks by directly interconnecting IP routers with different wavelengths.

Adding Intelligence to the Optical Layer

Capacity is no longer the differentiating factor for service providers in the optical networking field. DWDM is commonly used and delivers massive amounts of raw bandwidth. Although DWDM allows multiple wavelengths to be transmitted over a single fiber, each wavelength must be terminated on an individual network element, typically a SONET/SDH ADM or IP router. This leads into an excessive amount of required equipment and a very complex and rarely manageable network architecture.

Transforming the optical layer from being a group of point-to-point pipes to a resilient, manageable optical network introduces devices called ***wavelength routers***. Carriers can build a new optical layer that provides dynamic provision-



ing, reconfiguration for optimizing network resources, and protection and restoration at the wavelength level.

Wavelength routers are huge cross-connects, switching wavelengths through a meshed optical network (see Figure 1–25). Through the wavelength-routing capability implemented in the wavelength routers, end-to-end wavelength paths can be dynamically provisioned and restored. By managing the network at wavelength granularity, intelligent optical network elements scale to meet upcoming demands. Furthermore, carriers are able to assign different QoS levels or service types to wavelength paths.

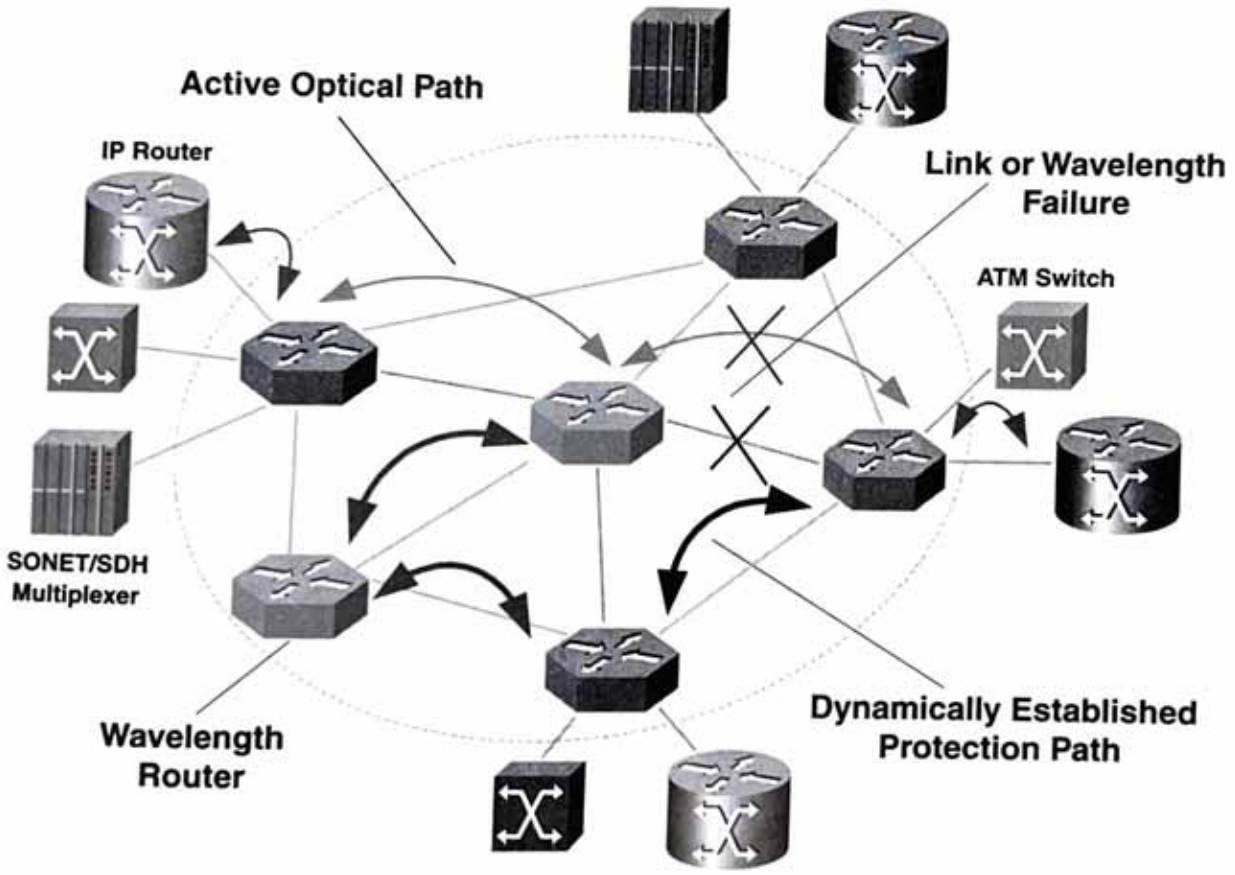


Figure 1–25 Wavelength provisioning using wavelength routers

Wavelength routers have a lot of additional benefits, such as rapidly increasing connectivity between service nodes, allocating bandwidth as dictated by service demands, and restoring connectivity between service nodes in event of failure. In addition, they allow providers to allocate restoration bandwidth according to business drivers, to simplify service network architectures, to offer

next-generation bandwidth services with varying QoS at wavelength granularities, and to scale and introduce new technologies easily into their optical core.

A very important fact that must be mentioned is that the growth of service capacity will lead to higher costs if bandwidth is not manageable. It must be possible to manage the massive amount of new capacity provided by DWDM technology and to provision appropriate customer services. Basically, there are two approaches for managing a network. Centralized management ensures rapid provisioning; a single action propagates reactions throughout the network to activate a particular service or initiate rerouting across an alternative route. With distributed management, intelligence is put into every network element, and the network has "self-healing" capabilities. Therefore, network elements can react to failures without operator intervention and provide real-time survivability and reduced operational costs.

Summary

In this chapter, we presented an overview of how traditional carrier networks have been designed until today. These designs have resulted in a complex network architecture, consisting of multiple network layers, where WDM provides optical transmission capacity, SONET/SDH delivers resilience against fiber cuts, ATM enhances the bandwidth utilization through statistical multiplexing and provides certain QoS functionality, and IP (on top of that architecture) delivers the interface to the service providers' data service offerings.

In contrast to these past implementations, next-generation carrier networks will be optimized for delivery of IP services while providing up to a terabit of capacity in the most scalable and flexible way. They will also provide easy, seamless, end-to-end provisioning. The key to this transition is to transform the old multilayer network into a two-layer network, meaning transporting IP directly on top of the optical layer. Having discussed these trends, we can now move on to the current and most popular standardization activities in this industry.

Recommended Reading

[CIENA-1] CIENA Corp. Whitepaper, *Fundamentals of DWDM.*

[CIENA-2] CIENA Corp. Whitepaper, *The Evolution of DWDM.*

[CIENA-3] CIENA Corp. Whitepaper, *The New Economics of Optical Core Networks*, September 1999.

[CNET-1] CANARIE Inc., *Architectural and Engineering Issues for Building an Optical Internet*, draft, Bill St. Arnaud, September 1998.

[CSCO-1] Cisco Systems Inc. Whitepaper, *Scaling Optical Data Networks with Wavelength Routing*, 1999.

[CSCO-4] Networkers Conference—Session 606, *Advanced Optical Technology Concepts*, Vienna, October 1999.

[CSCO-5] Cisco Systems Inc. Whitepaper, *Cisco Optical Internetworking*, 1999.

[NFOEC-1] Technical Paper, *"IP over WDM" the Missing Link*, P. Bonenfant, A. Rodrigues-Moral, J. Manchester, Lucent Technologies, A. McGuire, BT Laboratories, September 1999.

[NFOEC-2] Technical Paper, *Packet over SONET and DWDM—Survey and Comparison*, D. O'Connor, Li Mo, E. Catovic, Fujitsu Network Communications Inc., September 1999.

[NFOEC-3] Technical Paper, *Position, Functions, Features and Enabling Technologies of Optical Cross-Connects in the Photonic Layer*, P. A. Perrier, Alcatel, September 1999.