

**1 YEAR UPGRADE**  
BUYER PROTECTION PLAN



# ASP [Application Service Provider]

# Configuration

## Handbook

### **Become an Application Service Provider!**

- Step-by-Step Instructions for Converting an ISP from Standard Bandwidth Provisioning to Providing Complex Services
- Hundreds of Common ASP Terms Defined, Types of ASP Firms Identified, and Best Platforms Revealed
- Complete Coverage of Application Outsourcing, Business Process Outsourcing, and Platform IT Outsourcing

**Gary Palmatier,**  
Vice President of Business Development and Solution  
Architecture, EngineX Networks, Inc.

**Foreword by Dale Booth,**  
Chairman & CEO, EngineX Networks, Inc.

**Sean Thurston** Technical Editor

s o l u t i o n s @ s y n g r e s s . c o m

With more than 1,500,000 copies of our MCSE, MCSA, CompTIA, and Cisco study guides in print, we continue to look for ways in which we can better serve the information needs of our readers. One way we do this is by listening.

Readers like yourself have been telling us they want an Internet-based service that would extend and enhance the value of our books. Based on reader feedback and our own strategic plan, we have created a Web site that we hope will exceed your expectations.

**Solutions@syngress.com** is an interactive treasure trove of useful information focusing on our book topics and related technologies. The site offers the following features:


- One-year warranty against content obsolescence due to vendor product upgrades. You can access online updates for any affected chapters.
- “Ask the Author”™ customer query forms that enable you to post questions to our authors and editors.
- Exclusive monthly mailings in which our experts provide answers to reader queries and clear explanations of complex material.
- Regularly updated links to sites specially selected by our editors for readers desiring additional reliable information on key topics.

Best of all, the book you’re now holding is your key to this amazing site. Just go to [www.syngress.com/solutions](http://www.syngress.com/solutions), and keep this book handy when you register to verify your purchase.

Thank you for giving us the opportunity to serve your needs. And be sure to let us know if there’s anything else we can do to help you get the maximum value from your investment. We’re listening.

[www.syngress.com/solutions](http://www.syngress.com/solutions)

SYNGRESS®



**1 YEAR UPGRADE**  
BUYER PROTECTION PLAN



# ASP [Application Service Provider]

# Configuration

Handbook

Syngress Publishing, Inc., the author(s), and any person or firm involved in the writing, editing, or production (collectively “Makers”) of this book (“the Work”) do not guarantee or warrant the results to be obtained from the Work.

There is no guarantee of any kind, expressed or implied, regarding the Work or its contents. The Work is sold AS IS and WITHOUT WARRANTY. You may have other legal rights, which vary from state to state.

In no event will Makers be liable to you for damages, including any loss of profits, lost savings, or other incidental or consequential damages arising out from the Work or its contents. Because some states do not allow the exclusion or limitation of liability for consequential or incidental damages, the above limitation may not apply to you.

You should always use reasonable care, including backup and other appropriate precautions, when working with computers, networks, data, and files.

Syngress Media®, Syngress®, and “Career Advancement Through Skill Enhancement®,” are registered trademarks of Syngress Media, Inc. “Ask the Author™,” “Ask the Author UPDATE™,” “Mission Critical™,” and “Hack Proofing™” are trademarks of Syngress Publishing, Inc. Brands and product names mentioned in this book are trademarks or service marks of their respective companies.

KEY	SERIAL NUMBER
001	KKHJH87875
002	GSRUY6456
003	ALERKTPD53
004	E458XPS368
005	5ERPTK348A
006	NV49533JFE
007	Q8URVNA394
008	BHU89FE2MP
009	Q2WPMKA843
010	CDFUU8Z922

PUBLISHED BY  
Syngress Publishing, Inc.  
800 Hingham Street  
Rockland, MA 02370

#### **ASP Configuration Handbook: A Guide for ISPs**

Copyright © 2001 by Syngress Publishing, Inc. All rights reserved. Printed in the United States of America. Except as permitted under the Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of the publisher, with the exception that the program listings may be entered, stored, and executed in a computer system, but they may not be reproduced for publication.

Printed in the United States of America

1 2 3 4 5 6 7 8 9 0

ISBN: 1-928994-26-1

Technical Editor: Sean Thurston  
Co-Publisher: Richard Kristof  
Cover Design by: Michael Kavish

Freelance Editorial Manager: Maribeth Corona-Evans  
Copy Editor: Beth A. Roberts  
Index by: Robert Saigh  
Page Layout and Art by: Shannon Tozier

Distributed by Publishers Group West in the United States.



# Acknowledgments

We would like to acknowledge the following people for their kindness and support in making this book possible.

Richard Kristof and Duncan Anderson of Global Knowledge, for their generous access to the IT industry's best courses, instructors, and training facilities.

Karen Cross, Lance Tilford, Meaghan Cunningham, Kim Wylie, Harry Kirchner, Bill Richter, Kevin Votel, Brittin Clark, and Kent Anderson of Publishers Group West for sharing their incredible marketing experience and expertise.

Mary Ging, Caroline Hird, Simon Beale, Caroline Wheeler, Victoria Fuller, Jonathan Bunkell, and Klaus Beran of Harcourt International for making certain that our vision remains worldwide in scope.

Anneke Baeten, Annabel Dent, and Laurie Giles of Harcourt Australia for all their help.

David Buckland, Wendi Wong, Daniel Loh, Marie Chieng, Lucy Chong, Leslie Lim, Audrey Gan, and Joseph Chan of Transquest Publishers for the enthusiasm with which they receive our books.

Kwon Sung June at Acorn Publishing for his support.

Ethan Atkin at Cranbury International for his help in expanding the Syngress program.

Joe Pisco, Helen Moyer, and the great folks at InterCity Press for all their help.





# Contributors

**Dale Booth** is EngineX Networks' Chairman and CEO. He has more than two decades of telecommunications industry experience, including the role of Senior Vice President and Chief Operating Officer at Fujitsu Network Services, a division of Fujitsu Network Communications. Dale created the vendor-neutral services division and powered it to a 373 percent compound annual growth rate. He also served as Chief Information Officer of Fujitsu Network Communications. Prior to that, he held various technical and management positions at InteCom, a PBX start-up specializing in converged voice and data platforms. Dale has an engineering degree from DeVry Institute and did post-graduate work at the University of Pennsylvania, Wharton School of Business. He serves on numerous boards and councils, including Daisytek International, the Texas Quality Foundation, and the International Engineering Consortium.

**Gary Palmatier** is Vice President of Business Development and Solution Architecture for EngineX Networks Inc. He is responsible for defining markets, creating technical service solutions, and managing customer, partner, and alliance relationships. He has more than 18 years of telecommunications industry experience. Previously, he was Director of Advanced Network Services at Fujitsu Network Services, where he established the professional services business unit and directed the division's global professional services efforts. Gary has a master's degree in IS management from Aurora University, in Aurora, IL and a bachelor's degree in industrial technology from Southern Illinois University. Gary serves as Chairman of the Systems Integration Committee of the National Convergence Alliance. He is certified as an Oracle Master Systems Analyst and is a member of the IEEE. Gary has been a featured speaker at Argonne National Laboratory, professional societies, and industry symposiums on topics such as software engineering, innovation, and management systems. He has also taught electronics technology, programming languages, and networking at the college level.



**Matt Lyons** (CCIE #1133) is the Director of Solution Architecture for EngineX Networks Inc. Matt has over 20 years of experience working with networking and IP. He has worked for such early network pioneers as Ungermann Bass, SynOptics Communications, Network General, and Nestar Systems, the first company to network the Apple II. After more than 7 years at Cisco Systems, and leveraging his international consulting experience, he has joined EngineX Networks to build a Solution Architecture group that is focused on large-scale carrier network issues and design. Matt lives in Fremont, CA.

**Kevin Murphy** (CCNA, CSE) is the Director of Business Development for EngineX Networks Inc. Kevin is responsible for developing the partnerships and alliances required to support EngineX Networks' advanced designs and deployments. Previously, Kevin served on the EngineX Design team as a Solution Architect and Engagement Manager. His areas of focus include VoIP, content networking, and VPN. He has also held various business development and sales positions within the Infrastructure segment. Kevin holds a bachelor's degree in business administration from the University of Southern California with an entrepreneur emphasis. He lives in San Francisco, CA.

**Aaron Davidson** (CCNA) is a Solution Architect with EngineX Networks Inc. He creates infrastructure and data center design solutions. His specialties include security, load balancing, and implementation of various vendors' equipment. Aaron's background includes designing and securing several Silicon Valley dot.com ventures and working for numerous Internet service providers. Aaron lives in the East Bay and is working on his certifications and a music career.

**Mark Egan** (CCNP, MCSE, MCP+I, CNE) is a Senior Solution Architect for EngineX Networks Inc. He provides technical leadership for the Solution Architecture team as well as support for the design and implementation of customer networks. Mark's background includes working for Sprint-Paranet as a Technical Analyst and as a Consultant for Exxon. Mark Lives in Dublin, CA.

**Ben Tsui** (CCNP, CCDA, MCSE, MCP+I, MBA) is a Senior Network Engineer with EngineX Networks Inc. He designs telecommunications infrastructures and implements network devices as well as the provision of local loop for transport deployment. His specialties include research, documentation, and implementation. Ben's background includes positions as Design Engineer for PacBell and a Network Engineer Specialist for SBC Datacomm. Ben lives in Fremont, CA. He is pursuing his master's degree in telecommunications.



# Technical Editor and Contributor

**Sean Thurston** (CCDP, CCNP, MCSE, MCP+I) is a Senior Solution Architect with EngineX Networks Inc. He provides network and data center design solutions for large-scale deployment. His specialties include implementation of multi-vendor routing and switching equipment and XoIP (everything over IP) installations. Sean's background includes positions as a Technical Analyst for Sprint-Paranet and the Director of a brick and mortar advertising dot-com company. Sean is also a contributing author to Syngress Publishing's *Building a Cisco Network for Windows 2000*, ISBN: 1-928994-00-8. Sean lives in Renton, WA. He is currently pursuing his CCIE.

# Contents

## Choose the Best Platform for Your ASP

ASP s take advantage of existing Internet connectivity to offer corporations the opportunity to outsource not only peripheral applications but also mission-critical applications. This trend will continue to escalate as customers discover that outsourcing firms can deliver mission-critical applications that meet their demands for service licensing agreements (SLAs). For ISPs and ASPs, these application-hosting responsibilities require the choice of a platform that can deliver the correct balance of performance, scalability, upgradeability and manageability.

<b>Foreword</b>	<b>xxix</b>
<b>Chapter 1 An Introduction to ASPs for ISPs</b>	<b>1</b>
Introduction	2
Why This Book Is for You	3
What This Book Can Do for You	4
Whom This Book Is Written For	5
Definitions of Common ASP Terms	5
What Is an Internet Service Provider?	5
What Is an Application Service Provider?	6
The Pure ASP	6
What Is Information Technology Outsourcing?	6
Application Outsourcing	7
Business Process Outsourcing	7
Platform Information Technology Outsourcing	7
The Elements That Make an ASP Viable	8
Life Cycle for the Cost of Ownership	8
The Initial Cost of Hardware Acquisition	9
Hardware Maintenance and Associated Costs	10
Initial System Software Package Acquisition	10
Initial Application Software Package Acquisition	10
Implementation	10
The Cost of Hardware Upgrades	11
The Cost of System Software Upgrades	11

The Cost of Application Software Upgrades	12
Network Administration Resources	12
Other Support (Training, Help Desk, Etc.)	12
Possible Business Models and Offerings	12
Types of ASP Firms	13
Professional Consulting	14
Project-Based Service Providers	14
Outsourcing Providers	15
Staff Augmentation Providers	15
Education and Training Providers	16
Value-Added Resellers	16
The OSI-ISO Seven Layer Model	16
Layer 1: The Physical Layer	17
Layer 2: The Data-Link Layer	18
Layer 3: The Network Layer	18
Layer 4: The Transport Layer	18
Layer 5: The Session Layer	19
Layer 6: The Presentation Layer	19
Layer 7: The Application Layer	19
The Upper Layers	21
The Lower Layers	21
The Pseudo Layers	21
Layer 8: The Political Layer	21
Layer 9: The Religion Layer	22
Layer 10: The Financial Layer	22
Choosing the Best Platform for Your ASP	22
Hardware	23
Servers	24
Hewlett-Packard	24
Sun Microsystems	25
Compaq	25
Network Equipment	25
Data Traffic Explosion	25
Alcatel Networks	26
Cisco Systems	26
Extreme Networks	26

F5 Networks	27
Foundry Networks	27
Juniper Networks	27
Lucent Technologies	27
Nortel Networks	28
Cache Appliance Makers	28
Akamai	28
Intel	28
Inktomi	28
Software	29
Load-Balancing Software	29
BEA Systems	30
Hewlett-Packard WebQoS	30
IBM	31
Microsoft	32
Resonate	32
Segue	33
Business Drivers for the Conversion to ASP	34
Business Factors That Impact the ASP Model	34
Enabling Technologies	35
Technical Factors	36
Barriers to the ASP Business Model	37
ASP Business Model Strategies	38
System Integrators and Implementers	40
Internet Service Providers and Telecommunication Companies	41
Independent Software Vendors	42
Independent Software Vendor Companies	43
Why All the Mergers?	45
Performance Issues	45
Amount of System Uptime (Five Nines)	46
Failover	47
Clustering	47
Sun Microsystems	47
Hewlett Packard	47
Compaq	48

Problems That Could Arise from Conversion	48
Major Issues in the Implementation of an ASP Model	49
What Is Needed to Sell Your Services? Necessary Components	50
Summary	51
Solutions Fast Track	51
Frequently Asked Questions	57

**Review ASP Decision  
Criteria**

Current Analysis published the results of their survey of ASP customers that ranked the major decision criteria they used to choose an ASP provider. Major factors included:

- Support Capabilities
- Hosting and Facility Experience
- Cost and Pricing Structure
- Reputation and Client Reference
- Service Level Agreement
- Past Performance
- Scalability and Completeness of Solution

**Chapter 2 The Business Case 59**

Introduction	60
ISP Market Conditions	61
The Onset of Commoditization	63
Broadband—The Enabling Technology	64
Service Provider Business Requirements	67
The New Model	68
Customers' Demands	69
Investor Demands	71
The Evolving ISP	72
The Steps Necessary to Offer Value	73
Deployment of Services	75
Value-Added Services and Core Competencies	78
The Service Provider of the Future	80
The Finances Involved	82
The Case for Application Service Provider Conversion	82
Market Factors	84
ASP Customer Value Proposition	86
ISP Value Proposition	88
ASP Services Also Enable Future Migration Up the Value Chain	91
ISP to ASP: The Perfect Fit?	91
Critical Success Factors	94
Business Models	94
Determining Your Offerings	96
Customer Issues	99

Summary	103
Solutions Fast Track	104
Frequently Asked Questions	106

## **Chapter 3 Server Level Considerations 109**

### **Understand How Servers Work in Your Environment**

ASP's need to support content, database information, storage area networking (SAN), and file servers to truly provide well-rounded application service offerings. These "server farms" must be maintained and supported with the priority that the application demands. Remember that the design of a system and implementation of an infrastructure that will work for your company and satisfy your customers will require exceptionally careful planning and forethought.


Introduction	110
Implementation, Where to Begin	111
Server Hardware	111
Central Processing Unit	112
Symmetric Multiprocessing	114
Random Access Memory	115
Mass Storage	116
Network Adapters	118
Software Solutions for Your ASP	133
System Software	133
Application Software Types	137
Web Applications	137
Database Applications	141
Middleware Software	142
Server Redundancy	143
Shared Device	144
Shared Nothing	144
Under- and Over-Subscription	145
Network Service Considerations	147
Network Storage	147
Network File System Protocol	148
Server Message Block Protocol	149
Common Internet File System	159
Data Backups and How They Can Affect You	159
Software Selection	162
Virus Scanning Suggestions	168
Thin Client Solutions	171
ICA Protocol	172
Maintenance and Support Issues	174
Planned Upgrades	174
Break/Fix	176
System Monitoring	177



Summary	179
Solutions Fast Track	180
Frequently Asked Questions	184

**Chapter 4 Performance Enhancement Technologies 187**

**Content Delivery Networks and Streaming Media**



Content Delivery Networks can provide more efficient service, as they can store and forward multiple formats and bit rates for streaming media. A number of companies now offer streaming media appliances:

- Midstream—  
[www.midstream.com](http://www.midstream.com)
- Network Engines—  
[www.networkengines.com](http://www.networkengines.com)
- Vingage—  
[www.vingage.com](http://www.vingage.com)
- Vividon—  
[www.vividon.com](http://www.vividon.com)

Introduction	188
What Is Web Caching?	189
What Is Load Balancing?	189
What Is Content Routing?	189
Web Caching and How It Works	189
What Is Data Caching?	190
The Benefits of Data Caching?	191
What Happens With and Without a Solution in Place	193
How to Reduce Bandwidth Usage	195
Key Requirements for a Caching Solution	195
Deployment Models for Data Caching	197
Forward Proxy	198
Transparent Caching	198
Reverse Proxy	198
Cache Locations and Placement	199
Cache Hierarchies	199
What Are Cache Appliances?	201
Cost Effectiveness	201
Ease of Installation and Management	201
Fault Tolerance	202
Scalability and Flexibility	203
Performance and Speed	203
Load Balancing in Your Infrastructure	204
Localized Load Balancing	204
Distributed Load Balancing	204
Comparing Different Load-Balancing Systems	205
Software-Only Solutions	206
Switches	207
Routers and Caching Systems	207
Clustering	208

Network Appliances	208
Criteria You Should Look for in a Superior Load-Balancing Solution	209
Dependability	209
Quality of Service	210
High Availability	210
Can Load Balancing Enhance and Extend Your Network?	211
Vendor Credibility and Their Support Infrastructure	211
Load-Balancing Solutions from F5	212
First-Generation Load-Balancing Solutions	213
What Takes a Site Down?	213
Guaranteeing Availability to Your Client	214
Cisco Systems' LocalDirector	215
Scaling a Server Farm	215
High Availability	217
Managing Your Server Connections	218
Security with the LocalDirector	219
LocalDirector Configuration Samples	219
Multiple Virtual Servers and One Real Server	223
Multiple Virtual Servers and Multiple Real Servers	226
Foundry Networks' ServerIron	228
Content Delivery Networks	230
Today's Content Delivery Landscape	231
Functional Components of a CDN	232
How Do CDNs Work?	232
Who Needs CDNs?	233
Content Providers	234
What Do Content Publishers Require from CDNs?	235
CDN Service Providers	237

What CDN Service Providers Require	238
CDN Deployment Basics and Considerations	239
Network Service Providers	239
Satellite-Based Network Service Providers	240
What Network Service Providers Require from CDN Service Providers and CDN Component Product Makers	240
CDN Product Manufacturers	240
Enterprises	241
Consumers	241
The CDN Services Landscape	241
Industry Standardization Efforts	241
The Content Alliance	242
The Content Bridge Alliance	243
CDN Solutions from Various Vendors	244
Inktomi Content Delivery Suite	244
Inktomi Content Distributor	246
Inktomi Content Manager	247
Cisco System's Content Delivery Networks and Next-Generation Content-Based Services	247
Cisco's CDN Group	248
Akamai and F5 Networks' Combined Offerings	249
Akamai's Solution	249
F5 Products	250
Summary	251
Solutions Fast Track	252
Frequently Asked Questions	255
<b>Chapter 5 Storage Solutions</b>	<b>257</b>
Introduction	258
Upfront Concerns and Selection Criteria	259
Concerns for Your Storage Devices	259
Host Independence	259
Mixed Vendor Support	260

## Choose the Best Storage Solution for Your Company



Six major concerns and criteria should be taken into account before deciding on the storage solution that best fits your requirements:

- Host independence
- Mixed vendor support
- Security
- Legacy support
- System availability
- Price versus performance

Security	260
Legacy Support	261
System Availability	262
Price versus Performance	262
Directly Attached Storage in Your Infrastructure	263
Network Attached Storage Solutions	264
Quality of Service	266
Location of NAS in Your Network	266
Storage Area Networks	267
The Need for SAN	267
Benefits of SAN	268
SAN Virtualization	270
Multihost Arrays	270
Logical Unit Number Masking	271
In-Band Virtualization	272
Storage Domain Servers	273
NAS versus SAN	274
Comparing Fiber Channel to SCSI	275
The Benefits of Fiber Channel	276
What Are the Limitations of SCSI?	277
All Fiber versus Mixed Solutions	277
SAN Management	280
Capacity Management	281
Configuration Management	281
Performance Management	281
Availability Management	281
Scalability and How It Affects Your Business	282
Storage in Your Infrastructure	282
Wire Speed and How It Can Help You	284
One versus Many	287
Fault Tolerance Features and Issues	288
Shared Resources	289
Data Backup	289
Remote Mirroring	290
Synchronous	290
Asynchronous	291

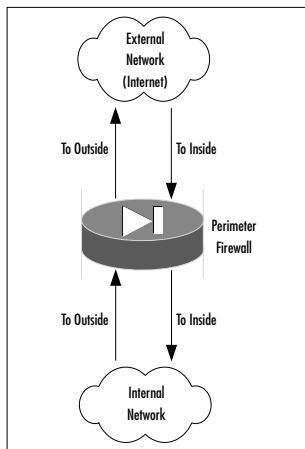
Redundant Array of Inexpensive Disks	291
RAID-0	293
RAID-1	293
RAID-2	293
RAID-3	293
RAID-4	294
RAID-5	294
RAID-6	294
RAID-10	294
RAID-53	295
SAN Solutions Offered by Various Vendors	295
IBM's SAN Solution	295
The IBM SAN Strategy	295
Summary	297
Solutions Fast Track	298
Frequently Asked Questions	301

## **Chapter 6 ASP Security System Provisioning**

**303**

Introduction	304
Security Policy	306
Developing a Security Policy	306
Privacy Policy	308
Security Components	309
Authentication	309
User Authentication	309
IP Addresses and Spoofing	312
Confidentiality Protection	313
Key Length	314
Types of Algorithms	315
Further Cryptographic Considerations	316
Incident Response	317
Security Auditing and Risk Assessment	319
Security Technologies and Attacks	320
Virtual Private Networks	320
Perimeter Firewalls	321
Stateful Inspection	323

## Security Components Discussed from an ASP's Perspective



Perimeter Firewall

Packet Filtering	324
Embedded Firewalls	328
Intrusion Detection Systems	330
Types of Attack	332
Applications Attack	332
Denial-of-Service	333
Buffer Overflow Attacks	334
SYN Attacks	335
IP Fragmentation Attack	336
Smurf Attack	336
Fraggle Attack	338
Physical Attacks	338
Distributed Denial of Service	339
Trinoo	341
Tribal Flood Network	341
Tribal Flood Networks 2000	342
Stacheldraht	342
Prevention Techniques	343
Filtering RFC1918 Address Spaces	344
Ingress and Egress Filtering	346
Rate Limiting	348
Flood Attacks	350
SYN Attacks	351
TCP Intercept	352
TCP Intercept Mode	353
TCP Intercept Timers	354
Drop Mode	354
Aggressive Mode Thresholds	354
Capturing Evidence	355
Syslog	356
Packet Capturing	357
Summary	358
Solutions Fast Track	359
Frequently Asked Questions	361

**Learn the Three Main Areas of Service Level Agreements**

- **Planning** Determining the wide area network (WAN) service levels.
- **Verification** Monitoring the service levels to guarantee fulfillment.
- **Troubleshooting** Isolating issues when service levels are not delivered.

<b>Chapter 7 Management and Monitoring</b>	<b>363</b>
Introduction	364
The Effect of Outsourcing	364
Service Level Agreements	365
Some Common SLA Guarantees	365
What Are the Basic Components of SLAs for Frame Relay Circuits?	366
What Service Levels Should the Service Provider Consider?	368
Network Availability	369
PVC Availability	369
Average Network Delay and Average PVC Delay	370
Effective Throughput	370
Response Time	371
Time to Resolution or Repair	371
The Realities of Customer Compensation	371
What Will Your Customers Look for in Their Implemented SLA?	372
What Are the Guidelines for Implementing the Monitoring Necessary to Handle These Tasks?	372
Where Is Your Weakest Link?	373
Network SLAs	374
System Level SLAs	374
Application SLAs	374
Making Your Company More Customer Oriented	375
How Service Providers Have Responded to Acceptable Performance	376
The Added Bonus	377
The Operation Support System Model	377
What Are the Basics of OSS?	378
The Workflow Engine	378
Ordering	379
Inventory and Allotment	379

Engineering and Provisioning	379
Activation and Service Management for the Field	379
Network Management and Support	380
What Is OSS Interconnection, and What Does It Mean?	381
What Are the Challenges Facing Interconnection?	382
Upgrading the OSS	382
Efficiencies in Your OSS	383
Remaining Flexible	383
API Functionality and Gateways	383
Supporting Your Data Services	384
Provisioning Data Service	385
Activation of Data Services	386
Broadband Access Changes the Market	386
Getting Access to the Masses	386
Quality of Service	387
Management Systems for Your ASP	388
The TMN Outline	388
TMN Standards	389
The Building Blocks of the TMN Model	391
How the OSI Functions in the TMN Model	392
Manager and Agent Roles	393
The Standard Interfaces	393
The Logical TMN Model	394
What Tools Do You Need to Automate TMN?	396
The ASP Transformation	397
Industry Examples of Successfully Deployed ASP Management Tools	398
ASP Infrastructure Operations	399
Network Operating System	400
Pricing Models and Billing	401
Billing	404
Managing Billing with Partners	405



	Summary	407
	Solutions Fast Track	408
	Frequently Asked Questions	414
	<b>Chapter 8 Designing the Infrastructure</b>	<b>415</b>
<b>Answer Your Questions about Capacity Planning</b>	Introduction	416
	Design Considerations	417
	Getting Started: The Design Process	418
	Data Center, WAN, and Remote Links Defined	419
	The Design Process—Getting Down to Business	420
	Site Considerations	421
	Physical Equipment Space	421
	Network Equipment Basics	424
	Designing with the Hierarchy in Mind	425
	Scalability of Hierarchical Internetworks	426
	Manageability of Hierarchical Internetworks	427
	Optimization of Broadcast and Multicast Control Traffic	427
	Possible Types of Topology Design	428
	Star Topologies	428
	Fully Meshed Topologies	429
	Partially Meshed Topologies	430
	Broadcast Issues	431
	Performance Issues	432
	Frame Relay Internetwork Design Considerations	432
	Hierarchical Design for Frame Relay Internetworks	433
	Hierarchical Meshed Frame Relay Internetworks	434
	Hybrid-Meshed Frame Relay Internetworks	436
	Regional Topologies for Frame Relay Networks	437
	Star Topologies	437
	Fully Meshed Topologies	437
<b>Q:</b>	I have agreed to a Committed Interface Rate (CIR) with a client. How do I charge them if they "burst" over this agreement?	
<b>A:</b>	Most ISP/ASP companies charge what is called the 95 percentile. This is where traffic is measured in timed intervals. You will be charging the customer for whatever their usage is for 95 percent of the time.	

Partially Meshed Topologies	438
Broadcast Issues for Frame Relay Networks	439
Creating a Broadcast Queue for an Interface	440
Committed Interface Rates	440
Capacity Planning for Your Infrastructure	442
Connection and Expansion	442
Best Practices	442
Protocol Planning Concerns	444
Routing Protocols	444
Interior Gateway Protocols	444
External Protocols	448
Choosing the Right Interior Protocol	448
Route Selection	449
Addressing Considerations	450
Topology	451
Application and Network Services	453
Designing the Data Center Network	454
Terminal Data Centers	454
Application-Aware Networking	455
Traffic Detection and Classification	455
Admission Control	455
Traffic Classification	456
Congestion Avoidance	457
Scheduling	458
Scalability Considerations	458
Scaling Bandwidth	458
Scaling Considerations	458
Multimedia Services	460
IP Multicast	461
Virtual LANs and Emulated LANs	462
Policy in the Core	463
WAN Link Considerations	464
Routing and Scalability	464
Planning for the Future Growth of Your Company's Infrastructure	465
Even More Network Scalability	465

Layer 2 Switching	466
Layer 3 Switching	466
Layer 4 Switching	467
Bridged Protocol Needs	467
Bridging in the Multilayer Model	468
Security in the Multilayer Model	468
High-Availability Design	469
High Availability	469
Things to Consider When Implementing High-Availability	469
Summary	472
Solutions Fast Track	473
Frequently Asked Questions	477

## **Appendix A Sample Configuration for an Application Service Provider Network 479**

Introduction	480
The Test Network	481
The Logical Network Overview	481
The Access Layer	481
The Distribution Layer	482
The Core Layer	484
Configuration with Cisco Systems Commands and References	485
Configuration for a Cisco Systems 7200 Router That Is Located within the Core Layer	486
Configuration for a Cisco Systems Gigabit Switch Router Router That Is Located within the Distribution Layer	509
Configuration for a Second Cisco Systems Gigabit Switch Router Router That Is Located within the Distribution Layer	522
Configuration for a Third Cisco Systems Gigabit Switch Router That Is Located within the Distribution Layer	532

Configuration for a Cisco Systems MGX Router That Is Located within the Access Layer	537
Summary	553
<b>Appendix B ASP Configuration Handbook Fast Track</b>	<b>555</b>
<b>Index</b>	<b>585</b>



# Foreword

For the modern Web-enabled community to grow and prosper, we need to create and deploy communication solutions that can collate an assorted collection of data and applications, while incorporating legacy solutions still in use and allowing for painless migration to future technologies.

For the Application Service Provider (ASP) to grow and prosper they must provide reliable access and integration of independent information, applications, and data stores into dynamic, interactive solutions that delight their customers.

## The Web of Change

Technology historian James Burke refers to events that lead to technological advancement as the “Web of Change.” In his numerous books, including *Connections*, *The Day the Universe Changed* (both of which were fantastic series on BBC and The Discovery Channel), and *The Pinball Effect*, Burke brought together disparate and apparently unconnected advancements in engineering, chemistry, and technology to show how these developments lead to unforeseen and unpredictable consequences.

When linking these events together, Burke always brought a sense of logical advancement to technology, while highlighting the haphazard nature of efforts to produce breakthrough results.

As an example of this, Burke explains that there was a need for early American rail systems to coordinate trains. The need was desperate because only a one-track line connected one point to another. As a result, there were numerous accidents from trains traveling in both directions on a single track, as it’s really difficult for two trains speeding along a single set of rails in opposite directions to pass each other (something about two objects cannot occupy the same space at the same time).

To prevent these accidents, railroad companies began deploying telegraphs and creating schedules, which helped to delegate authority for the management of day-to-day business. This created divisions and departments that were far removed from central headquarters and were able to handle things in a geographically closer area. The railroads were able to handle the delivery of goods and products, by ensuring timely delivery. In many cases, different products had different priorities. Perishable products had to be delivered in shorter times than dry goods, for example.

The ability to coordinate the delivery of goods, Burke hypothesized, allowed for the creation of the department store. These new stores were able to offer an array of goods to customers, and the products were now available from far and wide, as geographic constraints were removed. These stores started to use the communications and management techniques that were pioneered by the railroads for managing this increasingly complex inventory of products. Since inventory was likely to include items that had a certain level of timeliness attached to them, department stores sought to manage delivery schedules and set priorities for the arrival of particular items. They did this so that they had competitive advantage by being able to deliver popular items before rival stores could. This led to guaranteed delivery that created strong loyalties among big-spending but impatient customers.

As you can see, there is a close parallel between the events that led to the arrival of the department store and the future of internet service providers (ISPs) and ASPs.

Like the railroad companies, ISPs control the means of transportation; the access to the network, if you will. The railroads could only offer access onto the network of rail lines that they had constructed across the country. With the arrival of more finely tuned services (such as express rail, more reliable schedules, high levels of “uptime,” etc.), railroad customers saw limited benefits.

Today’s ISPs also have to deal with a limited set of goods to offer customers. Initially, bandwidth alone was their stock in trade, and that was the norm because no competitors were offering anything but bandwidth, and no customer was demanding more than that. However, as things progress, companies are requiring a larger array of offerings from their ISP.

## Changing the Business

As we look into the deployment of new technologies and how they impact traditional ISPs, it is essential to stress that the Web itself is constantly changing. The future manifestations of the World Wide Web will drive the demand for new ISP businesses such as the ASP model. These new models will drive the changes of the

Web from within. The Web is becoming increasingly frictionless. The Internet is able to spread information to clients that is personalized to meet their specific needs and interests.

In the early use of the Internet as a viable business model, there were no differentiated classes of service to applications. Applications that made money for the company were given the same priority as lesser or nonrevenue-generating applications. In order to counteract this flawed model, the application that generates revenue should have “always on” high-availability status that allows it to meet customer demands. This also allows for the prioritization of different classes for customers, so the businesses that are willing to spend more will gain more robust access than the window shopper.

## The Electronic Economy

The electronic economy has provided ISPs with a challenging and demanding Web environment. In 1994, the Internet was mainly used for the publishing of information and related marketing activities of a company. Now, people routinely use the Internet for information gathering on any and all topics. Many companies (even most nontechnical companies) established a presence on the World Wide Web. It was almost like a validation for their existence in a market that was too vast to understand.

Information in these Web sites about offerings and prices of goods and services allowed for a modicum of stickiness. People started to e-mail each other about various sites, and these spikes in attention created traffic, and so on.

In 1997, a new technology was incorporated into the Web. The ability to perform transactions was introduced and configured to work with vendors' Enterprise Resource Planning (ERP) systems, which provided seamless integration with backend systems. End-users were able to buy products and services through their Web browsers. This was a boon and a pariah at the same time.

In this day and age, very few Web sites occupy markets without some form of competition. This is fairly understandable as there is very little barring multiple sites from inhabiting the same markets. Price points and services are the differentiators for these sites. Online vendors are always looking for ways to draw users to their sites, and keep them there. For the most part, goods were sold from fixed-price lists, which is part of the reason that many of the brick and mortar companies were unable to jump directly into the Web economy. As an example, you may go to a store that is part of a nationwide chain. Depending on where you are, the price for an item may be higher or lower. There are several factors as to why this works, such as geograph-



ical economics and relative need in an area. Most Web-enabled vendors have distribution points that are located throughout a region, which allows them to charge a standard rate for the items that are purchased.

The most recent wave of the electronic economy is the movement to the hosted and managed application model. In these models, the issues of scale and reach are less dependent on location and more dependent on the ability to access business applications in a timely manner.

The electronic economy allowed for greater flexibility in the implementation and monitoring of hosted Web servers and gave ISPs the ability to offer a range of services to their customers. This has a ripple effect in that it will in turn create newer business models, which will spawn even more ripples. The intersection of customer demand and new technologies in the electronic economy will allow for more flexible hosting options that will create the same type of explosion in economic activity that consistent, predictable rail service had on dry-good merchants in the nineteenth century.

## New Opportunities for Service Providers

Internet service providers are finding new opportunities in the hosting of online transaction sites. These hosting opportunities allow ISPs to offer a full suite of online sales and services, by connecting online commerce databases to their clients' core business software applications.

As it stands now, the simple transfer of applications from the intranet to an ISP as an outsourced service is not the only business model for ISPs to pursue. There is also a trend toward a more granular and complex price-for-service matrix that can extend beyond hardware, software, and access packages that are currently offered by ISPs. The ASP will look for more sophisticated prospects in the form of processing power and transaction-per-second service level agreements (SLAs) for their clients. This book will help ASPs to focus their attention on the optimization of their application environment, while competing on a price for performance matrix with other ASPs.

In conclusion, the future of a well-managed and maintained ASP is bright. The ASPs' best days may yet be ahead of them. ASPs will be able to offer better services, retain and grow their customer base, and generate higher margins and profits.

One of the key elements in creating this next-generation ASP will be the ability to extend its offerings in the uncontrollably changing environment of the World Wide Web. The ASP that can react quickly and efficiently to customers' needs will be the ASP that thrives in the coming years.

Just as Burke hypothesized that railroads brought a whole new array of products and opportunities to the retail merchants of the nineteenth century, ASPs will give ISPs an entirely new range of opportunities. ISPs will be the nexus for a wide variety of services, not only for current customers, but also for those seeking ways to compete and survive in the next generation of the Internet.

—Dale Booth, Chief Executive Officer  
EngineX Networks, Inc.



## An Introduction to ASPs for ISPs

### Solutions in this chapter:

- Why This Book Is for You
- Definitions of Common ASP Terms
- The Elements That Make an ASP Viable
- Possible Business Models and Offerings
- Types of ASP Firms
- The OSI-ISO Seven Layer Model
- Choosing the Best Platform for Your ASP
- Business Drivers for the Conversion to ASP
- Performance Issues
- Problems That Could Arise from Conversion
- Major Issues in the Implementation of an ASP Model
- What Is Needed to Sell Your Services
- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

# Introduction

Internet service providers (ISPs) have customarily been suppliers of bandwidth and connectivity to their clients. ISPs are now finding that they are faced with a progressively narrowing margin brought on by fierce competition, with little available differentiation in their offerings. In order to reverse this decline in profitability, ISPs have begun to expand services beyond bandwidth and connectivity, by adding remote hosting services such as outsourcing applications and electronic commerce. ISPs will need to offer new sets of services and tools to manage and reduce costs. These services will give ISPs a distinctive advantage over other ISPs, while creating the opportunity to entice a whole new type of client. The ISP market is not dying; it's undergoing a transformation... Welcome to the world of the application service provider (ASP).

As noted earlier, ISPs were primarily created to supply Internet connectivity (dial-up, dedicated, always on access), electronic mail (e-mail), and Domain Name Services (DNS). With a conversion to an ASP, there is a need to house fully redundant server-based services, high-speed data switching, and load balancing that will allow for greater levels of application service and a superior class of service to their customers.

ASPs (not to be confused with Active Server Pages, also known as ASPs) are starting to appeal to businesses by offering a variety of Web-hosted applications that allow businesses to offload functions they do not want to internally maintain, while operating on their main business strategies.

ASPs provide software programs that include e-commerce, communications, project management, financial, word processing, and human resource applications. ASPs can offer inexpensive use of software, as the price is usually based on a per-use plan rather than a licensing fee. ASPs also allow users to share utilities from multiple locations.

Different types of ASPs can also take care of the complicated functions ranging from Virtual Private Branch eXchange ((V)PBX) systems, Storage Area Networking/Network Attached Storage (SAN/NAS), virtual private networks (VPN), and Network or Remote Operations Center (NOC/ROC) services, as well as a multitude of other services that are covered in this book. Generally, it is more cost effective to leverage these ASP services rather than maintain them "in-house," which will often stretch the capabilities of an Information Technology (IT) staff. ASPs are rapidly developing as the new Internet business in which to be.

Many ISPs are changing their business model to become ASPs. The most successful of these will deliver applications in a secure, highly available infrastructure

that will provide network support, and implementation and maintenance. For a conversion to be successful, you need to understand the different areas of ASP operations, the infrastructure changes that will occur, and the business issues that will present themselves in this new venture.

## Why This Book Is for You

The ASP industry is still in the developing stages of its life cycle, despite all the hype and attention that it has received since 1998. According to the International Data Corporation (IDC), worldwide spending for outsourcing services should reach approximately \$142 billion by the year 2002. IDC also estimates that the application outsourcing (AO) market, which comprises both application maintenance outsourcing (AMO) and application service providers (ASPs), will grow to approximately \$16.2 billion in 2003.

Forrester Research estimates a substantially more aggressive growth pattern for the AO market, stating that it will reach approximately \$21 billion by 2001. This estimate of the AO market is commonly confused with the total income potential of ASP opportunity. According to IDC, the management of enterprise applications by an ASP, a large portion of the total ASP market, is estimated to be \$2 billion by 2003.

The ASP market began capturing the interest and commitment from a large number of venture capitalists and the telecommunications industry in the late 1990s. Some of the industries that established a presence in the ASP market included “pure” ASPs, ISPs, independent software vendors (ISVs), and IT service providers. At the time, the ASP concept was a formidable choice when compared to traditional business models. As a result, many companies formulated strategies for this emerging market.

Early ASPs tended to target small to medium-sized enterprises (SMEs). Forrester Research estimated in 1999 that there were 300,000 emerging enterprises in the United States with revenues between \$40 million to \$500 million, and IT budgets of \$5 million or less. Based on those projections, less than 5 percent of those emerging enterprises in the United States were needed to use an ASP solution to allow it to become a viable market.

Due to the enormous growth and earning potential of ASPs, this book addresses the emerging trends shaping the ASP market and the long-term implications that need to be considered for the service provider industry. We will define the ASP market and its dynamics, and evaluate several business models that can benefit established ISPs and help them grow in the lucrative ASP marketplace.

Network and Web development companies have, until recently, been separating themselves from other project-based companies as a way to gain a higher margin. A major concern that exists in this market is the ability to leverage resources and people who have the skills necessary to manage and maintain these businesses. Since it is hard to find the right people to fill positions and pay them accordingly, outsourcing and IT consulting have become very popular ways to maintain your budget and still have a highly capable staff. The same logic holds true for other resources, such as applications that are too costly to purchase a license and keep your overhead low.

The ASP concept is the advent of a new computing era, with small to medium-sized companies searching for IT alternatives, and a gradual acceptance among larger enterprises.

## What This Book Can Do for You

This book can help you understand the intricacies that are involved in the migration of an ISP to an ASP. Many benefits and pitfalls will be encountered along the way.

With the convergence of software and IT infrastructures, there is a trend toward the Internet or net-centric environment that has enabled the ASP concept to emerge. Software applications have evolved from proprietary, custom-coded applications to prepackaged and net-centric suites. Net-centric software assists in Web-enabled e-commerce, communication, and the management of information.

The IT infrastructure has evolved from a self-contained environment to a distributed computing model and now toward a net-centric infrastructure that links multiple areas of operation. The ASP concept is now attainable due to the availability of relatively inexpensive hardware, efficient communication links, and a robust economy.

Long since past are the days of predictable local area network/wide area network (LAN/WAN) utilization spikes and expected enterprise growth rates. In as such, there will always need to be advances made in software, bandwidth aggregation, and availability to further propel growth in the ASP market.

The following paragraphs provide examples of what this book can help you to be aware of.

As information travels through the network, users act in response; in many cases, in concert. For every action, there is an equal and opposite reaction. Users, by reacting, create what can be compared to traffic jams on Web sites. These traffic jams can overwhelm a site's servers to the point where only a few customers can make use of the application or access the Web site. The experiences of

those who are able to connect are so tainted that they will be reluctant to return to the site.

No amount of planning can truly make your site or application an undefeatable titan, but proper planning can make your service a juggernaut. Due to the frictionless nature of the Internet, there will constantly be a mass of users who will have the ability to trigger hit storms on a much more regular basis. These hit storms will become a way of life on the Internet, and will require technology that can transparently respond to such conditions in a predictable, reliable fashion.

With this book helping you to understand issues such as this, you should be well on your way to making your ISP into an ASP.

## Whom This Book Is Written For

This book will assist the technical executive who either currently runs an ISP or is working with an ISP and wants to know what it will take to convert an ISP to an ASP. This book will help if you are looking for different ideas on how to upgrade your business model as well as your business, and what it will take in terms of investment; types of personnel and timeframes complete the process. The intention is to give the executive a better understanding of what it is going to take to migrate to this new model.

This book will also help the engineer who works for an ISP that is in the process of converting to an ASP model. This book will go into the technical handling of issues that you will need to consider in order to convert an ISP from standard bandwidth provisioning to providing complex services. The objective is to address the technical issues with services covered, and what obstacles, changes, and concerns that will crop up when converting to an ASP.

## Definitions of Common ASP Terms

Here are some working definitions and categorizations for analyzing trends and developments in the ISP-to-ASP industry. These are merely suggestions, as readers often have their own definitions. Pure-play ASP (which is defined later in the chapter) examples are hard to find, so this book will use the following definitions to give perspective in depicting critical developments within the service provider industry.

## What Is an Internet Service Provider?

An *Internet service provider* (ISP) is an organization that provides access to the Internet. ISPs can provide service via modem, or dedicated or on-demand access.



Customers are generally billed a fixed rate per month, but other charges may apply. Some ISPs allow Web sites to be created and maintained on the ISP's server. This along with e-mail allows smaller organizations to have a Web presence with their own domain name. Some larger ISPs also provide news servers, chat environments, and miscellaneous other services (such as Domain Names Services (DNS), among others) in addition to Internet access.

## What Is an Application Service Provider?

The ASP Industry Consortium, an alliance of companies formed to promote and educate the IT industry, offers the following definition: “*An ASP manages and delivers application capabilities to multiple entities from a data center across a wide area network.*”

There are variations of this definition, and sometimes the definition and meanings are confusing. To simplify this definition, an ASP is a third-party service firm that deploys, manages, and remotely hosts a software application with centrally located servers in a “rental” or lease arrangement.

An ASP is a mediator that facilitates remote, centrally managed “rent-an-application” services between a client and an independent software vendor (ISV). The client does not own the application or the responsibilities that are associated with initial installation and ongoing maintenance. The client, through a personal computer (PC) thin client or an Internet browser, can access centralized computer servers that host the application. The client then manages the results from these external applications locally.

### The Pure ASP

The definition of a *pure* ASP is an ASP that joins with a particular ISV, and performs the initial application implementation and integration. In doing this, the ASP manages the data center and provides continuous connectivity and support. The ASP manages client relationships by acting as a complete end-to-end solution provider.

It is possible for an ISV to bypass an ASP and work directly with the client, and it is feasible that another company exists between the ASP and the end user. As an example, Concentric Networks and Exodus Communications manage the data center infrastructure for Corio; this is considered a “pure-play” ASP.

## What Is Information Technology Outsourcing?

*Information technology (IT) outsourcing* is the transfer of an organization's internal IT infrastructure, staff, processes, or applications to an external resource provider.

Outsourcing can encompass anything from the simplest to the most sophisticated IT infrastructure, processes, or applications. Usually, outsourcing contracts are created to handle non-core information technologies or processes.

The outsourcing market can be divided into three main groups:

- Application outsourcing (AO)
- Business process outsourcing (BPO) and information utilities
- Platform IT outsourcing

## Application Outsourcing

*Application outsourcing* (AO) is comprised of ASP and application maintenance outsourcing (AMO), both of which are subcategories of the AO market. The application provider is responsible for the management and maintenance of software applications. The difference between an ASP and an AMO is who actually owns the application.

An ASP remotely hosts and delivers packaged applications to the client from a centralized location. The client is effectively “renting” the application on a per-user or per-use basis. An AMO provides management for proprietary, packaged applications from either the client side or the provider side.

## Business Process Outsourcing

*Business process outsourcing* (BPO) and information utilities providers are primarily concerned with economic and efficient outsourcing for the highly sophisticated but repetitive business processes. These processes can be as complex as accounting and finance, or more recurring processes such as payroll. The provider is responsible for all of the processes associated with the business process.

## Platform Information Technology Outsourcing

*Platform IT outsourcing* offers an array of data center services, such as facilities management, onsite and offsite support services, data storage and security, and disaster recovery. The main differentiation for this type of outsourcing is the transfer of facilities and resources from the client to the provider.

The ultimate intention of an ASP is to allow the client to interact only with the ASP for the services involved. The main elements for this integration are providing the hardware, software, integration, testing, a network infrastructure that is secure, reliable data center facilities, and qualified IT professionals who can manage and maintain these services.

The most critical portions of the ASP channel are the ability to include software vendors, systems implementation, integration, and ongoing support. These components encompass the responsibilities that are necessary to effectively create and administer an ASP solution. These responsibilities help define the development of ASPs. Because of this, there are new opportunities for IT service providers to establish themselves in these markets and still differentiate their service offerings.

An ASP is capable of delivering any type of software application, from e-mail and instant messaging applications to an enterprise resource planning (ERP) system that can manage, control, and report on the multiple facets of the enterprise. The ASP should be able to provide prepackaged applications, support services, and the ability to tailor these packages based on client needs. Generally, the ASP would like to keep these alterations down to a minimum, as customization adds to complexity and the associated support issues. Several of the larger ASPs have publicly stated that there is a lack of customization and they have limited their implementations to core applications. Part of the reason that ASPs do this is because they have negotiated short-term, nonexclusive licensing terms with ISVs, which helps to minimize overhead costs.

## The Elements That Make an ASP Viable

What do you need to check to see if the conversion to an ASP is a viable option to you? There are several factors:

- Is there a reasonable demand either presently or in the immediate future for your possible service offerings?
- Can the model that you plan to use support the possible growth that may be unexpected?
- What can you expect for a return on investment (ROI)?

Several of these questions can be answered by planning the life cycle for the cost of ownership. This is also a good way to gain potential customers, if you can explain that their output would be economically unfeasible, and it would be more cost efficient to use your services.

## Life Cycle for the Cost of Ownership

What are the elements of the life cycle for cost of ownership? This section indicates the items that must be incorporated into the internal cost of ownership

model, and the methodology that is used to determine the values associated with those components.

## NOTE

This is as complete a list as I could come up with, but it is not completely comprehensive. If there are any items that I did not explicitly include, then you should assume that they are excluded from this cost analysis. Things that were not taken into consideration were energy consumption costs, the depreciation of equipment, the variation in administration costs between different environments and/or locations, and infrastructure costs that are considered to be external to the desktop or server

Elements of the life-cycle cost included in this analysis are:

- The initial cost of hardware acquisition
- Hardware maintenance and associated costs
- Initial system software package acquisition
- Initial application software package acquisition
- Implementation
- The cost of hardware upgrades
- The cost of system software upgrades
- The cost of application software upgrades
- Network administration resources
- Other support (training, help desk, etc.)

## The Initial Cost of Hardware Acquisition

This is an average selling price based on common discount levels available for products from value-added resellers (VARs). In some instances, there may be no volume discount applied. Also, keep in mind that there is the possibility to acquire less expensive equipment from local outlets, but you must make sure that they have the same quality of components or the complete package support of a national reseller.

## Hardware Maintenance and Associated Costs

The standard warranty for most vendor equipment is between one and three years, and should cover most issues with problematic gear. The purchase of additional service for the years following the hardware warranty period is added into the hardware maintenance category.

### Designing & Planning...

#### **Hardware Service Contracts**

Hardware warranties differ for each vendor and device in the infrastructure. Manufacturers may offer these extended warranties as part of their purchase plan, or there may be some type of agreement wherein the vendor may cycle in new equipment based on the timeframe involved. You should review what impact these and other service scenarios will have on your business.

## Initial System Software Package Acquisition

The initial purchase price of system software such as a Unix platform or a Microsoft Windows platform and their licensing are considered part of the initial system software purchase. This category could also include software packages that are necessary to run the applications on each machine. For example, WinFrame for Windows Terminals operating system software would fall into this category.

## Initial Application Software Package Acquisition

Initial application software acquisition is any application that assists in the productivity of the organization. This could be an ERP package or some customer relationship management (CRM) suite that will assist the company in management and billing for its applications.

## Implementation

This category represents the cost associated with initial implementation and configuration of hardware and software, as well as costs associated with the ongoing installation of expected upgrades.

Some of the categories you can use to find out the initial implementation cost include:

- The amount of time and resources that are necessary to install and configure the equipment
- The amount of time and resources that are necessary to install and configure the applications for the client base

### NOTE

---

There is also the category of ongoing installation costs that are associated with upgrading the hardware to support anticipated network growth, as well as the application and operating system upgrades that are necessary to maintain the customer's happiness.

---

## The Cost of Hardware Upgrades

Hardware upgrade costs are associated with obligatory improvements to hardware when your company will need to support expanded applications databases, and a more robust operating system. These clients will require these upgrades to grow and improve performance and usability.

### NOTE

---

In order to gauge the cost of the upgrades, you should estimate expected upgrades to the infrastructure for an extended period of time. This period of time will depend on the nature of your application and infrastructure base.

---

## The Cost of System Software Upgrades

Operating system software may need to be upgraded to support newer, vigorous applications. In addition, the software packages that are needed to run your infrastructure may need to be upgraded so that they can handle more efficient management and monitoring tools. This category is based on the platform that you are using and the software suites that you plan to run for standardization.

## The Cost of Application Software Upgrades

Applications are constantly being improved due to customer and client demands. When these packages are in a stable revision, the client will usually insist on the latest and greatest of these application packages. That is where the cost of application software upgrades category is used. A company should try to maintain its return on investment while keeping its client base happy with well-built application packages that can handle all of their needs.

## Network Administration Resources

To determine network administration costs, take stock of what equipment you have, and what it would take in human resources associated with configuring, monitoring, managing, and maintaining the infrastructure.

## Other Support (Training, Help Desk, Etc.)

This category includes help desk support, and training associated with the maintenance of the equipment and applications. This is also the area where other items that were not categorized can be calculated.

## Possible Business Models and Offerings

ASPs host services work on an extensive array of hardware, so at any given time that hardware will have a substantial amount of its processing power idle. The ASP will find that this ability to provision and partition that extra horsepower can be the basis for a very valuable and profitable differentiation service offering. If an ASP can allow its client to be able to respond in real time to hit storms and processing power through management and monitoring software, new opportunities to provide “overdraft protection” for high-traffic applications could be a hot commodity.

The ability to offer different types of service to different types of clients is an incredibly valuable way for ASPs and ISPs to provide granular and real-world service degrees of difference. For example, a machine that generates multimillion-dollar revenues per year is maintained much differently than one that can barely pay for itself.

The more granular control of hosting services will give ASPs and ISPs an opportunity to offer service level agreements (SLAs) that relate to availability of applications in the application-hosting role. The ASP will be able to deploy the necessary “horsepower” to an application in an on-demand fashion, in much the same way electric companies draw additional power from sources through arranged agreements.

This creates the opportunity for the ASPs and ISPs to charge higher rates for on-demand horsepower to applications, allowing them to sell services and guarantees in much the same manner that public utilities do. In doing this, an ASP and ISP are able to realize a demand-based pricing model, which will afford the company higher margins.

An added benefit of this higher granularity is that it will drive down the cost of ownership for service providers, as hardware, software, and administration costs will drop based on the greater flexibility that these organizations are able to offer.

Remember that there is no such thing as an ideal solution, so you want to offer good service, at good rates. For example, it is not considered economically viable for service providers to over-provision their services and computing power as a way to stop hit storms and traffic jams. There is no way to justify the expenses involved with over-provisioning, as there will be too much idle time on the infrastructure when these hit storms aren't plaguing the network.

The ideal solution would be able to deliver overload protection while classifying users and their levels of access without forcing over-provisioning. This would provide protection against unforeseen client demands and high availability in extraordinary conditions.

Classification of services for both users and applications will give service providers new pricing models. Instead of flat fees for bandwidth, service providers be able to sell “units of work” to different types of clients. They will then be able to prioritize and price such units based on user privilege, application priority, and time of day.

This model is able to distinguish that applications are not equal and that the network environment is not democratic. These solutions require the deployment of controls over system policies and how they handle classes of applications as well as clients. This is the area where the user experience is defined. With these tools and policies in place, service providers will be able to offer Quality of Service (QoS) options that will also create new pricing model opportunities.

## Types of ASP Firms

There are several types of ASP-enabled firms. These organizations can be separated into professional consulting, project-based service providers, outsourcing providers, staff augmentation providers, education and training providers, and value-added resellers.



## Professional Consulting

Professional consulting firms focus on corporate-level business and strategic engagements. This can be broken down into three subcategories:

- **IT consulting** These firms focus on high-level consulting projects that are aimed toward strategic IT engagements. These project scopes often entail a companywide evaluation of client needs, essential processes, existing platforms, available technologies, and design solutions. What separates these companies from their project-based service provider counterparts are their efforts to structure their IT initiative as the main component of a strategic business process design.
- **Strategic management consulting** These firms provide advice that centers on the client's corporate objectives and competitive position. The strategy to this method is the creation of a unique positioning that allows for sustainable advantages. Project scope usually involves topics such as market trend analysis, marketing efforts, business and customer mix, and capital structure.
- **Business process consulting** These firms provide consulting expertise for operational effectiveness at either the functional or business unit level. Their best practices and processes will allow a company to use its resources more effectively, so that it can generate the highest operational effectiveness at a minimal cost.

## Project-Based Service Providers

Clients that select these providers for projects are opting for well-defined tangible deliverables and scopes. Contract designs range from a billable-hours approach to fixed-price engagements for components and entire projects. These companies focus on industry expertise, either in specific technologies or industry applications.

- **Application and systems development** These companies specialize in the customization and software development that handles specific needs of clients, and for proprietary systems. Some of the deliverables include modules or components, upgrades to existing systems, and original application development.
- **Integration and implementation** These firms focus on the deployment of complex enterprise software packages such as ERP. For implementation, these companies will integrate the new software, ensuring

that hardware, network, and software components work together. These companies also specialize in integration technologies, interface development, database management, and other technologies that enable dissimilar systems to share information.

- **Network and Web development** These businesses develop client/server and Web-enabled technologies that link businesses together through LAN and WAN facilities, as well as other Internet-based solutions. The projects may require vendor and supplier management, customer communications, overflow management, billing, and receivables.

## Outsourcing Providers

Outsourcing providers are organizations that provide process automation services, facilities management, and operations for clients who require an assortment of technical answers and can be divided into three categories:

- **Application outsourcing** These companies manage and maintain software applications, and assume the responsibilities associated with these applications. AO can be further be subcategorized:
  - An ASP remotely hosts and delivers packaged solutions to clients from an offsite location.
  - An AMO provider manages proprietary, packaged applications from either the provider or client's site.
- **Business process outsourcing (BPO) and information utilities** These firms focus on economic and efficient outsourcing solutions for multifaceted but tedious daily business procedures. The provider assumes responsibility for the business process.
- **Platform IT outsourcing** These organizations offer a range of data services that include hardware facilities management, onsite/offsite support, data security, and disaster recovery competence. These contracts typically involve the transfer of IT staff and/or resources.

## Staff Augmentation Providers

Staff augmentation organizations specialize in providing IT professionals, on a temporary or long-term contract basis, to clients who need specific skill sets and support for internal systems and development projects.

- **Pure IT staff augmentation** These firms create the majority of their revenues from their core IT staffing business. These companies use strategies that are usually defined by geographic location, expertise, or technology focus.
- **Transitioning firms** These companies are traditionally viewed as being in the IT staffing business, but have attempted to redirect their business model toward higher value-added and higher-margin projects, through amalgamations of mergers and acquisitions, divestitures, and internal growth.
- **General staffing** These services provide professionals with a vast array of skills that can include finance, accounting, network design and engineering, and so forth, and have an IT staffing division that draws significant revenues. Many companies in this sector are building professional services divisions, through internal growth and by acquisition.

## Education and Training Providers

Education and training companies provide training and help desk consulting to firms that have implemented custom-designed or packaged software products. These services can include onsite training or off-campus programs following new installations, or as a skills development seminar for certain technical applications.

## Value-Added Resellers

Value-added reseller (VAR) organizations are solution-oriented vendors who can provide integration for hardware and software systems. These firms usually include consulting, design, and implementation services. Traditionally, these VARs have operated under specific hardware and software vendors, due to agreements, though recent trends are toward vendor-neutral representation.

## The OSI-ISO Seven Layer Model

One of the things that enabled the rapid growth of the Internet is the adherence to the *Open System Interconnection* (OSI) communication model. The model allowed companies to create pieces of software and hardware that could easily integrate along each of the layers. With this, companies were able to add incremental value, as they were able to work on their strengths and create packages that now allow service providers the ability to differentiate themselves.

This book will use the seven-layer OSI model, as well as three added pseudo layers, to describe what's happening in each of the chapters. This will assist in the understanding of the topics and where they relate in the scheme of network infrastructure. Throughout the book, there will be references to these OSI layers and how they pertain to the topics covered.

The OSI reference model is a conceptual model composed of seven layers, each specifying particular network functions. The International Organization for Standardization (ISO) developed the model in 1984, and it is the primary architectural model for intercomputer communications. The OSI reference model describes how information from a software application in one computer moves through a network medium to a software application in another computer.

The OSI model divides these communications involved with the moving of information between networked computers into seven smaller, more manageable layers. Each of these layers is self-contained, so that the responsibilities assigned to each can be implemented independently from other layers. This enables the results offered by one layer to be implemented without adversely affecting the other layers.

Information that is being transferred from a software application in one system to a software application in another must pass through each of the OSI layers.

The following list details the seven layers of the OSI reference model:

- Layer 1—Physical layer
- Layer 2—Data Link layer
- Layer 3—Network layer
- Layer 4—Transport layer
- Layer 5—Session layer
- Layer 6—Presentation layer
- Layer 7—Application layer

## Layer 1: The Physical Layer

The *Physical layer* defines electrical, mechanical, procedural, and functional terms for activating, maintaining, and deactivating the physical link that exists between networks. Physical layer specifications define voltage levels, timing of voltage changes, physical data rates, physical connectors, and maximum transmission distances. Physical-layer implementations can be categorized as either LAN or WAN specifications.

## Layer 2: The Data-Link Layer

The *Data-Link layer* provides transmission of data across a physical network link. Data-Link layer specifications define network and protocol characteristics, such as physical addressing, network topology, error notification, sequencing of frames, and flow control. The physical addressing (in contrast to network addressing) defines how devices are addressed at the Data-Link layer.

A network topology is comprised of Data-Link layer specifications that define how devices are physically connected, as in a bus or a ring topology. Error notifications alert the upper-layer protocols that a transmission error has occurred, so that the data frames can be reordered to transmit in sequence. Flow control monitors the transmission of data so that the receiving device is not inundated with more traffic than it can handle.

## Layer 3: The Network Layer

This *Network layer* provides routing and functions that allow multiple data links to be combined into an internetwork. This routing is accomplished by the logical addressing (as opposed to the physical addressing) of network devices.

The Network layer supports both connection-oriented and connectionless service from higher-layer protocols. Network-layer protocols are typically routing protocols, but other types of protocols can be implemented at this layer as well.

Some of the common routing protocols include external protocols such as Border Gateway Protocol (BGP), or internal protocols such as Open Shortest Path First (OSPF) (which is a link-state protocol developed for use in TCP/IP networks) and Routing Information Protocol (RIP) (which is a distance vector protocol that uses hop count as its metric).

## Layer 4: The Transport Layer

The *Transport layer* creates reliable internetwork data transport services, which are transparent to upper layers. Some Transport layer functions include flow control, error checking and recovery, multiplexing, and virtual circuit management.

Flow control also manages the data transmission between devices so that the transmitting device does not send more data than the receiving device can process. Error checking involves the creation of mechanisms that can detect transmission errors, while error recovery involves the retransmission or the requesting that data be retransmitted, to resolve errors that may have occurred.

Multiplexing enables data from multiple applications to be transmitted through a single physical link. Virtual circuits are established, maintained, and terminated by the Transport layer.

## Layer 5: The Session Layer

The *Session layer* establishes, manages, and terminates sessions between Presentation-layer (Layer 6) connections. These sessions consist of requests and responses that occur between applications, which are located on different network-enabled devices. These requests and responses are then coordinated by protocols implemented at the Session layer.

## Layer 6: The Presentation Layer

The *Presentation layer* provides the ability to convert functions that are applied to Application-layer data. These functions help to ensure that the Application layer of one system will understand information sent from the Application layer of another system. These conversions include the conversion of character representation formats, common data-compression schemes, and common data-encryption schemes.

A common data-representation format allows the usage of standard images, sounds, and video formats, enabling the interchange of application data between different systems. Using different text and data representations, such as EBCDIC and ASCII, systems can exchange information. Standard data-compression schemes enable data that is compressed at the source device to be decompressed at the destination device. Standard data-encryption schemes allow data that is encrypted at a source device to be properly decoded at its destination.

Presentation-layer implementations are generally not associated with a particular protocol. There are some standard programs for video, which include QuickTime and Moving Pictures Expert Group (MPEG). QuickTime is the Apple Computer specification for video and audio, and MPEG is a generic standard for video compression and coding. Some well-known graphic image formats include Graphics Interchange Format (GIF), Joint Photographic Experts Group (JPEG), and Tagged Image File Format (TIFF). GIFs, JPEGs, and TIFFs are the standard for the compression and coding of graphic images.

## Layer 7: The Application Layer

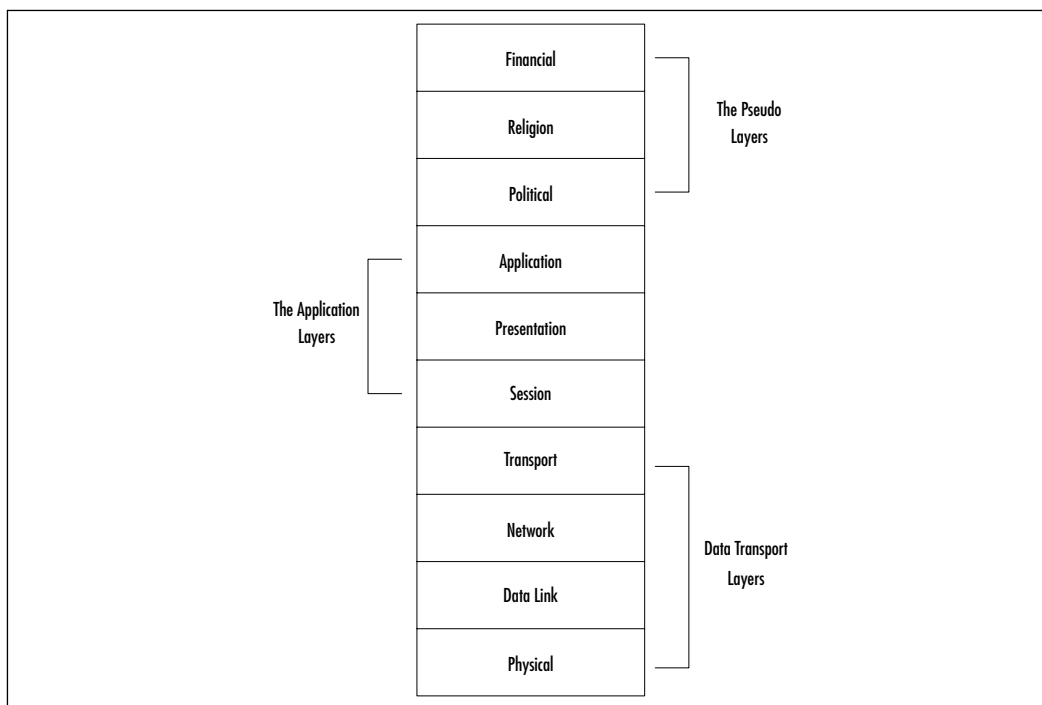
The *Application layer* is the OSI layer closest to the end user. This means that both the OSI Application layer and the user interact directly with the software application.

This layer interacts with software applications that implement communication components. These application programs are considered outside of the OSI model. Application-layer functions typically include the identification of communication partners, determine resource availability, and the synchronization of communication.

To identifying communication partners, the Application layer determines the identity and availability of the communicating parties for an application with data to transmit. When determining a resource's availability, the Application layer must make a decision as to whether there are sufficient network resources for the requested communication to occur. The synchronization of communication requires cooperation between applications that are managed by the Application layer.

Any layer in the OSI layers is able to communicate with three other OSI layers: the layer that is directly above it, the layer that is directly below it, and its peer layer in other networked computer systems. For example, the Network layer (Layer 3) in System A can communicate with the Data Link layer (Layer 2) of System A, the Transport layer (Layer 4) of System A, and the Network layer (Layer 3) on System B (Figure 1.1).

**Figure 1.1** The OSI Seven-Layer Model



The seven layers of the OSI reference model can then be further subdivided into two categories: the upper layers and the lower layers.

## The Upper Layers

The *upper layers* of the OSI model handle application issues and are generally implemented in software. This upper layer, also known as the Application layer subcategory, is comprised of the Application layer (Layer 7), the Presentation layer (Layer 6), and the Session layer (Layer 5). The users and Application layer interact with software applications, which contain a communications component. The term *upper layer* can also be used to refer to any layer above another layer in the OSI model. For the purposes of the book, the Application layer subcategory will be differentiated from the Application layer (Layer 7) with the layer noted in parentheses.

## The Lower Layers

The *lower layers* of the OSI model are also known as the Data Transport layers. The Physical layer (Layer 1) and Data-Link layer (Layer 2) are implemented in hardware and software. Layers 3 and 4 are implemented only in software. The Physical layer (Layer 1), is closest to the physical network such as network cabling, and is responsible for actually placing information on the medium.

## The Pseudo Layers

These *pseudo layers* are not actual OSI model layers, but they will directly influence the way in which you will implement your equipment and policies. Each of these layers—Political, Religion, and Financial—will affect how, what, and when you implement your infrastructure. The pseudo layers are also shown in Figure 1.1.

## Layer 8: The Political Layer

The *Political layer* is the first of the pseudo layers. This layer is where companies implement policies and procedures. This is also one of the barriers that must be taken into account for the successful deployment of your internetwork infrastructure. Depending on the structure of the company, a network design and implementation may have to go through numerous panels or committees in order to pass companies' stringent needs. This consumes time, and therefore slows the implementation process.



## Layer 9: The Religion Layer

The *Religion layer* is based on the unswerving loyalty that a client has to one specific customer. This could be a good thing or a bad thing, as it tends to lock the client into a rigorous mold. In some cases, it leads to an end-to-end solution that is easily managed and monitored. In other cases, it may not allow for the “best of breed” equipment that will benefit the client in speed and functionality.

Both sides have their benefits as well as drawbacks. There is no one correct way to go about implementing your infrastructure. Take, for example, the deployment of company C’s equipment. They offer name recognition, a strong market presence, and the ability to offer a client a complete solution from soup to nuts. Many companies see these factors and purchase company C’s equipment exclusively, even if there is a competing product on the market that could do the required task in a faster, more efficient manner.

Then there are some clients, who are referred to as ABCs (anything but C) and will not use company C’s equipment. These companies will go out of their way to find equipment that will work “at least as good as company C,” whether they state it that way or not. These companies look for the newest technologies that are hitting the market and incorporate them into their infrastructure. This usually allows for great speed and power, but may take away from the manageability and interconnection of the network.

## Layer 10: The Financial Layer

The *Financial layer* could be the most disruptive and least technologically controlled area of the layer model. This layer deals directly with outside factors on the purchasing and deployment of a company’s network infrastructure. Based on a number of cost factors, companies may not be able to implement the most efficient or robust solutions. As an example, due to a large cash output in the front end of a network build-out, a company may look to piecemeal an infrastructure deployment, or use substandard products. This becomes more costly in the long run as the support bills will mount and the customer satisfaction levels will decline.

## Choosing the Best Platform for Your ASP

ASPs take advantage of existing Internet connectivity to offer corporations the opportunity to outsource not only peripheral applications but also mission-critical applications. This trend will continue to escalate as customers discover that

outsourcing firms can deliver mission-critical applications that meet their demands for service level agreements (SLAs). For ISPs and ASPs, these application-hosting responsibilities require the choice of a platform that can deliver the correct balance of performance, scalability, upgradeability and manageability. This section is covered more comprehensively in later chapters of the book.

## Hardware

Early in the 1960s, an IBM mainframe salesperson identified a need among mid-sized companies. They needed the same types of computing capabilities and applications that were used by larger companies, but they lacked the resources to buy and maintain these complex mainframe systems and applications. There was born the idea to lease unused mainframe compute cycles and application space from larger customers, to sell as outsourced data processing to these mid-sized companies. H. Ross Perot had founded Electronic Data Systems (EDS).

The current Internet movement toward application service provisioning is an array of both old services and new functionality that sits on top of a new sleeping giant, the Internet. Virtual private networks (VPNs), intranets, and extranets have been around for some time, and their value proposition remains the same, but the means of their implementation have changed.

ASPs usually integrate software, hardware, networking technologies, and related IT consulting services into an outsourced platform. The ASP then leases these services to its consumers, who typically enter into multiyear contracts. Charges are usually based on the number of users or transactions. These services tax the ASP; so one of the major challenges that faces an aspiring ASP is the choice of the right hardware platform.

In an effort to meet its customer demands, an ASP must have a hardware platform that will support clustering, fail over at the data link, application, network node, and system level, and the maintenance and response times of the applications themselves. The ultimate solution to meet an ASP's needs requires the following capabilities:

- There must be an excessively high level of performance to meet both current and future needs.
- It must have a high degree of scalability that will support the ability to add, interchange, and evolve components and subsystems when necessary, while offering a high return on investment (ROI).

- It needs to have high availability (HA) features and redundancy, which allow for the ability to reconfigure, repair, and replace components without requiring system downtime.
- It should be able to adapt effectively and include load balancing to handle peak traffic loads, and prevent excessive system overload to avoid performance degradation.
- The platform should allow the ability to put a lot of power (processing and throughout) in a small space.

## Servers

Several useful criteria assist in evaluating hardware platforms for ASP applications. Performance is probably rated very highly. However, it is almost impossible to compare vendors against each other, since product evaluations are often based on what each company's strong point is and how it is marketed. Relative comparisons can provide a frame of reference as long as you keep in mind that your needs might affect the ratings.

Companies cannot ignore the price-to-performance value proposition. Therefore, this book focuses on a trio of mid-level symmetrical multiprocessing (SMP) systems, whose offerings can be clustered to performance, scalability, as well as present a high degree of investment protection:

- Hewlett-Packard (HP)
- Sun Microsystems
- Intel based (Compaq, IBM, etc.)

HP and Sun usually are implemented with the more performance-driven Unix operating system solutions, while Intel-based servers are more suited to (but not exclusively) Microsoft Windows NT solutions.

## Hewlett-Packard

Hewlett-Packard's (HP) servers offer a range of service-provider class systems. HP's midrange N-Class systems, and its high-end M-Class systems, are available with multiple, high-speed RISC processors, and can be upgraded to faster processors as they become available.

HP also includes its WebQoS software with its M-Class and N-Class server packages. WebQoS identifies and prioritizes access to Web-based applications and

services according to the type of customer, the application requested, and the particular combination of rights and privileges determined by business policy.

## Sun Microsystems

Sun has a good reputation with ISPs and more recently ASPs, and is likely to be on most ASP networks. In Sun's midrange server line, the Sun Enterprise servers have a fantastic price-to-performance ratio, although Sun has targeted these systems as departmental-level servers. Sun servers have the ability to support external Sun StorEdge products, which helps in their scalability and high availability.

## Compaq

The Compaq ProLiant series is the pinnacle in Compaq's Intel-based high-availability line. It uses Intel microprocessors and can support multiple processors in an SMP configuration. More information on SMP can be found in Chapter 3.

## Network Equipment

Because service providers must be flexible to handle new market demands when clients require new applications in a timely and efficient manner, their business strategies and infrastructures must be able to change quickly. Generally, these new applications require a more flexible, adaptable infrastructure to support them, so in order to remain a leading service provider, you will need to aggressively pursue the Internet and the opportunities it offers.

## Data Traffic Explosion

Traditional ISPs are experiencing an explosion of data traffic on their networks. The Internet and its dramatic growth have fueled the need to satisfy this increasing data demand by forcing the migration towards multiservice network platforms.

The convergence of data, voice, and video traffic onto a packet-based infrastructure, coupled with global deregulation, has given rise to two events:

- Telecommunication providers need to move to a data-based infrastructure, while recovering their huge initial investments in legacy-switched voice technology. What they are finding is that they are struggling when trying to push data traffic across these less efficient circuit-switched networks. In order to compete, traditional carriers are forced to reduce prices, create a migration strategy to IP, purchase and build-out new IP networks, or identify and create differentiated services to lock in customers

- ASPs have a unique opportunity to enter this growing market by building out more efficient (and hopefully less expensive) IP-based networks that can offer data transport services at a fraction of the price that telecom providers can. These providers are also looking to create value-added services that will lock their customers into their distinctive suite of services.

Since implementation of specific types of equipment are discussed in later chapters, I have included a brief overview or history of several of the larger networking companies. These companies consist of Alcatel Networks, Cisco Systems, Extreme Networks, F5 Networks, Foundry Networks, Juniper Networks, Lucent Technologies, and Nortel Networks. This is not a complete list, as other vendors and new players are coming into the market space.

## Alcatel Networks

Alcatel adapts to its customers' needs by developing a complete offering of innovative products. Alcatel can create communication networks by combining the necessary components such as network systems, access, and transmission solutions that can be tailored to businesses and end users. The company can handle anything from basic telephone services to the most complex multimedia networks. They have expertise in network integration and management, which enables them to provide "end-to-end" services. Alcatel can provide integrated private networks, and basic and advanced services such as facilities management.

## Cisco Systems

Cisco Systems is the leading supplier of networking solutions worldwide and a leading vendor of IP-based solutions. Cisco Systems has positioned itself to help service providers in increasingly competitive environments. Cisco is also a recognized leader in the design, development, implementation, and operation of Internet business solutions.

## Extreme Networks

Extreme Networks' broadband solutions can adapt quickly and transparently to handle rapid provisioning of new IP services and applications across service provider networks. Extreme arms service providers with an "end-to-end" solution that can clearly differentiate them from the competition. Because broadband services are only as fast and reliable as the infrastructure on which they run, some of

the world's largest service providers rely on Extreme Networks for their business and mission-critical applications and services.

## F5 Networks

F5 Networks, Inc. is a pioneer in the field of high availability, and Internet traffic and content management. F5 develops solutions that increase the availability and performance of IP-based servers and network devices such as firewalls, routers, cache servers, proxy servers, and more. Their products act essentially like air traffic controllers, with the ability to reroute client requests to the most available server and direct requests away from servers not responding.

## Foundry Networks

Foundry Networks is a performance and total solutions leader for “end-to-end” routing and switching solutions. Their portfolio includes Internet routers, Layer 2/3 LAN switches, and Layer 4-7 Internet traffic and content delivery switches. They are implemented in some of the world's premier ISPs and enterprises, portals, search engines, e-commerce sites, and universities.

## Juniper Networks

Juniper Networks Internet backbone routers are built to meet the unique challenges of the new IP infrastructure. These routers are used in critical applications such as core management, dedicated access, peering, and data center hosting by the world's largest service providers. One of the benefits to running a Juniper solution is the fact that all of their platforms share common JUNOS Internet software.

Designed specifically for the specialized needs of high-growth, high-speed service providers, Juniper routers enable you to scale networks to meet your bandwidth and service needs. These platforms perform at wire rate over the most flexible combination of interfaces with the highest port densities per rack inch on the market today.

## Lucent Technologies

Lucent's Full Circle Program delivers services in an open, converged communications network. This allows developers to implement multimedia applications that can operate on a variety of networks and platforms. Service providers can choose from a collection of merchants who offer open-network-based products and services that they need to differentiate themselves and generate revenue.

Full Circle is a standardized, open, programmable network architecture with an open application-programming interface (API). This technology allows for open development environments that will provide distinctive and potent capabilities.

## Nortel Networks

Nortel Networks' Internet Data Center Solutions support service provider requirements, such as fast time to market, efficient network and service management, and the ability to grow yet remain flexible. All Nortel components are built to support the critical need for security in today's eBusiness environment. Nortel Networks' hosting infrastructure solutions are designed to enable high-performance Web hosting services. Nortel Networks can design, assemble, and integrate the products for an "end-to-end" solution, and also offer a full range of services including site surveys, facilities build-outs, design, implementation, and even facilities operation.

## Cache Appliance Makers

Several appliance makers service this area of the market, including Akamai, Intel, and Inktomi.

### Akamai

Akamai's content delivery solutions can substantially increase Web site performance and reliability by solving one of the major problems on the Internet: Web congestion. Akamai delivers content and applications from the edge of the Internet. Their technology is scalable and allows you the freedom to create more dynamic and profitable Web sites.

By providing the ability to download faster, and create greater site reliability and content customized to your unique users, Akamai solutions can deliver positive user experiences that will increase traffic, ensure customer loyalty, reinforce your competitive position, and increase the power of your brand.

### Intel

Intel makes the Intel NetStructure Cache Appliance that can store, retrieve, and serve not only Web pages, but also pieces of Web pages, and can thus provide optimum bandwidth savings.

### Inktomi

Inktomi makes an appliance that they call the Inktomi Object Store. This package is a custom-designed Web object database that has been fully optimized for

caching. It uses raw disk input/output to achieve optimal storage and retrieval of content, which results in much higher speeds than can be obtained by conventional file systems. The way that this is handled is that the most frequently requested objects are cached in random access memory (RAM) and are maintained so that hot objects can be read from high-speed memory instead of from disk. All objects in the content are indexed according to their URL and associated headers for faster access.

## Software

Reliability is one of the most important considerations to make when choosing a server platform, but it is likely that ASPs should be most concerned about the operating system. This often boils down to a choice between some form of Unix, Linux, or Microsoft Windows platform.

On the high end of application hosting, Unix seems to be the system of choice. Unix is primarily a 64-bit architecture that is already way ahead of the 32-bit Windows platform in terms of scalability. However, because there are so many companies that are Microsoft centric, it looks like NT is gaining in the low- and mid-tier application server space. With a number of business class applications now available on NT (and Windows 2000), it is very likely that IT departments will use it to outsource at least some of their programs.

Most ASPs are already working in a mixed environment where both Unix and NT are used. It is not hard to imagine a time when both platforms will be deployed as a single integrated solution.

## Load-Balancing Software

Meeting the varied needs of different classes of users will push software and hardware load-balancing requirements for the ASP model. With the ability to deliver Internet applications predictably to a nearly infinitely elastic user community, there is no limit to the potential growth of an ASP.

The ubiquitous need for such solutions is readily apparent, as both NT and Unix server schedulers work in a fair-share model, meaning that they work on a first-come, first-serve basis.

There is one way to meet and exceed every demand that is made of a network, and that is to throw an unlimited amount of hardware at it. This is both expensive and unnecessary. It is better to leverage load-balancing technology in which traffic requests are routed to different servers based upon:



- Who is the user?
- What application is being requested?
- What priority level is assigned to the user?

Both HP and Sun created have load-balancing software that allows companies to identify and prioritize access to Web-based applications and services. This prioritization can be based on peak usage, management, monitoring, user policy, and service policy.

HP created the WebQoS product and introduced it in 1998. HP has recently introduced updated versions of its WebQoS platform. In 1999, IBM created class-of-service differentiation with its revision of WebSphere Performance Pack. Several smaller companies such as Atrieve (purchased by Inktomi), BEA Systems, Bluestone (purchased by HP and integrated with their WebQoS software), Bright Tiger (purchased by Allaire and later merged with Macromedia), Hewlett-Packard, IBM, Microsoft, Resonate (created an alliance with Inktomi), and WebManage (purchased by Network Appliance and integrated with their storage solutions) created Java agent-based tools and management systems that assisted in Web application class-of-service differentiation in load balancing.

## BEA Systems

BEA Systems built a solid reputation as the creator of large e-commerce sites such as E\*Trade and Internet banking sites. Their ability to leverage dynamic load balancing and security are integral parts of the BEA Web server deployments. Its class-of-service management software is more solution specific, rather than wide ranging and general purpose, and it requires that application developers create BEA-specific versions of their products.

As an example, if a trading application is under a high load, BEA software will reroute stock quote requests to a server with a lower priority in a server farm, and actual stock trades will be routed to the higher-priority servers. Its load balancing is application specific, and it assumes that there is an adequate array of server farms, instead of having a “safety valve” that would channel the excess requests to a server that is considered on-demand.

## Hewlett-Packard WebQoS

Hewlett-Packard’s WebQoS was designed for businesses with high-volume Web sites that have volatile demands. WebQoS is server-based software that isn’t a true load balancer. It manages periods of peak demand, resulting in larger and more completed transactions.

WebQoS delivers a reliable and controllable infrastructure that allows ASPs to aggressively leverage the Internet to their advantage. WebQoS allows for the management of your Web site and how it performs for customers, resulting in higher customer satisfaction and loyalty, higher revenues, lower risk of site failure, and reduction of support costs.

This package is actually a product line that encompasses the ability to control peak usage or hit storms, while providing differentiation of performance levels based on a per-user, per-application basis. It is able to monitor and control throughput and availability of Web applications by maintaining a fine-grained security model that can exchange information with other models that already exist in the market. One of the main selling points for this package is that it can solve the interference problems that are associated with hosting multiple applications on the same server.

HP offers a package aimed at ASPs that has a built-in service resource management facility. This ensures that one application on a multihosted server does not consume all of the resources at the expense of performance from another application on the same server. This also allows ASPs to set, monitor, and manage application performance and scalability requirements. By allocating resources to applications as a service thread with an overall performance level that can be written into a guaranteed contract, an ASP will be able to garner more customers.

## IBM

IBM addresses two of the aspects of privilege-driven computing with this package. The package can load balance and offer tight security at the same time. The software implements load balancing for firewalls, Web page caching servers, and e-mail servers across Windows NT, AIX, and Solaris servers.

WebSphere has two security models:

- The AFS file management that is used in conjunction with Kerberos access control for low-security applications
- The DFS model that is based on the DCE (Distributed Computing Environment) middleware package for high-security

While it does address load balancing and security issues, the IBM WebSphere software also address the differentiation of classes of application service. However, there is currently no “safety valve” that would enable a site to recognize a situation in which the entire server farm is overloaded so that it could channel excess client requests to an outside server on demand.

## Designing & Planning...

### WebSphere Commerce Suite, Service Provider Edition

WebSphere Commerce Suite Service Provider Edition (SPE) permits service providers to deliver a range of store models to meet their customers' growing needs. This package offers two templates with which the customer can work.

- **Basic Store** offers a Store Creation wizard that can create a fully operational e-store that contains populated product directories and includes offline or online payment processing.
- **Advanced Store** offers unrestricted-sized catalogs, a catalog editor, and the customer can design store flows and have the ability to include custom-designed HTML pages. Tax and shipping calculations can be defined at the product category level and be based on particular jurisdictions. SPE provides the ability of snap-on-commerce, which allows current Web sites to be commerce enabled with Buy buttons, shopping cart catalog search, and customer service functions. This product package will meet the needs of Web designers, developers, and administrators for inexpensive, creative, and rapid application development, integration, deployment, and maintenance.

## Microsoft

With its Wolfpack applications added to Windows NT, Microsoft has created the fundamentals that are necessary for the development of load-balancing, Web-based applications. Through thread management, and the possibility of failover and high-availability deployment schemes, an ASP could create homegrown tools to handle its internal business applications. The ability for the Web application service class management to guarantee Web application performance is not currently addressed in this package or by third-party development companies.

## Resonate

Resonate has two software packages that can be used in the software management and monitoring realm. The first package is designed for the WAN and is called Resonate Global Dispatch. It is a multisite, service-level solution that

provides high availability and optimal performance for geographically dispersed applications. It enables points of presence (POPs) to act as a single system and routes user requests to the site that is best able to handle the client's requests. It can also redirect users to the POP that is closest to their location to save on WAN costs or to enhance users' experiences by routing them to the site where the content that is relevant to their needs is located.

It can synchronize information with other Resonate products that are deployed within an ASP, locally and globally, so that you can ensure that service levels are maximized across your locations.

Their second package, Resonate Commander, monitors the status and health of the multiple layers that make up these sophisticated applications, and then automatically takes the appropriate action, in real time, to ensure maximum service levels. It allows the client to define thresholds for taking action to prevent problems from occurring. Commander takes immediate action to ensure that traffic is rerouted around the potential bottlenecks, so that the user maintains a positive experience.

This product also provides historical statistics so that you can do trend analysis and capacity planning, and can synchronize information with other Resonate products deployed within an ASP, much the same as the Global Dispatch package.

## Segue

Segue Software picks up the Web application measurement idea where Microsoft leaves off. Segue's Silk family of products are a set of end-to-end application testing tools for the functionality and testing of distributed Web applications. Several models of the Silk product line do support load balancing, performance testing, Java Virtual Machine environments testing, automated program defect tracking, database access and verification testing, Web link, and Web page load testing.

Segue technology places intelligent agent applets across the layers of a distributed Web application for reporting purposes, and is primarily an application development and debugging tool for the programmer. It does not allow for real-time application response management in an automated runtime environment, so it is unable to provide Web application management and service class distinction in a deployed application.

# Business Drivers for the Conversion to ASP

The ASP Industry Consortium, founded by 25 technology companies in May 1999, was developed with the purpose of promoting the ASP industry to educate the marketplace by developing common definitions, facilitate industry discussion, stimulate research, and encourage open standards by promoting industry “best” practices. Since that time, other companies have joined the ASP Industry Consortium with membership of more than 120 companies by the end of 1999.

## Business Factors That Impact the ASP Model

The new value proposition that is offered by various services using Internet and intranet technologies is the ability of the ASP to free its customers from having to develop, maintain, and provide services for themselves. It helps to establish customer independence from the types of hardware that are required to run the outsourced applications. Usually, all the client needs is a browser to use the software that the applications require, and not the specialized hardware and servers.

There is also the added benefit of using an ASP’s application management expertise. Over the lifetime of an application, such as Enterprise Resource Planning (ERP), an ASP can estimate that software licensing, hardware and basic infrastructure costs will account for less than one-fifth of the total cost of ownership (TCO) over a five-year period. The remaining four-fifths is consumed in software management and human capital overhead that is associated with the application. An ASP can more effectively use its resources to its advantage in this arena with the ability to provision multiple clients.

For clients, these outsourcing applications and services represent a palpable savings of time to market, and more effective ways to use IT personnel. It is drastically cheaper to use the services of an ASP than maintain similar services themselves. Companies that lack the technical means to deal with new technologies, software architectures, and the rapid release of updates to products regard the availability of ASP as a savior.

A company can focus more on the things that will make it successful by outsourcing tasks that are not part of its core competencies. An ASP can leverage its personnel, resources, and expertise to implement applications in a fraction of the time it would take the customers’ organization.

These same companies can alleviate the burden of buying expensive, management and maintenance intensive, rapidly obsolete hardware. By leasing its services, ASPs save companies the substantial cost of assembling, monitoring, and

supporting these computer systems. The company also absconds itself from the responsibility to either create or maintain the specialized software that is usually associated with high-end applications. Using an ASP has the following advantages:

- It helps minimize the TCO. By using an ASP, a company can typically factor in a 30-percent to 50-percent annual savings, depending on the complexity of the application.
- It can allow for cash flow that is more predictable. There can be a degree of predictability by eliminating the uncertainties of after costs and software-related expenditures, as the ASP usually mitigates these issues.
- It allows the company to focus on their core competencies and strategic planning. The transfer of the implementation and management of an application to a third-party helps the company to focus on developing its core aptitudes.
- It helps improve internal IT staff. By eliminating application management, the company is able to help the IT staff develop processes and systems, and leverage core competencies.
- It also can improve coordination efforts on a global scope. The ASP model helps organizations use the latest tools and systems that can coordinate internal and external global business.

## Enabling Technologies

The reach of intranets and the Internet to virtually every desktop in a company is creating a crisis for organizations. On the positive side, it is bringing all users into the shared network, but it also increases expectations for what the network is able to accommodate.

Companywide access creates massive amounts of stored data, pushing the need for access to legacy data stores. All of this information, once it is gathered, assembled, and stored, is considered useful to users across the enterprise.

Companies are faced with the chore of managing, organizing, and distributing numerous forms of data to browser-enabled desktops. Most companies cannot afford to handle all facets of this endeavor, as the pressure to provide these capabilities is intense; however, a shared and organized data collection can give an organization a significant competitive advantage. By using outsourced resources, companies can become more efficient with their internal business processes, and

that can make the difference between success and failure in this intensely competitive market. Here are some of the technologies that enabled the ASP concept:

- **The widespread usage of the Internet.** The migration from internal application management to a hosted application solution has become feasible due to the availability of the Internet and the constant development of Web-enabled solutions.
- **The declining cost of bandwidth, along with the increase of available capacity and the ability to easily access the Internet.** The combination of the declining cost of bandwidth and the increase in the accessibility of the Internet at higher speeds has enabled a hosted solution distributed over the Internet or through thin-client computing to become viable.
- **The increased use of shared applications in a client/server environment.** Users are now more accustomed to the usage of remote access for client/server technologies. This assists in making the shift to an ASP model fairly painless.
- **Browsers are now considered acceptable as a graphical user interface (GUI) for applications.** The acceptability of browsers as a GUI has increased with the popularity of thin-client and Web-enabled computing.
- **The added ability to include security and reliability to e-commerce and e-business solutions with management and maintenance of third-party companies or software.** E-commerce and e-business solutions share the same business and technical concerns for security and reliability. The ability to resolve these issues will help influence the acceptance of hosted applications.

## Technical Factors

Many technical factors can add to the practicality of the ASP concept. With the cost of equipment dropping and the amount of qualified technical personnel at an all-time low, there are serious possibilities for the ASP to become a viable alternative to inhouse hiring of personnel.

The following are technical reasons that make the ASP model work:

- **A shortage of skilled IT labor** Most organizations, particularly smaller companies, cannot afford the time and expense associated with recruiting, training, and retaining highly skilled IT employees.
- **The acceptance of emerging technologies and “best of breed” applications** The ASP model, with its favorable economics, allows smaller organizations to employ more complex applications such as Supply Chain Management (SCM) and Customer Relationship Management (CRM), which up until recently had only been affordable and manageable by larger enterprises.
- **The ability to accelerate application deployment** The average duration for an ERP deployment is over 12 months. The implementation periods can be measured in weeks and days in the ASP model compared to years and months that are associated with traditional channels.
- **The ability to adjust to the rapidly changing and more complex technologies** Internal IT departments struggle with the rapid IT development, as it is hard to maintain the infrastructure and also be on top of the latest and greatest available technologies. The ASP model helps to resolve these issues by assuming the application responsibilities of keeping up with the new technologies.
- **Ease of obtaining technical expertise** Many ASPs focus on particular markets, business functions, or application types. This approach is very valuable to an organization that is searching to solve particular needs.
- **The ability to transfer risk** The IT sector has traditionally been concerned with the acceptance of an application among its clients. Due to these concerns, organizations are wary of their ability to deploy the next “big thing.”

## Barriers to the ASP Business Model

Some concerns and issues are impeding the growth of the ASP model. Some of these challenges confronting the budding ASP market are due in part to its relative youth. Here are some issues that will need to be resolved or addressed before market acceptance can be realized:

- **The ability to secure information** One of the largest challenges to ASP acceptance is the uncertainty of the security of proprietary information. Companies are very apprehensive about endangering sensitive



information. Companies are demanding very stringent security standards from ASPs than would normally be required internally. The integrity of mission-critical information is an important benchmark for the success of the ASP model.

- **Quality of service and support** Performance concerns generally include availability, scalability, bandwidth capacity, and network redundancy. Service level agreements (SLAs) are contracted agreements that bind an ASP to a predetermined level of service and performance. These agreements obligate that performance standards and measurements will be maintained. A typical arrangement would require an ASP to provide 99.999-percent (also known as “five nines”) total service availability and uptime, which guarantees all but constant uptime. The ASP’s quality of service is evaluated by the ability to ensure that there is no single point of failure; they can accommodate increasing network traffic spikes, and the perception that the system is based locally.
- **Breadth and depth of services** There is a tradeoff between breadth and depth for ASPs. These demands require expertise from an ASP on the front end of the application, as well as the back end where the ASP has control of the implementation and infrastructure requirements. This is further convoluted by requirements that a company will ask of an ASP to meet their unique needs. The point of contention will become whether the ASP or the application has the ability and flexibility to accommodate these constantly evolving demands. This leads into the adaptability of software.
- **The adaptability of software** Most software is not Web enabled. To be more efficient, existing ERP software applications are evolving toward a net-centric model that is capable of using the Internet; therefore, greatly increasing accessibility, gathering information from multiple destinations, and reducing maintenance needs. Future applications will be developed with modular components so that they can be upgraded more efficiently for improved functionality.

## ASP Business Model Strategies

ASP participants have used multiple strategies with regard to the type of applications they will host. The configuration of ASP channel components helps target their core market. There is no dominant business model that has proven to be the

standard by which others are measured. This book outlines some ASP relationships and strategies that are developing within these industry sectors—“pure” ASPs, traditional systems implementers and integrators, telecom companies and ISPs, and the ISVs.

### *The Pure ASP*

USinternetworking (NASDAQ: USIX) was an early adopter of the ASP model, and became publicly traded in late May 1999. The investment community allotted a billion-dollar market valuation to the company with a 12-month trailing revenue base of \$15 million. The company offered enterprise-hosted solutions. USinternetworking stressed shortened implementation intervals and attempted to narrow the gap between standard design configurations and clients' customized needs. The company financed, developed, and managed four data centers located in Annapolis, Milpitas, Amsterdam, and Tokyo, and continues to expand its network facilities. Their management believed that the ownership of data centers was a fundamental component of being an ASP.

This philosophy underwent a change when, in June 1999, they announced that a strategic business alliance had formed with The Hunter Group (a subsidiary of Renaissance Worldwide). The Hunter Group had managed the planning and integration of PeopleSoft Financial Management and HR systems for new client engagements.

USinternetworking had originally amassed the internal resources to manage these responsibilities, but the late September acquisition of Conklin and Conklin, a company with systems integration and Lawson Software expertise, followed the August announcement that introduced Lawson Software to its own enterprise-hosted solutions. These developments are important, as they shows the challenges that are involved with management of multiple channel responsibilities.

Corio, a company that was founded in 1998, is a privately held organization that also received notoriety as a pioneering ASP. Their management was based on the establishment of third-party partnerships and focusing on select applications to provide a competitive advantage. Corio focused exclusively on PeopleSoft and Siebel enterprise applications, and then partnered with Concentric Networks and Exodus Communications to manage its data centers. Corio deliberately built strategic relationships with third-party ISVs by partnering with preeminent enterprise application solution providers. The company also developed strategic relationships with other channel partners, but is now a single-source solution provider for its clients.

FutureLink Distribution (NASDAQ: FLNK) is an ASP that coined itself as “the world’s first computer utility company.” The company offered four services: application service provider, IT outsourcing, business practices consulting, and facility management. FutureLink positioned itself as an end-to-end ASP solution by internally managing its ASP channel. The company tailors its hosted applications to the needs of the customer, rather than offer a limited application-hosting portfolio. FutureLink has established software vendor relationships with Great Plains Software, Applix, Galleon Distributed Technologies, Microsoft, and Onyx.

Telecomputing ASA is an ASP based in Norway. It is quite possibly the longest operating ASP in the world. Its humble beginnings can be traced back to 1995. The company can provide a hosted application solution and has several applications specific to the client market. The company also boasts one of the most complex yet proven client networks with its customers using more than 70 different applications throughout Europe. The company had been focused on the European market, but moved its worldwide headquarters to Fort Lauderdale, Florida in 2000.

AristaSoft has a noticeably different strategic objective compared to other ASPs. AristaSoft claims that it is the first industry application service provider (IASP) focusing on using J.D. Edwards software. Their management believes that its business and technical knowledge gained on a single product focus will simplify implementation by 80–85 percent. The company uses subcontractors for integration and implementation to service its clients, and has contracted with Exodus Communications for external infrastructure services.

ServiceNet is a joint venture between Accenture (formerly Anderson Consulting) and GTE Internetworking. This company is based on bringing together software vendors and partners with experience in desktop support and maintenance services to supplement its Lotus Notes expertise.

Global Recruiting Solutions is a human resource applicant tracking and hiring process management vendor that is entirely Web enabled. Headquartered in Raleigh, North Carolina, this company is based on a completely outsourced solution, as it is designed to replace internal client/server products.

## System Integrators and Implementers

CIBER Enterprise Outsourcing created an ERP outsourcing division so that they could become an application solution provider. In February of 1999, they purchased Paradyme HR Holdings and assisted in the development of its ERP outsourcing and hosted solutions. The company built a hosted application model that is based on multiple back-office ERP solutions.

CIBER hosts applications from PeopleSoft, Lawson Software, SAP, Baan, J.D. Edwards, Lotus Notes, and Microsoft Exchange. The division acts as a single-source solution provider, and internally manages their systems development, maintenance, and data center infrastructure.

Metamor Worldwide started in June 1999 when it formed its Enterprise Operations (E-Ops) business unit. Metamor built a single-source ASP solution, by striving for a integrated ASP model that was capable of delivering complete services that are specific to each client. The company uses Baan, PeopleSoft, and SAP enterprise solutions. E-Ops also developed relationships that focus on ERP, SCM, and CRM solutions.

Breakaway Solutions is a strategic consulting and systems integration company that was founded in 1992. The company positioned itself as a single-source solution for application hosting, consulting, and systems integration. The company internally develops hosting relationships that span Web sites, e-commerce, CRM, and database applications.

Breakaway focuses on front-end enterprise applications, and disregards back-end enterprise applications such as ERP and SCM. The company has established over 20 client-hosting relationships and 75 CRM strategic-consulting relationships that can be leveraged into hosting relationships. The company remotely manages its application-hosting development through its “solution centers.” The company went public in October of 1999.

## Internet Service Providers and Telecommunication Companies

So, here is the business section on which this book focuses. Several telecommunication companies and ISPs feel that the ASP market is a logical progression of their Web-hosting knowledge. These telecoms and ISPs are expanding their business offering into other hosted business solutions and the ability to offer infrastructure management to other ASPs.

Qwest Communications International is a large Internet-based communications company. They entered the ASP market by forming Qwest Cyber.Solutions, a joint venture with KPMG. The company licensed agreements with SAP, Oracle Business OnLine, and Siebel. The company owns a fiber network that spans over 20,000 miles. Their management fully controls the data center infrastructure, and partnered with several companies to facilitate the other aspects associated with ASPs. The company acquired Icon CMT in January 1999 to add over 400 IT consultants and Web application-hosting experts. The company then partnered

with KPMG to provide systems implementation and integration services for its ASP clients.

Interliant became public in July 1999, when it transitioned to expand its Web hosting to include a variety of hosted applications such as e-mail, messaging, and CRM. The company cross-sells its Web hosting services with application hosting. Interliant also manages an extensive Web hosting customer base. Interliant broadened its CRM-hosted applications by acquiring Sales Technology; a UK-based IT firm that specializes in implementing CRM solutions. The company manages its own data centers, but increases its capacity with colocation support by UUNET, a subsidiary of MCI WorldCom.

Exodus Communications and UUNET aggressively market their colocation infrastructure hosting capabilities to ASPs as well as standard bandwidth provisioning. They facilitate data center management relationships for ASPs and avoid developing their own hosting solutions. Providers such as NaviSite use a direct approach with the development of their own ASP services.

## Independent Software Vendors

Some of the most influential backers of the ASP model are the enterprise software vendors. The ASP concept is a departure from the standard business models that are familiar to software vendors. ISVs give up large, upfront license payments to receive smaller annuity payments that accrue over long periods of time.

Because of this, it is astounding that the largest enterprise software vendors have accepted this business model so readily. There are several benefits for the software vendors make this route viable.

- **New market opportunities** ASPs generally target small to medium-sized enterprises (SMEs). This market was typically ignored by larger enterprise applications because of the complexities and high costs associated with the customization that is needed. The growth of ERP implementations among Fortune 1000 companies stagnated as the Tier 1 market became. The ASP model created new software channels and assisted in the distribution to lesser-exposed markets.
- **The ability to be first, and all of the advantages that come with it** Software vendors are motivated by being the first mover to delve into this market. By establishing an early presence in the ASP market, a company can establish itself as the dominant provider and potentially create barriers to the entry of other companies. The earliest software vendors, including PeopleSoft, Oracle, SAP, J.D. Edwards, Great Plains Software,

Broadvision, and Siebel as a collective group, controlled three quarters of the ERP market in 1998.


- **Some predictability within the ASP business model** There is a level of revenue predictability within the ASP model. Three-quarters of quarterly revenue for software vendors traditionally occurs in the last weeks of each quarter. The ASP model helps with the creation of a revenue backlog.
- **The learning curve involved in the new economy** ISVs are learning from their ASP partners, which helps in internalizing their own ASP infrastructures. With this turn in the market, it is conceivable that the software vendors will start bypassing the external ASPs and develop their own ASP model. This is a motivating factor for aggressive ISV participation.

## Independent Software Vendor Companies

So, what services were outsourced and enabled the ASP model to be viable? Primarily, there were four major companies whose applications emerged as the largest segment of hosted application services—Baan, Oracle, PeopleSoft, and SAP (this group can also be called the BOPS). The complexity of installing, configuring, and maintaining these complex ERP and Supply Chain Management (SCM) applications made them very attractive ASP offerings. The hosting of ERP, SCM, and other specialized applications (such as Great Plains) made it possible to connect groups everywhere within the supply chain to provide goods and services much more effectively.

To give an overall view of the types of value-added services that were involved in the transition from ISP to ASP, Figure 1.2 is a chart that was developed based on projections in 1999 by several research companies, which projected that Web site hosting, bandwidth guarantees, and security management would top the list of value-added services that were to be offered by ISPs as they grew into ASPs.

Oracle is an ISV that has embraced the ASP concept. The company has taken a different approach from other software vendors with the introduction of its own ASP services through its Business OnLine service. Management acknowledged that its application-hosting division would provide a source of revenues for its own company. Oracle also started building other ASP relationships by launching a \$100-million venture capital fund that focused on companies with ASP relevant technologies. This capital is being allocated to companies that base


**Figure 1.2** ISP Value-Added Services in 1999

their ASP services on the Oracle 8i database. In addition, Oracle launched iHost in July of 1999 to encourage partnership relationships between software vendors and service providers. They created the “e-certified” program, which is granted to e-business applications built upon Oracle platforms. Oracle estimated that its Business OnLine service would represent half of its total revenues by 2001. Figure 1.3 shows projected ISP Value-Added Services in 2001.

PeopleSoft created its own certified outsourcing partner program in October 1998. Since then, PeopleSoft has been licensing its technology to ASPs as well as enterprise clients.

SAP America joined the American ASP market in February 1999 by entering the hosting market with its R/3 financial suite. They also announced a nonexclusive alliance between SAP and EDS. This agreement divided the software and services between SAP and EDS. EDS handles systems implementation, integration, and infrastructure, and can contract with third-party providers. This German software giant was forceful in licensing its ERP systems to hosting providers in United States and Europe in order to increase its market share in the SME market.

Great Plains Software entered into the ASP model in October 1998 when the company created the Alliance Hosting Partner (AHP) program for its ASP partners. The company estimates that its hosted application revenues will account for approximately one-fifth of its total license revenues by 2001 and could reach half of their revenues by 2004.

**Figure 1.3** Projected ISP Value-Added Services in 2001

## Why All the Mergers?

Why are these companies merging and acquiring other ASPs? In reviewing some of these mergers and acquisitions, these partnerships reveal a strong tendency to enhance their domain and industry expertise while strengthening their ASP competencies.

Many of the companies that demonstrated domain expertise and/or industry expertise have been targeted for purchase or partnership. Several ASP companies are trying to arm themselves with systems and application integration consultants and experts in PeopleSoft, Oracle, Sybase, and Lawson Software application (et al), as well as systems integration experience. Additionally, these resources should provide industry exposure. Remember that there are enormous challenges associated with managing multiple aspects of the ASP channel; consequently, many ASPs are forced to seek external skills expertise through mergers and partnerships.

## Performance Issues

With the adoption of Internet technologies and widespread connectivity that is now available, corporate networks are increasing the demand for universal, transparently accessed, Internet-based applications. As ASPs become a more viable alternative for corporate IT, the demands that are placed on service providers to



deliver high-level 24x7 service will continue to escalate. This places the burden of reliability and scalability on the ASPs' underlying system platforms.

There must be a reliable infrastructure to support your network and its future growth (expected and unexpected). Whenever an application is to be used on the Internet, it must be accessible all day, every day. This is crucial when your Web site's downtime is something that everyone in the world can see. Any transactions that cannot be completed in a timely, efficient manner will likely have financial repercussions for you and your clients. Due to this, outsourcing customers have zero tolerance for downtime (again, expected or unexpected).

## Amount of System Uptime (Five Nines)

Some of the points you need to consider when evaluating the high availability of your systems include: "Is your system robust enough to handle your application if it is based on a single system?" This is when your company will be measured by its availability, and will be charted by something they call the "five nines," as in 99.999-percent uptime. *Five nines* means that in a year's time, a system will be "down" or offline for no longer than five minutes. Since there is no system that is capable, no matter how well built, of providing the *five nines*, it is considered the Holy Grail of high-availability systems makers.

Hewlett Packard was the first company to write a *five nines* guarantee into its hardware contract, and Sun followed with a similar guarantee on its top-end servers. Sun offers a 99.95-percent availability guarantee on a single node.

In 1998, IBM created a metric that estimated what downtime of key business applications can cost on a per-minute basis. E-commerce applications were calculated at over \$10,000 a minute, and ERP applications were valued at over \$15,000 a minute.

Hardware vendors approach the elusive *five nines* by building their systems to accommodate for changes and allow them the ability to respond to problems with monitoring and management solutions. HP, Sun, and Compaq all support high-availability features like hot swap and hot plug disk drives, peripheral cards, replaceable power supplies, fans, and so on.

One of the main separators in the approach to hardware versatility is the ability to start with a single-processor system and allow for the addition of processors as performance needs increase. Sun's SMP systems and those of several other OEMs handle this growth with the additional processors installed on an open I/O slot.

HP supports scalability with its central processing units (CPUs), input/output (I/O), and random access memory (RAM) without requiring that the ASP sacrifice

one for the other. This is important for ASPs who have higher than normal I/O requirements. ASPs who face these demands should take note of a system's maximum bus bandwidth in their evaluation process. Generally, hosted applications are very I/O intensive, and the system's bus can be a potential bottleneck. Other things to consider are the type of error protection and correction.

## Failover

High availability is much more attainable from a multisystem perspective. With the ability to cluster these multiple servers, you gain inherent failover capabilities.

*Failover* allows servers to be mirrored and then brought online in the event of a catastrophic system failure. These clusters can also be used to deliver workload balancing; if a particular system in the cluster (node) fails, the applications that are running on the failed unit are redistributed throughout the remaining nodes on the cluster automatically. There is a range of cluster implementations, from dynamic load balancing to those that require manual intervention, based on the vendor.

This should be a major consideration for ASPs and ISPs that host potentially mission-critical applications for their customers, as SMP-based systems that can be clustered for reliability and scalability will be necessary for future growth.

## Clustering

*Clustering* is the combination of multiple servers that will allow for failover and data reclamation from storage in case of a catastrophic occurrence. This is necessary for the high availability of your network, as it will transparently access the redundant serve if there is a failure on the primary server.

## Sun Microsystems

Sun uses its Sun Cluster software, a major part of its Full Moon initiative, to handle failover and provide for parallel database functionality. Implementing the Sun Cluster package allows the clustering of up to four nodes that can be located up to 10 kilometers apart. However, in order to receive the system uptime guarantees for the clustering solution, you will have to purchase the top-of-the-line servers to support inter- and intra-domain failover.

## Hewlett Packard

HP's cluster solution, Multi-Computer/ServiceGuard (MC/ServiceGuard), will automatically detect and react to failures in system processors, memory, LAN media, network adapters, system processes, and application processes.

## Compaq

The Compaq ProLiant uses Microsoft Windows NT nodes and Compaq's StorageWorks fiber channel storage system to cluster their high-end servers. When set up in a dual-loop configuration, the system removes the single point of failure that is associated when running under Microsoft Cluster Server.

There are definitely downsides in choosing a system that lacks sufficient power and growth to meet evolving application requirements. The other side of the coin, though, is that by overinvesting in a system platform, which is generally not as large an issue for ASPs as it might be for corporations, an ASP can dynamically reapportion its under-utilized systems and resources to other tasks.

Applications hosted by ASPs will develop an escalating need for systems resources and higher performance. This will also increase the pressure that is placed on an ASP by its customers through service level agreements (SLAs) that will spell out what are acceptable application performance and latencies while creating a binding contract.

Hosting applications for delivery through the Internet will tax the resources of application service infrastructures of both ISPs and ASPs. There are greater expectations imposed on Internet-based services that are pushing the level of dependability and scalability equal to that of what is expected from the telecommunication industries.

## Problems That Could Arise from Conversion

ISPs that are converting to ASPs face an assortment of hurdles in trying to break into their chosen markets. Perhaps the greatest obstacle is the acquisition, training, and retention of intellectual property, all of which will allow an ASP to offer stellar implementation, service, and support. There is also the significant capital investment needed to purchase new equipment and facilities and maintain them.

One of the hardest things to do is to be able to scale network operations and support to meet customers' requirements and unanticipated demands, while maintaining and hiring qualified personnel. An ASP that is spreading its investment across many clients can more easily justify the hiring and retention of specialized personnel who focus on the maintenance of a particular service.

# Major Issues in the Implementation of an ASP Model

IBM did a study in the 1960s to determine how long a user would wait for an application screen to refresh before becoming impatient, and found that users became irritated after approximately two seconds. Now in the era of the Internet, where applications are running remotely, and network delays and propagation are added to application-processing delays, there are concerns about the latency involved with outsourced applications.

The contractual assurances that an ASP must make to its clients is usually some form of negotiated contract that specifies acceptable levels of service, availability, security, and performance collectively called a service level agreement (SLA). SLAs are used to spell out what an ASP agrees to deliver to clients, but also set penalties (typically financial) for failure to deliver the contracted for QoS.

QoS doesn't just refer to a controlled, guaranteed level of bandwidth; it is also a software mechanism that assists service providers in proactively managing and controlling connections to ensure that their commitments are being met at the application level, not just the bandwidth allocation.

The software applications must conform to a company's business guidelines by being able to discriminate between customers, partners, and suppliers and provide the best business value, and return on a company's investments (ROI) in time and resources.

## What Is Needed to Sell Your Services?

An ASP must draw together resources that traditionally have operated independently of one another. Since the ASP is a hybrid of technologies and functions, it must integrate content from several functional organizations, including:

- **Sales and marketing** This organization has traditionally been charged with the creation and management of the outbound corporate communications and identity. This group will be driving the majority of the Web-based application materials to your clients.
- **Information technology** This organization holds the expertise in computers and networking. Their involvement helps to ensure that applications remain operational and available, that network resources are sufficient to keep up with demand, and that access to resources is limited to authorized users.

One of the challenges involved with the creation of an ASP is the integration of these two units based on the fact that these groups are usually not experts in each other's function and may not understand the rationale as to why tasks are performed in certain ways. Each organization's role in developing and marketing applications must leverage its existing functional expertise and not significantly alter its workflow process.

## Necessary Components

To successfully deploy a dynamic and interactive application, you will need to integrate several components, while providing access to other network resources. Although the number of potential application types is nearly limitless, the following content is the most likely what needs to be implemented to create a successful ASP.

- **Marketing copy** This is necessary because without name recognition and a story, there will be no sales for your applications. Most marketing material is likely to be mass-produced with graphical content, and will be used in the development of sales.
- **The application itself** This is the product itself, which will be shown to, and hopefully purchased by, the customer.
- **The infrastructure** This is where the platforms and applications will run. This will entail the network and all of the necessary application storage.
- **Method of access** What is necessary for clients to use your application? Do they need to have a high-speed link, and can you support it if they do?

## Summary

As with any business undertaking, you should first scope out the needs of the market you wish to move into. A strong business plan should keep you from running off into the weeds, and it should allow you minimal “red” time while you ramp up to the next generation of service providers.

There will be a need for ASPs to assist companies in leveraging their best skills, while outsourcing those items in which they have little competency. This will create organizations that are more efficient, yet still allow for individuality with modifications for applications on a per-company basis.

There are unknown factors that will make or break a company—Karma, if you will. Sometimes, it really does take being in the right space (or market) at the right time. You must find a need and fulfill it—therein lies the rub.

This chapter is an overview of the rest of the book. This chapter pointed out some of the common definitions and acronyms that you will run into when implementing and marketing an ISP and ASP business model. We discussed what elements allowed for the creation of a viable ASP model. The chapter discussed what are some of the possible offerings that you can make to your clients, as well as the types of firms that are associated with these offerings.

On a hardware level, the chapter discusses the OSI seven-layer model and how it will affect you, and the probable platforms that are available to your company’s infrastructure build-out. We discussed the probable performance issues and gains that are inherent within the ASP infrastructure, and the potential problems and issues that could occur with the conversion.

Finally, we talked about the business drivers that will assist you in the conversion to an ASP, and what services you will need to use your services.

## Solutions Fast Track

### Why This Book Is for You

- ☑ According to the International Data Corporation (IDC), worldwide spending for outsourcing services should reach approximately \$142 billion by the year 2002.
- ☑ The ASP market began capturing the interest and commitment from a large number of venture capitalists and the telecommunications industry in the late 1990s.

- ☑ The ASP concept is the advent of a new computing era, with small to medium-sized companies searching for IT alternatives, and a gradual acceptance among larger enterprises.
- ☑ The IT infrastructure has evolved from a self-contained environment to a distributed computing model and now toward a net-centric infrastructure that links multiple areas of operation.

## Definitions of Common ASP Terms

- ☑ An Internet service provider (ISP) is an organization that provides access to the Internet. ISPs can provide service via modem, or dedicated or on-demand access.
- ☑ The ASP Industry Consortium, an alliance of companies formed to promote and educate the IT industry, offers the following definition: “*An ASP manages and delivers application capabilities to multiple entities from a data center across a wide area network.*”
- ☑ The definition of a *pure* ASP is an ASP that joins with a particular ISV, and performs the initial application implementation and integration.
- ☑ Information technology (IT) outsourcing is the transfer of an organization’s internal IT infrastructure, staff, processes, or applications to an external resource provider.
- ☑ Business process outsourcing (BPO) and information utilities providers are primarily concerned with economic and efficient outsourcing for the highly sophisticated but repetitive business processes.
- ☑ Platform IT outsourcing offers an array of data center services, such as facilities management, onsite and offsite support services, data storage and security, and disaster recovery.

## The Elements That Make an ASP Viable

- ☑ The initial purchase price of system software such as a Unix platform or a Microsoft Windows platform and their licensing are considered part of the initial system software purchase.
- ☑ Initial application software acquisition is any application that assists in the productivity of the organization.

- ☑ Hardware upgrade costs are associated with obligatory improvements to hardware when your company will need to support expanded applications databases, and a more robust operating system.
- ☑ Operating system software may need to be upgraded to support newer, vigorous applications.
- ☑ Applications are constantly being improved due to customer and client demands.

## Possible Business Models and Offerings

- ☑ ASPs host services work on an extensive array of hardware, so at any given time that hardware will have a substantial amount of its processing power idle. The ASP will find that this ability to provision and partition that extra horsepower can be the basis for a very valuable and profitable differentiation service offering.
- ☑ The ability to offer different types of service to different types of clients is an incredibly valuable way for ASPs and ISPs to provide granular and real-world service degrees of difference.

## Types of ASP Firms

- ☑ There are several types of ASP-enabled firms. These organizations can be separated into professional consulting, project-based service providers, outsourcing providers, staff augmentation providers, education and training providers, and value-added resellers.
- ☑ Professional consulting firms focus on corporate-level business and strategic engagements. This can be broken down into three subcategories: IT consulting, Strategic management consulting, and Business process consulting.
- ☑ Clients that select project-based service providers for projects are opting for well-defined tangible deliverables and scopes. Contract designs range from a billable-hours approach to fixed-price engagements for components and entire projects. These companies focus on industry expertise, either in specific technologies or industry applications.



- ☑ Outsourcing providers are organizations that provide process automation services, facilities management, and operations for clients who require an assortment of technical answers.
- ☑ Staff augmentation organizations specialize in providing IT professionals, on a temporary or long-term contract basis, to clients who need specific skill sets and support for internal systems and development projects.
- ☑ Education and training companies provide training and help desk consulting to firms that have implemented custom-designed or packaged software products.
- ☑ Value-added reseller (VAR) organizations are solution-oriented vendors who can provide integration for hardware and software systems.

## ISO-OSI Seven Layer Model

- ☑ The OSI reference model is a conceptual model composed of seven layers, each specifying particular network functions. The OSI model divides these communications involved with the moving of information between networked computers into seven smaller, more manageable layers.
- ☑ The *upper layers* of the OSI model handle application issues and are generally implemented in software.
- ☑ The *lower layers* of the OSI model are also known as the Data Transport layer.
- ☑ These *pseudo layers* are not actual OSI model layers, but they will directly influence the way in which you will implement your equipment and policies.

## Choosing the Best Platform for Your ASP

- ☑ ASPs take advantage of existing Internet connectivity to offer corporations the opportunity to outsource not only peripheral applications but also mission-critical applications.
- ☑ Traditional ISPs are experiencing an explosion of data traffic on their networks. The Internet and its dramatic growth have fueled the need to satisfy this increasing data demand by forcing the migration towards multiservice network platforms.

- ☑ Reliability is one of the most important considerations to make when choosing a server platform, but it is likely that ASPs should be most concerned about the operating system. This often boils down to a choice between some form of Unix, Linux, or Microsoft Windows platform.

## Business Drivers for the Conversion to ASP

- ☑ The new value proposition that is offered by various services using Internet and intranet technologies is the ability of the ASP to free its customers from having to develop, maintain, and provide services for themselves.
- ☑ There is also the added benefit of using an ASP's application management expertise. Over the lifetime of an application, such as Enterprise Resource Planning (ERP), an ASP can estimate that software licensing, hardware and basic infrastructure costs will account for less than one-fifth of the total cost of ownership (TCO) over a five-year period.
- ☑ By using outsourced resources, companies can become more efficient with their internal business processes, and that can make the difference between success and failure in this intensely competitive market.

## Performance Issues

- ☑ As ASPs become a more viable alternative for corporate IT, the demands that are placed on service providers to deliver high-level 24x7 service will continue to escalate. This places the burden of reliability and scalability on the ASPs' underlying system platforms.
- ☑ *Five nines* means that in a year's time, a system will be "down" or offline for no longer than five minutes.
- ☑ Clustering is the combination of multiple servers that will allow for failover and data reclamation from storage in case of a catastrophic occurrence.

## Problems That Could Arise from a Conversion

- ☑ ISPs that are converting to ASPs face an assortment of hurdles in trying to break into their chosen markets. Perhaps the greatest obstacle is the acquisition, training, and retention of intellectual property, all of which will allow an ASP to offer stellar implementation, service, and support.

## Major Issues in the Implementation of an ASP Model

- ☑ The contractual assurances that an ASP must make to its clients is usually some form of negotiated contract that specifies acceptable levels of service, availability, security, and performance collectively called a service level agreement (SLA).
- ☑ The software applications must conform to a company's business guidelines by being able to discriminate between customers, partners, and suppliers and provide the best business value, and return on a company's investments (ROI) in time and resources.

## What Is Needed to Sell Your Services

- ☑ An ASP must draw together resources that traditionally have operated independently of one another.
- ☑ To successfully deploy a dynamic and interactive application, you will need to integrate several components, while providing access to other network resources.

# Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the “Ask the Author” form.

**Q:** Now that I am currently running an ASP, where do my revenue projections come from?

**A:** This depends on the market you are trying to corner. What you want to create is name recognition and product satisfaction. These will lead to the revenue streams that will allow you to advance your marketing and create a company story.

**Q:** How can I make this ASP a necessary component for businesses?

**A:** Oracle saved over 1-billion dollars by internalizing their information technology needs in 1999 and 2000. What you will be able to offer is a higher competence, and therefore more efficient applications that will assist your clients in saving time and resources.

**Q:** What are some of the pitfalls to the ASP model?

**A:** Starting in April 2000 and continuing until today, there has been a massive downturn in the market. This has had a ripple effect on technologies-based companies. At the time of the writing of this book, companies mentioned in this chapter were being purchased, retooled, or even closed. This should not be a deterrent to your ASP growth. Part of the major issue that has faced many of these companies is that there was no strong, repeatable business model that earned money. This caused investors to shun even those companies that did have strong financials and good business practices.

**Q:** Where can I find information about ASPs?

**A:** The ASP Industry Consortium, located at [www.aspindustry.org](http://www.aspindustry.org). This should assist you in finding information and trends that may be more current than those at the time of the publishing of this book.



## The Business Case

### Solutions in this chapter:

- **ISP Market Conditions**
- **Service Provider Business Requirements**
- **The Evolving ISP**
- **The Service Provider of the Future**
- **The Case for Application Service Provider Conversion**
- **Critical Success Factors**
  
- ☑ **Summary**
- ☑ **Solutions Fast Track**
- ☑ **Frequently Asked Questions**

# Introduction

The proliferation of widespread and affordable Internet connectivity has revolutionized communications and the transmission of information as we know it. In a matter of years, we have changed the way we communicate with each other in both business and personal situations. Ten short years ago, our primary methods of communication generally centered on telephone calls and postal delivery services. Today, you are more likely to receive e-mails than all other methods of remote communication combined. With technologies like Voice-over IP (VoIP), Web Collaboration, and IP-Video Conferencing, many companies are using the Internet to replace previous methods of communication at a far lower cost.

The Internet also serves as a primary source of information in both business and our private lives. There is virtually limitless information available at our fingertips, and it can all be accessed with a simple Web browser (common with everything from personal computers to cellular phones). Intuitive search engines are able to provide us with instant access to a variety of sources of information. These capabilities have enabled the Internet to drastically change the pricing model for information. It has become far more difficult to charge for information in any form. Many traditional sources of information are now migrating to the Web. It is now possible to listen to the radio, search your favorite publication for a specific article, or get information on an upcoming television show, all for low or no cost on the Internet.

The Internet service provider (ISP) industry has been enabled and promoted by the ubiquitous adoption of a disruptive technology, the Internet. Many businesses have been forced to face these changes in one way or another. Industries never before conceived of have become components of everyday life, while long-standing business models have been torn apart in a matter of years.

The vast majority of these changes have empowered the user with access to information and new ways to communicate with associates and friends. These changes have also driven some of the more common inefficiencies out of supply chains everywhere, which helps in providing better products at reduced costs.

It is now time for ISPs to address the magnitude of the changes they have induced. Their customers have been empowered by the revolution the Internet has created. Users of Internet connectivity now have a vast number of alternatives sometimes at minimal or no cost. They also have information on these many options at their fingertips, thanks to the Internet connectivity provided by their ISP. To make the situation more precarious, the capital markets are quick to punish laggards and often times an entire market segment. To escape the economics

of a commodity reseller, ISPs must disintermediate the current providers of value-added services and continuously develop the applications that will attract and retain customers.

## ISP Market Conditions

The attraction to the ISP market is obvious. Users are adopting the technology faster than virtually any other advancement in history. Internet access reached 50-percent market penetration in less than eight years of existence. The growth rate in the United States is projected to be anywhere from 40 to 110 percent for at least the next few years. The growth rate is even more impressive when you measure bandwidth growth.

More importantly, much of the world has yet to be provided with Internet access, particularly some of the world's most populous nations such as India and China. These nations, which count their populations in hundreds of millions if not billions, have pent-up demand that is only increasing with the passage of time.

Even more attractive is the ever-increasing need for bandwidth. It has been demonstrated that the dial-up connection is only an introduction connection to the Internet. Users quickly lose patience with the slow speed of dial-up connections and long for broadband access. Applications such as digital photography, interactive content, and downloadable music only reduce the cycle-time for the inevitable upgrade. Demand for Digital Versatile Disk (DVD) quality video and other high-throughput applications has not even started its ascension, and this drives the demand for connection speeds far higher than the 1.544 Mbps that is now considered acceptable for a to medium-sized businesses.

Even residential users will require speeds exceeding those currently offered by Digital Subscriber Lines (DSL), cable modems, and the like as they begin to implement multiuser home networks, videoconference, and use collaborative applications for business and pleasure from the home. Recreational activities such as downloading feature films or efficiently trading entire albums will also drive the need for additional bandwidth. Consumers will not accept the trip to the movie store for that much longer, so the ability to access downloadable movies 100 times faster than anything that is currently available will be required to provide almost immediate access to the majority of existing films and shows.

Internet connectivity has become almost a requirement for any business and is quickly trending toward 90-percent penetration within the consumer market. As the power of convergence is fully implemented, Internet connectivity will become more of a necessity than connections to the Public Switched Telephone



Network (PSTN) are today. Access to telephone calls, high-quality television and radio, as well as a multitude of other services will all be provided by a single connection. The demand for value-added services is also increasing. Businesses and consumers are having their Web sites hosted, data stored, and applications provided across Internet connections.

If all of these reasons weren't enough, the tremendous pace of technological advance is providing faster and more reliable connections to meet the demands of the consumer. These advances are providing new offerings such as wireless broadband or private DSL-to-ATM (Asynchronous Transfer Mode) networks that solve a host of problems. Customers will want to upgrade to these new services, which will continuously push the revenue-per-user up for those service providers that are able to add these new technologies to their product offerings.

ISPs that thrive in this environment stand to profit enormously. Normally revenue-generating networks can become far more efficient at higher utilizations. Those players with the largest user base will likely be able to develop impressive economies of scale and develop barriers to entry that currently do not exist. Those dominant players should also enjoy the best margins in a commoditizing business. Figure 2.1 lists the services that are driving the demand for bandwidth.

**Figure 2.1** Services That Are Driving the Demand for Bandwidth

Consumer	Business
	Virtual PBX Service
	Video conferencing
	Voice over X
	ASP Services
	E-Commerce
	Managed VPN
	Distance Learning
	Managed Services
	Internet Access
Cable Quality Video	
Video on Demand	
Voice over X	
Digital Audio	
Interactive Gaming	
Internet Access	

## The Onset of Commoditization

The amazing potential of the ISP market did not go unnoticed. Competitors and investors alike flocked to the segment. According to *Boardwatch Magazine*, there are currently more than 7700 ISPs (early 2001) that are doing business in the United States alone. Many of these companies received access to significant amounts of capital both in the form of equity and debt. The business models that were developed with the low-cost capital tended to value market penetration over the profitability of the connections. The sheer number of competitors in the segment could have driven price competition, but the situation was exacerbated by the free access to capital in a highly speculative environment. The inevitable outcome was the commoditization of relatively undifferentiated products.

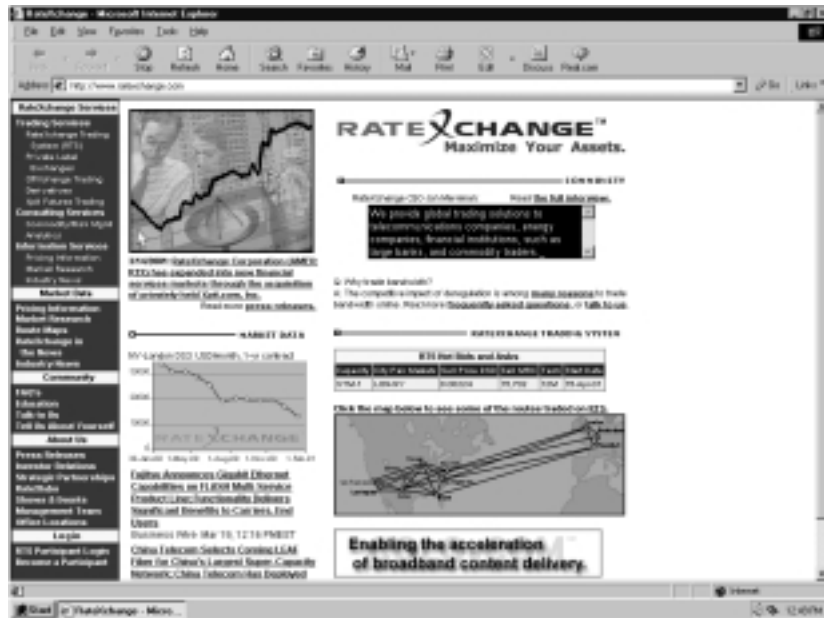
ISPs currently face stiff pricing pressure from competitors providing access ranging from inexpensive, approximately \$19.95 per month, to those that are free. While the free service model may be unsustainable, it has severely impacted the consumer's perceived value of Internet access, particularly dial-up services. The large number of competitors in the segment chasing a commoditizing business should keep average prices falling. The pricing pressure is being seen in virtually every segment from residential access to core transport.

The numerous competitors, continuous marketing efforts, and decreasing pricing have also impacted the industry's customer loyalty. Solomon-Wolff Associates suggests that the turnover rate of Internet users is approximately 25 percent per year. While turnover is most prevalent in the dial-up space, it will follow in broadband as well. As provisioning delays decrease and competitors deploy connections to customer premises by fiber, wireless, and traditional means, customer turnover, or churn within the broadband space, will increase. The business plans of Building Local Exchange Carriers (BLECs) and wireless providers center on supplanting the incumbent carriers in the most coveted multitenant locations.

Enormous increases in capacity coupled with largely undifferentiated service offerings has led to commoditization in all segments of the industry, including consumer, business, and transport. According to Sanford C. Bernstein & Co., these factors have led to an 80-percent decrease in the price of Internet connectivity during 2000. Definitive proof of the commoditization of bandwidth has also surfaced in the form of bandwidth exchanges and auctions.

Companies such as RateXchange ([www.ratexchange.com](http://www.ratexchange.com)) offer trading systems that allow telecommunications companies to buy, sell, and deliver bandwidth around the world. Their RateXchange Trading System allows participants to specify the amount of capacity, the route, and the duration of the contract (Figure 2.2).

Figure 2.2 RateXchange



A second offering, CustomAuctions, enables members to buy or sell a wide variety of telecommunication products through an online auction site. Items include bandwidth, dark fiber, and minutes of capacity. Options include English-style auctions (similar to eBay), reverse auctions (similar to Priceline.com), and sealed bid auctions. RateXchange even provides strategically located delivery hubs to facilitate participants' access to each other's networks. It is now possible to trade bandwidth and fiber with no more difficulty or differentiation than steel or chemicals.

The result of these factors is lackluster income statements and very difficult paths to profitability. The easily accessible capital that in many ways created the current situation has now flocked to safer havens. The capital markets, venture firms, and private investors that once courted the industry are now far more selective in both debt and equity investments. The valuations of both instruments have been severely impacted, virtually cutting off additional sources of capital for the service provider space as a whole. Existing shareholders now demand profitability in stark contrast to earlier requirements for market penetration.

## Broadband—The Enabling Technology

Initially, the growth of broadband seemed to be the way to escape the strong pricing pressures that dial-up providers faced. Significantly higher pricing was not

holding back explosive growth rates for broadband connections. Investors quickly took notice, and capital flowed into the broadband segment. For a while, companies such as Covad Communications, NorthPoint Communications, and Rhythms NetConnections seemed like the exciting evolution of the ISP. Unfortunately, as with dial-up access before, the realities of an undifferentiated product and strong competition drove pricing down and demonstrated the inefficiencies of their business model.

The reality of the DSL market is that providers must rely on the Incumbent Local Exchange Carrier (ILEC) for the all-important connection to the customer. This forces ISPs into the position of commodity resellers which puts them in direct competition with their suppliers. The extreme pricing pressure inherent in a commodity environment makes it difficult for new entrants into the DSL segment to provide the low prices required by the market while still retaining profitability. Additionally, many of the providers chose not to develop a sales force, but instead contract yet another layer of resellers to bring their product to market. The combination of these factors has already driven NorthPoint Communications into bankruptcy and has put many others in financial jeopardy.

To date, the cable industry has been able to keep virtually all competition out of their networks. It remains to be seen what the outcome will be, but all interested providers should carefully study the lessons of DSL providers. While the cable industry may not have to deal with direct competition on their infrastructure, they will not be immune to the competitive access costs of other mediums such as DSL, terrestrial wireless, or satellite. If they do not succumb to the pricing pressures of the industry as a whole, they will see massive turnover within their user base. They will face the same realities as all other providers and be required to add additional services to drive revenue.

While broadband connections seem to be following the same economic pattern as their slower counterparts, their significance should not be overlooked. Increasing broadband access speeds will be the foundation for the value-added services that will allow ISPs to differentiate their offerings. Bandwidth as a stand-alone technology will not provide profitability for service providers, but the capabilities of those connections and the advantages of packet-switched technologies will allow ISPs to add services that are highly profitable.

The inherent capabilities of high-speed packet-switched infrastructures will also perfectly position ISPs to capitalize on the shortcomings of legacy networks. In addition to offering traditional data services, ISPs will be capable of aggregating services that were previously provided by multiple disparate networks. Examples include local service, secure point-to-point circuits, long distance, and videoconferencing.

## Designing & Planning...

### Cogent Communications

It should not go unnoticed that even the very high end of the market is now coming under severe pricing pressure. Cogent Communications, an optical ISP located in Washington, D.C. ([www.cogentco.com](http://www.cogentco.com)) has offered 100 Mbps of “non-oversubscribed” Internet access for a flat rate of \$1,000 per month. The pricing does not require a long-term contract and covers unlimited usage. They have similar offerings for wide area network (WAN)/metropolitan area network (MAN) connectivity, and project that by 2002 they will be offering 1-Gigabit access for the same \$1,000 monthly price.

At \$1,000 per month, the connection costs roughly the same as a T-1 to the Internet, but the Cogent connections are 65 times faster than the 1.544-Mbps rate of a T-1. On a per-megabit basis, the connection is roughly 100 times more cost effective than traditional connections. Many consumers are gravitating to these offerings because they are handed off as a FastEthernet connection. By receiving WAN traffic in their native LAN protocol, users avoid conversion and the time delays that are associated with processor and memory requirements. Clients are usually able to significantly upgrade their connections without major hardware upgrades.

While the viability of the Cogent model must still be proven, other major players have also entered the market segment. SBC Communications threw its hat in the market with two new MAN solutions. Native LAN+ offers a 10/100 Ethernet connection for roughly \$5,000 per month, while their GigaMAN offering provides a gigabit connection between sites for roughly \$6,600 per month. These offerings are not currently available with integrated Internet connectivity, but as the market adopts high-speed MAN connectivity and becomes comfortable with the technology, we will see a strong demand for bundled or direct gigabit Internet connectivity.

The major players already have to address the number of requests for private OC-48 connections, since one OC-48 connection represents one-quarter of the capacity of an OC-192 transport link. The drawback is that it is extremely difficult for providers to efficiently add these connections to today’s infrastructure.

# Service Provider Business Requirements

While there is enormous potential in the Internet service segment, there is little opportunity in the business models as they have existed in the past. There is also very little room for tactical or technical error. ISPs must confront the realities of their current situation and develop strategies to overcome them. Virtually all of the 7700 Internet access providers in the United States are facing the same realities. Those providers that can successfully address these issues will be the market leaders of the future, and share the spoils of a far more lucrative converged communications segment.

In order to break out of the current cycle, many service providers and ISPs in particular will have to address these factors:

- **Commoditized offering** The current offerings of Internet connectivity and mail services are not enough to differentiate an ISP. Even service levels and service level agreements (SLAs) do not provide a distinct advantage with so many low- and no-cost providers. It has also been shown that increasing performance is not a sustainable advantage.
- **Significant pricing pressure** The high number of competitors in an undifferentiated segment can only lead to continued pricing pressure. Historically, average pricing has fallen as much as 80 percent annually. It also appears that advanced offerings such as DSL or 100/1000-megabit connections are facing significant pricing pressure at an accelerating rate.
- **High customer churn** Drastically falling prices, numerous connectivity options, and decreasing difficulty in switching providers have driven up churn. It is likely that broadband will now face increasing churn rates, similar to that which dial-up has experienced. Technologies such as terrestrial wireless, low-orbit satellite, and Ethernet WAN/MAN services are going to limit or remove the dependencies on the Regional Bell Operating Companies (RBOCs) for broadband connectivity, and make it faster and less painful to switch providers. New service provider models will also impact broadband churn. As BLECs bring up buildings, hotels, business parks, and residential developments, it will become both quick and easy for the most lucrative customers to switch carriers.
- **Drastically reduced valuations** Micro- and macro-economic factors have severely impacted the valuations of service providers as a whole. Additionally, the method for valuing ISPs has dramatically shifted.

Market penetration alone is not the method for determining value; the profitability of services and customers has returned as a guiding factor.

- **Restricted access to capital** Reduced valuations and a highly concerned investment community have severely impacted access to equity or debt instruments. However, significant amounts of capital are still looking for the right segment or model.

These factors are all highly interrelated. Valuations and access to capital are primarily dependent on churn, pricing pressure, and the other economic realities of a commoditized market. Churn and pricing pressure are partially the result of strong competition, but more importantly, are due to a lack of product differentiation. The benefit of this fact is that all of these problems can be addressed with a focused, coherent strategy. Do not overlook the reality that the interrelated nature of all critical success factors requires very good execution of your strategy going forward. There is very little room for error, and no way to buy time by addressing part of the list.

## The New Model

When developing a new business model, it is imperative that you thoroughly address the needs of your two most important external stakeholders: customers and investors. Look at these terms in the broadest sense. Customers include existing customers and potential customers of your current offering, as well as all customers you may be able to reach with expanded offerings. Don't limit your thinking. The greatest potential for the ISP segment is based on opportunities that are currently addressed by another industry. Examples include a variety of voice services, document management, data storage, and high-quality videoconferencing.

In the same way, you must address the needs of your current investors as well as all potential investors. Do not constrain your thinking here, either. Your traditional investors may not be the most important players in the evolving ISP space. The industry will undergo a pervasive transition from providing connectivity to providing services, in many cases adding professional services related to the design and deployment of these service offerings. Try to understand where the capital you require will need to come from, and who will understand the magnitude of the offering you will be going to market with. You are developing an infinitely scalable services model that can reach all corners of the globe. That is among the most attractive of all value propositions to an investor. Do not limit your thinking!

## Customers' Demands

Customers have become increasingly more discerning. They are currently focused on increasing performance and reliability at a better price. While it is likely that customer demands will always include those issues, it should be your goal to expand their view to include additional factors. If you can develop services on which they depend, functionality becomes more important than pricing.

It is imperative that services be developed that have a high perceived value and differentiation. America OnLine (AOL) has done this exceedingly well in the most commoditized of segments, dial-up access. It is very possible to provide a group of services that customers perceive as different from competitive offerings, whether factual or not. With this in mind, it is important to focus on the evolving demands of the customer:

- **Increasing performance** The majority of residential and business users are facing the need for increased data throughput rates. New technologies and services are driving this ever-increasing demand for bandwidth.
- **Increasing reliability** Businesses have come to a point where Internet connectivity is now a mission-critical application. E-mail and Web access have come to play such an important role in business that even short outages pose major problems. The mission-critical nature of an Internet connection only increases as companies begin to access additional services across that connection.

As services converge onto one connection, we will see the requirement for 99.999-percent (Five-Nines) reliability that is currently only required of voice networks. Residential connections are also demanding higher reliability percentages. Users are coming to rely on the convenience of the Internet, and people are routinely accessing business-related services from their home. In the near future, business-quality connections will be required in residential installations for home offices housing full-time remote users.

- **Security** Residential users will increasingly demand bundled security features from their provider. Those not complying will lose a significant piece of market share. Businesses will also look for monitored security offerings such as firewalls, intrusion detection, and corporate virtual private networks (VPNs). Security offerings are a major initial opportunity in the enterprise market. Act quickly to address these needs and continue



to develop your security offerings so that they do not become commoditized.

- **Improving pricing on commoditized offerings** Residential and enterprise users will continue to demand drastically improving pricing on all commoditized offerings. They will not have the same expectation of value-added services that are differentiated from competitive offerings. The key will be to bundle commoditized services with additional value-added services to develop higher perceived value and customer loyalty.
- **Additional services** Customers will continue to demand additional services. As new services are accepted in the marketplace, they will become required offerings to maintain customer happiness.
  - Services could include:
    - Application hosting
    - Inexpensive integrated long-distance services
    - Document management
    - Managed security
    - VPN services
    - Multicasting capabilities (for widespread transmission of voice and video streams)
- **Improved support** Customers are currently demanding better support for existing services, but be prepared for vast increases in the demand for high-quality support offerings as new services come to market. Troubleshooting of advanced mission-critical services will be far more urgent and complicated. Customers will look to their service provider for in-depth knowledge on applications and their implementation of it. The lack of qualified internal IT professionals is one of the main factors driving outsourcing to overtake the market. Do not overlook the increasing support requirements of the evolving ISP segment.
- **Fast installation and service upgrades** Customers will no longer accept four-to-eight-week installation delays. The dynamics of a monopoly model will no longer be accepted. Installation periods will need to be calculated in days or hours. Once connectivity is installed, the capability to deploy upgrades and additional services almost immediately is imperative. Some providers are currently offering immediate

bandwidth changes from Web-based interfaces. Immediate access to bandwidth and services will be a major differentiator in the future.

## Investor Demands

It was not long ago that the financial community was captivated by the seemingly insatiable demand for Internet access. Actually, there was an enormous demand for equity in all companies even remotely related to the Internet. Values in the ISP segment skyrocketed, and many financial partners appeared on the scene, happy to pay high multiples to join the game. It seemed that the focus had truly changed from sustainable business models and profitable offerings to market penetration. Some very smart people leading major service providers had gone so far as to ask analysts to judge them not on revenue or profitability, but instead on the number of lines they had installed. It is safe to say that those times are behind us.

The current demands of the financial community once again use traditional terms such as *differentiation*, *barriers to entry*, and *profitability*. The easy money is gone. Service providers of all varieties face difficulties selling equity or placing debt. Even vendor financing is becoming far more stringent. The segment is being punished for the excesses of the past. As difficult as today's realities are, the opportunities lurking within the turmoil are enormous. The ISPs that choose to embrace the evolution of the industry and update their current strategy will have access to capital. However, they must address the concerns of the financial community and clearly articulate why their model addresses these specific demands:

- **Profitable revenue growth** ISPs must be able to demonstrate that their strategy will allow them to increase revenues and penetrate markets, while providing profitable services. It is expected that new customers and new offerings will be quickly, and sustainably, profitable.
- **Increasing margin** It is also required that ISPs develop methods for improving current margins. In many instances, current providers have yet to prove that there is room for profitability in their services. Economic gambles like those of the past will not be looked upon favorably. Successful strategies will include methods for improving margins over time as ISPs differentiate their offerings and customize solutions for their users.
- **Differentiation** Capital will not be available to any service provider that has not thoroughly addressed a methodology to differentiate its offerings. In any market with strong competition surrounding an undifferentiated product, commoditization becomes the likely outcome.

- **Barriers to entry** The financial community is always looking for markets with strong barriers to entry. Barriers to entry increase the likelihood of recouping the capital invested. To date, capital has been the only major barrier to entry for ISPs.
- **Loyal user base** Investments will not flow into companies that have not addressed the high rate of customer churn. It is far more expensive to secure an additional customer than it is to retain an existing one. In addition, the number of additional services purchased tends to correlate with the amount of time a customer has been with a provider.

Understanding and addressing the needs of your most important external stakeholders is essential. You must address those needs in order to produce a marketable product and a sustainable business model. It is also vital to the success of the message you bring to the market and the investment community. While the list of items you must address is long and relatively diverse, value-added services provide opportunities to overcome each.

## The Evolving ISP

The evolving ISP must overcome the issues that are facing its core business, the demands of its customers, and the demands of the investor community. A few of these items will have to be addressed through a commitment to deploying leading-edge technology, but the vast majority of the issues require the addition of service offerings.

Deployment of the correct technologies to address issues of performance, reliability, and improved installation times will be no small feat, but overcoming these problems should only be seen as the first step. Developing the correct services, and deploying and marketing them correctly will be the difference between succeeding and failing (Figure 2.3).

**Figure 2.3** Chart of Issues and Demands



## The Steps Necessary to Offer Value

Among the first required steps to migrate to value-added offerings is to develop a highly reliable service model. Whatever the access method, connectivity must be maintained virtually 100 percent of the time. As additional services are deployed across that same link, reliability becomes exponentially more important.

Customers are coming to view existing applications such as e-mail and Internet access as mission critical. As additional business functions are accessed across that link, downtime will become increasingly unacceptable.

Developing customer confidence in the reliability of your service offerings will be an important prerequisite for the widespread acceptance of those offerings. Significant reliability will also need to be built in to each of your applications.

Existing applications such as MP3s, streaming audio, and streaming video are already driving demand for additional bandwidth, but the variety of services about to be offered by ISPs will continue to increase the demand for bandwidth. Vastly improved data rates will be required to support the variety of business applications offered.

Current implementations of hosted applications and Web sites are accessed across existing connections, sometimes with significant delay, but saturated links and latency will not be tolerated in the future. Applications and services will have to perform at least as well as today's LAN-based services. In the near-term, bandwidth constraints will be a major limiting factor of the market acceptance for remotely hosted applications and services. ISPs determined to develop this model will have to adopt high-performance access methods.

As performance and reliability improve throughout the industry, service providers must consistently bring that message to the public. Customers must gain greater confidence in the networks and experience the capabilities of remotely hosted services across intelligent high-speed links. They will quickly come to realize the great leap forward that has been taken. As the concerns of reliability and performance are quickly overcome, the remaining factors limiting acceptance of these services will be totally under the control of the existing ISPs.

The issues faced by the industry and the concerns of important external stakeholders, customers, and investors can all be addressed through the strategic deployment of the correct service offerings. It is important to also note that the right services alone will not guarantee success; they must be presented with the correct message and value proposition. In order to ensure long-term success, it will be necessary to deploy the right mix of services to address the market realities, customer demands, and investor demands outlined previously and in Figure 2.3.

## Configuring & Implementing...

### Gigabit Ethernet Brings Speed to the Edge

Over the past few years, major advances in technology have significantly increased the throughput of both core and LAN technology. While core technologies scaled toward 10-gigabit speeds and LANs experienced the availability of gigabit connections, the edge of the network has remained the domain of legacy Time Division Multiplexing (TDM) technology. DSL proved to be an exciting and inexpensive option, but holds little promise for major improvements in edge delivery bandwidth. The market is anxiously awaiting new solutions to address exploding demand for bandwidth and services.

High-speed IP over fiber appears to be one possible solution. Current implementations offer gigabit speeds, and 10-gigabit interfaces should be shipping from a number of vendors by the end of 2001. To put that in perspective, an OC-192 provides just less than 10 gigabits per second. Additionally, all vendors will likely announce adoption of the IEEE 802.17 standard for resilient packet rings. In addition to offering speed and improving reliability, IP-over-Photon architectures are excellent at supporting and managing value-added services. Current implementations offer bandwidth on demand, sometimes immediately accessible from a Web-based interface.

Surely, the demand for services is driving the adoption of high-speed Ethernet connections, but things will get really interesting when the bandwidth constraints at the edge of the network are relieved. High-speed connections should enable value-added services to reach their potential, because remote services will be provided across links of the same speed that users have come to expect from their LAN.

Proponents claim the IP-over-Photon architecture will effectively extend the LAN. Demanding services such as hosted applications, data storage, video on demand, and document management will be available as hosted services with no discernable difference from the LAN-based solutions in use today. Latency-sensitive applications such as voice and videoconferencing also enjoy improved service because of the additional bandwidth and the ability to be prioritized in their native protocol from end to end.

A variety of providers are now offering high-speed links using this IP-over-Photon architecture. Companies such as FusionStorm, Yipes Communications, Telseon, Cogent, and even SBC are offering connections with data rates ranging from 100MB to 1GB. MAN connections

Continued

and Internet access are available, sometimes on the same circuit. Predictably pricing pressure is already being felt. It is apparent that the economics in this segment will be no different from those that came before. Commoditization will quickly set in unless providers can cease selling the product as bandwidth and begin to sell the services to which such connections give access.

Initially, ISPs will have to address the issues of commoditized offerings, pricing pressure, and customer churn. These issues are in a large part responsible for the financial issues facing ISPs are highly dependent on the demands of customers/investors. It is clear that bandwidth services are highly impacted by commoditization, and that new bandwidth related offerings are quickly impacted by the same fate. DSL is an excellent example of this market reality. The only scalable alternative to this problem is the deployment of value-added services. ISPs must transition from selling the value of their bandwidth offerings to selling the value of the services that can be accessed across those connections. Competitors will have to clearly differentiate themselves based on the combination of services they offer.

#### NOTE

Some ISPs have postulated that network design services and/or hardware sales are the answer to the problem. While they may be the correct answer for some, they generally fail to address the underlying issues and simply mask the financial ramifications. Neither model forms a very good fit with the traditional core competencies of ISPs, and both require significant additions to the corporate structure and procedure. Additionally, neither model is particularly scalable, because they are highly dependent on adding resources to increase sales.

## Deployment of Services

Successful service providers will have to deploy services as the market begins to gain acceptance. Many offerings came too early and never reached their potential because more powerful or better-positioned competitors were able to improve on the offering as the product gained acceptance. Profitable offerings will address

problems of which the customer is acutely aware and willing to pay to solve. These services must be offered in a format and be accessed in a way that a significant number of potential customers are comfortable with.

ISPs will have to address two types of services:

- **Bundled services** These are the services that are included with existing offerings. They tend to offer limited benefits, but are used as differentiators. Services that we currently see used in this way are e-mail and Web hosting (often with a maximum allowable file size). These offerings are generally driven by competitive offerings already in the market and usually become commoditized very quickly. They do not directly increase margin or revenue, and they very seldom are capable of providing a sustainable competitive advantage.
- **Revenue-generating services** These are the services that we see ISPs focusing on to move forward within the market. These services will provide the potential for profitable revenue growth, and they are capable of bringing in additional revenue at a higher margin than traditional services. Revenue-generating services also allow ISPs to leverage their existing client base for additional revenue streams.

If developed properly they are excellent differentiators. Services such as application hosting or document management severely reduce the likelihood of customer churn as the customer becomes increasingly dependent on the applications and data controlled by the ISP. Services that we currently see showing revenue growth are managed VPNs, hosting services, and network monitoring.

Over time, we will see many offerings migrate from revenue-generating to bundled services as competitors strive to differentiate themselves by giving away desired and accepted services. These actions will perpetuate the problems faced by the industry, but must be expected. ISPs will have to account for this migration when developing offerings. It will be important to develop and deploy revenue-generating services that have some inherent resistance to the migration to a bundled service. Providers will also have to be cognizant of bundled services, as many will become required offerings as they spread throughout the industry and users come to expect them.

To combat the trend toward bundled services, ISPs must choose offerings that will not be quickly marginalized. Attractive services will have a degree of specialization, such as a focus on a vertical market or the capability to evolve as expectations increase. As with any product, it will be important to make sure that the

potential market for these differentiated services is large enough to support the initial offering and subsequent revisions.

It is also highly desirable to provide services that tend to increase their value as use continues. Document management services are a prime example, because as time goes on, it becomes harder and harder to migrate away from those services. Every day, the provider stores more data in their “proprietary system.” Additionally, over time, the provider is able to customize the front end for the customer to provide a better and more intuitive experience.

The best applications also employ technologies to solve problems in a way that is not available to the customer, or that is far more cost effective than existing options. Solutions that fit into this model are not priced based on the cost to provide the service, but as a percentage of the cost of traditionally managing the task. Again, document management is a prime example in this instance. Most small-to-medium-sized companies cannot afford high-end document management in-house, but they still create large volumes of documents. The value of outsourced document management in these instances is based on the cost of managing files of hardcopy documents and the benefits provided by digital storage methods (for example, the ability to search archives, ease of transmission, simultaneous use by multiple parties, etc.) Often times, the enhanced capabilities provided by the service will allow the provider to charge more than the customer is paying for the existing option.

ISPs should also target service offerings that are capable of evolving over time. If the service can continue to be developed, it will extend its differentiation and profitability. Application hosting provides the opportunity for almost unlimited evolution. As additional revisions and major upgrades of software become available, application service providers (ASPs) improve their services by implementing the upgrades. They also have the capability to offer alternative packages to replace existing applications if they fail to keep up with the industry. In this way, they can continue to evolve their service beyond the existing offering and promise the customer future-proof applications.

Evolving ISPs must also undertake the commitment to continuously improve their offerings. As they transition from selling bandwidth to selling services, the suite of service offerings becomes their value proposition. The makeup of those services is what will differentiate them from their competition, attract customers, and drive their revenues.

One thing that has been proven repeatedly is that technology will continue to advance and margins will be squeezed. Long-term success will be based on well-planned and continuous development of a suite of service offerings. The truly



evolved ISP will strive to develop a branded offering similar to today's portals. Providers will look to offer a suite of services through a single customized interface.

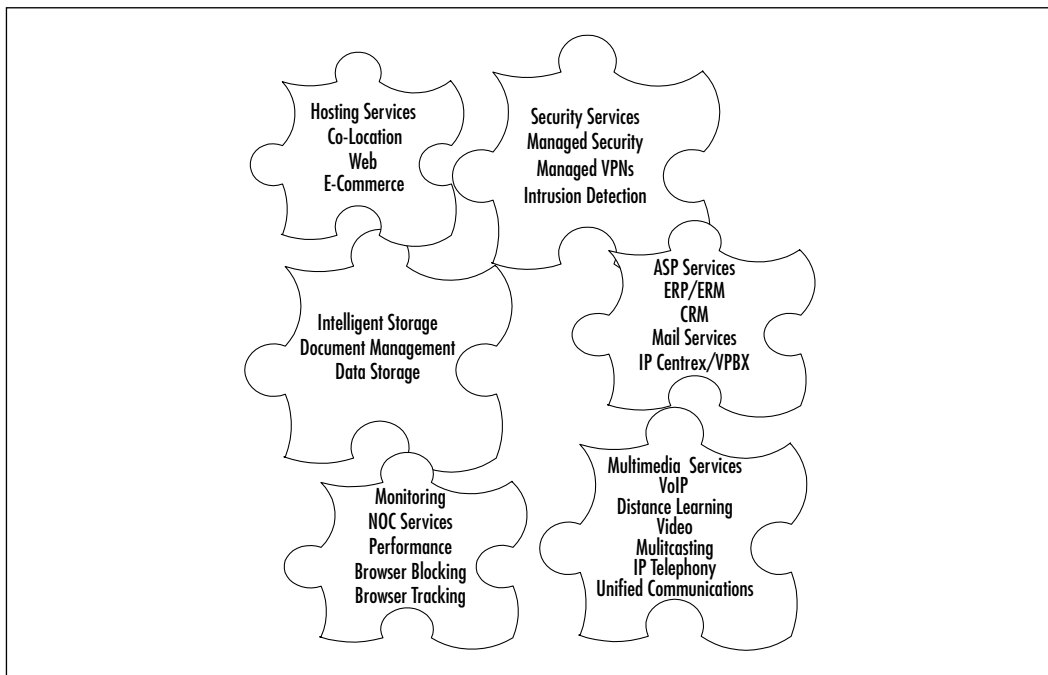
## Value-Added Services and Core Competencies

The high-level requirements for ISPs are relatively straightforward:

- Differentiate yourself from your competition
- Provide additional value, relative to the competition
- Develop additional high-margin revenue streams

It is also clear that the way to attain these goals is through the deployment of value-added services (Figure 2.4). What becomes more difficult is the development of a customized strategy for a specific ISP. In order to develop a successful strategy, each ISP must begin to assess a variety of internal and external factors. They must compare and contrast the various existing services and quickly come to understand the multitude of emerging services that have not yet gained mainstream acceptance.

**Figure 2.4** Potential Value-Added Services



The number of services available is almost limitless and will continue to expand rapidly. We prefer to allocate services into six major categories:

- Hosting services
- Security services
- Intelligent storage
- ASP services
- Monitoring
- Multimedia services

This list is by no means all-inclusive and does include overlap, depending on the implementation. For example, intelligent storage services are often times deployed in an ASP model. The special demands of intelligent storage have convinced us that they should remain separate from ASP services.

As a provider begins to understand the services available, it is important to understand the makeup of the organization. Each ISP will have its own core competencies and specializations. Success will be highly dependent on leveraging these skill sets effectively in the deployment of the initial offerings. It is also important to honestly assess areas that must be developed to attain future goals.

As the ISP continues to add additional services, it is important to continuously evaluate internal skills and determine those that must be developed to meet future requirements. The dominant companies of the future will continue to offer innovative services and address problems with new and varying skill sets.

Developing the correct offering will also require a deeper understanding of external factors and issues. It is important to understand your existing customer base in far greater detail than is conventionally required at this time. ISPs will have to develop an understanding of the needs and goals of the businesses they serve, and how to help these businesses succeed. They will want to understand what the demographics of their users are and how they are changing.

Successful product launches will generally be related to their relevance to the existing customer base. Existing users should be far more likely to add an additional service than new users will. Generally, the purchase of value-added services is correlated with the length of time that customers have been with your company.

Flourishing service offerings will also be dependent on thorough market research. While it is important to understand quantifiable factors such as potential market size, it is imperative to qualify the opportunity. When bringing new technology to market, it is always important to comprehend the willingness of the

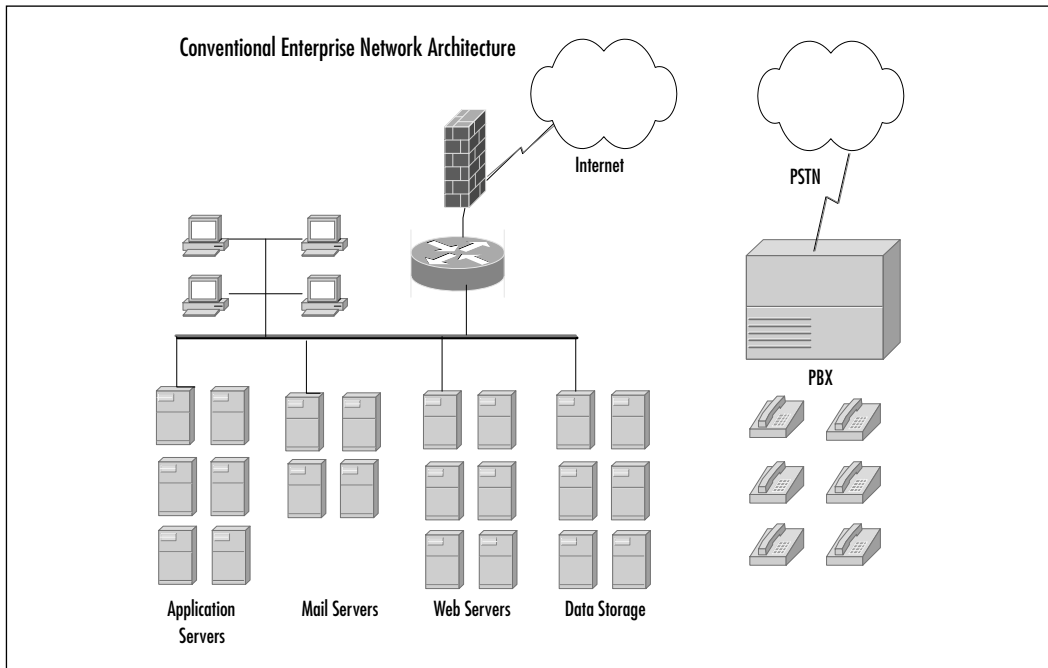
marketplace to accept such offerings. Successful offerings will address the requirements and concerns of potential customers.

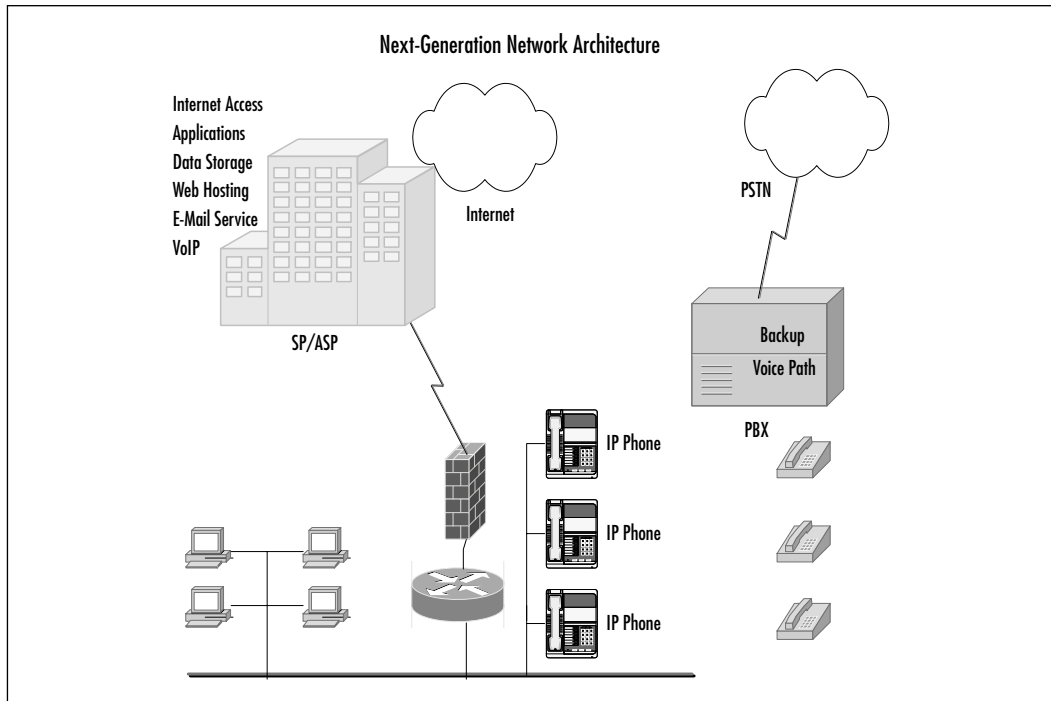
As an understanding of all of these factors are attained, it is important to develop an estimation of the cost to bring such an offering to market based on the current capabilities of your organization. To improve the likelihood of success, it will be important to thoroughly understand the internal, external, and economic factors involved prior to choosing a specific offering.

## The Service Provider of the Future

In the previous sections of this chapter, we focused on the current requirements faced by ISPs. While it is vitally important to understand the current realities of the marketplace, it is equally important to take some time to assess what the demands of the future will be. Without developing a short, intermediate, and long-term strategy, it is unlikely that your company will be able to escape the difficulties of the future. In addition to defining and addressing the current market requirements (Figure 2.5), ISPs must ask themselves what type of services they will need to have available in two, three, and five years in order to remain competitive and profitable (Figure 2.6).

**Figure 2.5** Current Enterprise Model



**Figure 2.6** Enterprise Model Using Next-Generation Services

As difficult as the current environment for service providers is, there is also more opportunity in this space than at any previous time. Technological developments and the proliferation of new services are providing a landscape that is ripe for consolidation and aggregation of services into a single medium.

These factors offer the opportunity for the most strategic players to steal market share and services from incumbent providers that are paralyzed by outdated technologies or fail to react to the opportunities available. Those who develop the correct offerings will be able to create strong barriers to entry, based on the size of their user base and the breadth of their services.

ISPs are well positioned to evolve into the dominant source of communications and data-related services to the home and business. All broadband providers that have the ability to use technological improvements to continuously increase available bandwidth and prioritize specific types of data streams will do well. This affords service providers the opportunity to aggregate various services into one connection. While it is unclear which technologies will dominate, either short or long term, many technologies are currently viable candidates, including cable, Ethernet, DSL, and wireless.

## The Finances Involved

The financial incentive to execute on this opportunity is enormous. The addition of the correct mix of value-added services allows for differentiation in a commoditized segment. Additionally, ISPs are on the verge of coming to market with a variety of technologies directly aimed at stealing market share from other industries such as long-distance, centrex, radio, cable, videoconferencing, and pay-per-view.

The combination of these services and data-related offerings such as application hosting, Web site hosting, and data storage will provide the opportunity to drive margins and revenue per user higher. By aggregating services previously offered across multiple connections to homes and businesses into one pipe, the strategic ISPs will also create strong barriers around their customer base. Customized user interfaces and portals, along with collaboration services among that growing user base, will create strong user dependencies and reduce user churn.

Those ISPs and service providers that do not grasp the value of the opportunity at hand, or those unable to adapt to that opportunity, will likely be marginalized over the next few years. As services converge onto effectively one connection, likely two or more physical connections for redundancy, only the providers offering all of the required services will have a viable offering.

Businesses and consumers will not purchase services from a provider that cannot include all required data, voice, and entertainment offerings. Over the next few years, providers who have not embraced new offerings and developed methods for continuously developing new offerings will not be facing commoditization, but extinction.

## The Case for Application Service Provider Conversion

The ASP offering is a revolutionary response to the inefficiencies in our current distributed computing environment. As we transitioned from mainframe architecture to a distributed model, there were great increases in productivity achieved by bringing computing power to every employee. This transition to the personal computer (PC) era enabled economic growth previously unheard of, and changed the ways in which we are able to transact business.

As corporate networks evolved, information became immediately available to authorized personnel throughout the world. As the business and technology evolved, we started to encounter difficulties. Largely enabled by widespread access to information, the business world now moves far faster than in it did in the past

and is capable of reaching all corners of the globe. Often, globalization becomes a requirement as progressive clients and partners demand support throughout the world.

The current speed of today's businesses and the extensive distribution of resources have uncovered some very important weaknesses that are inherent within distributed computing. As corporations add locations, the complexity of the required infrastructure increases exponentially, as centralized resources must be accessed, securely and reliably, from all locations.

Prime examples include Enterprise Resource Planning (ERP) or Customer Relationship Management (CRM) systems, with their complex back-end databases. In addition, today's business environment requires that new offices and short-term locations be set up quickly and efficiently. Employees who must travel also require consistent access to these resources. The current model is being severely tested by the conflicting realities of increasing complexity and shrinking deployment timetables.

Increasingly complex applications, infrastructures, and business realities are demanding improvements to the current distributed model without impacting its many beneficial aspects. Application hosting provides the answer to these concerns and a variety of others. By centrally locating these resources at an ASP, companies can leverage the many methods of access, both private and shared with security measures, and benefit from the economies of scale accrued by that ASP.

Application hosting enables ubiquitous availability of applications, while at the same time reducing costs and deployment times. In the near future, companies not using ASP services will be at a competitive disadvantage to their more forward-thinking competitors.

Forward-thinking ISPs should be very interested in the development and success of the ASP model. Internet service and various types of private connections will form the foundation of every ASP implementation. The potential for lucrative partnerships is enormous. The market is currently poised for explosive growth and widespread acceptance. The technology has proven that it can be reliable and effective, even in large-scale implementations. There is also a strong customer value-proposition based on reduced costs and increased capabilities.

Based on our research and the growth of some of the ASPs with which we worked, we recommend a more proactive strategy. Application hosting presents enormous potential for ISPs. It addresses many of the market realities that are currently plaguing the segment. Application hosting provides the opportunity to differentiate Internet connections and create additional high-margin revenue streams.

There are also significant synergies between the innate capabilities of ISPs and the requirements of application hosting. We feel that major opportunities exist within the segment for companies with established capabilities and familiar brands.

While many companies have entered the segment, the first-mover advantage has not been significant. There is still an enormous opportunity for established players from related industries to enter the market space and leverage their current offerings and reputation. The numerous early entrants have completed much of the customer development of the industry.

Consumers are aware of the segment and its various benefits. Success in application hosting will also position ISPs for further evolution up the value chain. The skills developed implementing an ASP model will provide access to various other lucrative service opportunities, ranging from professional services to managed services as those discussed in Chapter 1, “An Introduction to ASPs for ISPs.”

## Market Factors

The attention on the ASP industry over the past few years has been significant, numerous studies have been completed, and many projections have been made. Due to the early stage of the industry, those projections are varied. International Data Corporation (IDC) placed worldwide ASP spending at \$300 million for 1999 and estimated spending of \$7.8 billion by 2003 based on 92-percent compound annual growth. Many other companies have projected much higher figures.

Evans Marketing Services (EMS) and DataQuest project that worldwide spending on the ASP model and its services will be in excess of \$20 billion by the year 2003. Wherever the true number falls, the opportunity within the space is undeniable and will be largely dependent on the quality of the offerings and messages brought to market. One thing that is clear is that strong leaders will emerge within the industry between now and 2003. Potential ASPs must bring their offerings to market quickly, or face increasingly stiffer competition.

It is likely that the most avid purchasing of ASP services will be by the small and medium-sized business segment. According to the Yankee Group, this segment includes 10 million companies and accounts for 98 percent of businesses. This same group is responsible for half of the U.S. Gross National Product.

In the current environment, these companies have limited application options without enormous expenditures. The opportunity to bring affordable solutions to this market segment is enormous. The Yankee Group calculates small to medium-sized business IT spending at \$45 billion annually. The growth rate of that spending should increase significantly as those companies gain access to high-end solutions previously not available to them.

The ASP industry will be able to offer productivity-enhancing packages that were previously available only to large enterprises. The justification for increased spending on productivity-enhancing solutions as such sales force automation (SFA), financial packages, or document management should be persuasive.

## NOTE

Small businesses are generally categorized as those having less than 100 employees, while medium-sized businesses generally have between 100 and 1000 employees.

The opportunity within the small and medium-sized business market should mesh well with the existing user base of most ISPs. This market will be most interested in applications such as ERP, payroll, human resources (HR), SFA, and remote learning in the near term. The less-than-100-employees segment should also show a strong interest in messaging and backup/storage services.

Based on the projections of Forrester Research, the small- and medium-sized business segment should provide in excess of 80 percent of the spending on ASP services. While the majority of the revenue will come from the smaller companies, major opportunities will still exist within the larger corporations. We are sure that larger enterprises will show interest in remote learning, and various applications to support remote and branch offices as well as partners. The speed of deployment of applications and services to new sites will be major factors for the larger businesses.

As we discussed previously, the market appears to be ready for widespread adoption of ASP services. Increasingly complex applications and networks combined with ever-faster speed-to-market requirements, will help to position application hosting as a necessary offering. However, to date, à la carte providers have appeared to struggle in developing a successful sales model.

This is in part due to the infancy of the industry and the difficulty of educating the public. It can also be attributed to limited customer understanding of the technology and the new companies that are marketing the services. ISPs should have significant advantages over their predecessors based on their client relationships, established brands, and existing sales forces.

In some cases, ISPs that have added ASP services have reported increases in revenue of 50 percent or more. When you take into account that a simple suite of desktop applications in a hosted environment can cost roughly the same amount



per user as a business-class DSL solution, you can see the power of the economics. A customer who chooses to purchase such a product for 25 employees over an existing DSL connection will increase his or her monthly bill by 2600 percent.

As customers become used to the convenience of hosted applications, we should see a strong trend toward additional services and the bandwidth required to support it. Additional revenue opportunities are also available to those providers that offer professional services to their customers. Customization of applications and the integration of applications with existing in-house solutions will drive upfront and monthly costs significantly higher.

## ASP Customer Value Proposition

Some of the larger industry projections can seem unlikely, but as you come to understand the magnitude of the customer value proposition, they begin to look almost pessimistic. Today's companies are facing an unforgiving environment. Technology has increased the speed at which all companies must do business. Competitive advantages and barriers to entry tend to vanish at an alarming rate.

Investors have access to real-time information, so they act quickly at the first signs of weakness within a company's business model. While companies strive to address these realities with innovative deployments of technology, they are finding that there is a drastic shortage of skilled resources to design and implement these solutions. When these resources are available, it is at great cost, and retaining them tends to be as difficult as finding them.

The ASP value proposition does an excellent job of addressing a number of the issues that companies face today. We feel that the ASP model is an excellent evolution of the existing distributed computing model. As the industry begins to gain prevalence, it will become increasingly difficult for companies to *not* use the advantages provided by ASP services.

- **Focus resources on core business objectives** The ASP model allows companies to reduce their Information Technology (IT) related responsibilities. With the purchase of an ASP service, the provider controls complex applications and the required hardware. Additionally, the ASP provides all support for the application and hardware.

Hours previously devoted to system design, deployment, and support, are no longer required to be handled by the company's internal staff. Additional savings are also gained, as the resources devoted to hiring, training, and managing many IT professionals are no longer

required. By outsourcing these functions, the company as a whole can focus on meeting the core business objectives.

- **Overcome difficulties of IT staffing** The shortage of skilled IT professionals is having a severe impact on many layers of business. The search for qualified applicants is consuming a significant amount of internal resources. When applicants are found, they command high salaries and expensive ongoing training. The turnover rate in such positions is high, and very costly. While these realities are affecting all businesses, they particularly impact the small and medium-sized business.

ASPs enable businesses to reduce their internal IT staffing requirements and further focus the attentions of their existing resources. Employees are able to better target their efforts and training once complex applications are outsourced. The end user generally is also rewarded with improved system reliability and better support.

- **Speed in application development and deployment** ASPs are able to leverage experience and existing infrastructures to rapidly deploy new applications. In many instances, implementations that previously took a year or more are now available in months, if not weeks. Companies are able to leverage the advantages of new applications far more quickly and minimize the involvement of their staff in these previously draining deployments.
- **Access to the newest technologies** ASPs are able to leverage expertise and strong independent software vendor (ISV) relationships to make new technologies available to the market. Only the most capable corporate IT staffs can keep up with the multitude of new applications and upgrades that are available.

The low upfront costs and simple user-based billing of the ASP model allows far more companies to leverage these new technologies. Many companies, particularly small to medium-sized companies, would not have access to these options in any other way.

- **Improve time-to-market** By significantly reducing the required implementation time of mission-critical applications such as financial packages, ERP, or CRM, ASPs can enable companies to bring their offerings to market much faster than was traditionally thought possible. Advanced applications can be up and running in a fraction of the time

that in-house deployment takes. This is extremely important in both new and established companies.

- **Improve productivity** Advanced applications enable significant increases in employee productivity, but traditionally require large investments of resources and capital. The ASP model can provide these advantages in a much more reasonable timeframe and cost ratio. Implementation costs are significantly reduced, and ongoing costs are highly predictable. The ASP model can drastically reduce the required staff for design, implementation, and support, and can also provide application implementation and support. Small and large companies can easily leverage the productivity-enhancing capabilities or advanced applications with the ASP model.
- **Scalability and agility** ASPs allow their customers to efficiently and easily add users and locations. ASPs are able to offer far more scalability than in-house implementations can, allowing their customers to always be prepared for hyper-growth. The ASP model of Web-enabled applications also meshes perfectly with the needs of road warriors and remote users. Companies that use ASP applications can quickly and securely offer new locations and user access to important applications, without expanding the corporate network.
- **Cost advantages** The ASP model can provide both short-term and long-term financial benefits. The initial cost of deployment is far lower within the ASP model. Additionally, ASPs are capable of leveraging vast economies of scale to provide applications to companies, at more attractive costs than they would be able to achieve internally. The total cost to purchase, implement, and support an advanced application is generally far higher than most companies realize; particularly when you include the cost of hiring and training the required resources.

## ISP Value Proposition

In the current environment, it is very important that service providers address the market realities they face. A potential offering must be assessed in relation to its abilities to overcome market factors. It must also be understood that these factors may change or expand over time and will need to be understood on a continuing basis. As we discussed previously, the following are the most prevalent market factors facing ISPs:

- Commoditized offerings
- Significant pricing pressure
- High customer churn
- Drastically reduced valuations
- Restricted access to capital

Incorporating ASP services into the existing ISP model will enable providers to address these factors. Application hosting provides the opportunity to differentiate your services, command higher margins, and increase your customer's dependence on your services. Application hosting will also increase the demand for, and dependence on, your original bandwidth offerings. Bundled offerings will provide the opportunity to differentiate your current commoditized connections based on their integral role in the delivery of high-value services. As with all strategies, success will depend on the quality of execution.

Application hosting will immediately begin to differentiate your services from other ISPs that have not yet rolled out ASP services. Your sales teams, or channels, can now transition from selling bandwidth alone to offering the potent value proposition of hosted applications.

Ultimately, the connection will become a supporting offering to your value-added services. ASP services should enable differentiation of your connections during the sales process, even to those prospects who are not yet ready to purchase hosted applications. Such prospects should understand the differentiation of your offering once it includes the strong value proposition of hosted applications. Additionally, you can differentiate your services from existing ASPs because you control your own network infrastructure. Providing an end-to-end solution enables customers to rest assured that they will only have to make one telephone call in the event of a problem.

Differentiation will also provide the most effective way to combat pricing pressure. ISPs offering ASP services will have the best chance of fending off pressure on their existing offerings. They should also be able to leverage the value of an end-to-end solution, and command either a premium for the offering or an increased closure rate.

It is important to continually retool the corporate message to focus on the value of new offerings. Providing a persuasive combination of offerings will be key to sustaining advantages over competitors. The correct suite of products will command premium pricing in comparison to individual offerings. ASP offerings will also allow ISPs to improve margins with existing and new customers.

Existing customers will be marketed the new offerings, thereby increasing revenue and margin per customer. New customers that opt to include ASP services will also drive up revenue and margin calculations.

ISPs have to address customer churn, as it is only likely to increase if measures are not taken. The ASP model offers one of the most potent methods for reducing churn. Users quickly become dependent on their applications, particularly advanced applications like CRM or ERP. The familiarity user's build with a product and even limited customization prove to be strong reasons to continue with an application.

The accumulation of data retained within the network of the ASP is also a barrier to switching providers or application. It is difficult and expensive to migrate data and customize new applications. It is also extraordinarily expensive to retrain users and roll out new services in a way that is readily accepted by internal stakeholders. The quickly developed dependency on hosted applications is one of the most significant advantages of ASP services over other potential offerings.

Addressing the demands and realities of the capitals markets is no small undertaking. The preferences of investors can shift dramatically as signs of weakness appear. However, investors are always a critical key to success in capital-intensive industries such as those of ISPs and ASPs. Successful service providers in all segments will have to overcome the low valuations and limited access to capital that exist currently.

In order to do so, they will have to take on a number of factors, including profitable revenue growth, increasing margins, differentiation, barriers to entry, and development of a loyal user base. We touched on most of these factors previously in this chapter, but it is important to discuss one of the issues a little further. The ASP model allows ISPs the opportunity to form barriers to entry never before available. By forming strong relationships with ISVs, it will be possible to gain definitive barriers to entry. If providers are capable of creating a strong value proposition for the ISVs, they will be able to negotiate limits on the number of providers. Barriers can also be created based on the breadth of services offered by the early leaders. Those ISPs and ASPs that move quickly, and choose the correct suite of offerings, will be able to develop a lead with which later entrants are unable to compete.

ASP services may be the only product offering that truly allows ISPs to address many of the required factors. As technology develops, other categories may be capable of addressing these issues, but currently our research leads us to the conclusion that hosted applications provide the most persuasive single offering.

ASP services can also be developed to include applications from other service areas. Providers could add document management, storage services, or multimedia

offerings based on many of the same platforms and core competencies they developed for their initial ASP offerings. The ability to continuously evolve offerings and market them to existing and new customers is one of the main strengths of the ASP model.

## ASP Services Also Enable Future Migration Up the Value Chain

This helps ISPs to position themselves to overcome the continued commoditization of offerings. The skills that are used to design, implement, and support applications also provide an excellent foundation for entrance into professional services or the managed service provider (MSP) industry. ASPs are in the perfect position to offer their customers application deployment and customization.

Integration of various applications or existing in-house systems also offers highly profitable work. These professional services command premium rates in the marketplace and allow providers to build out their offerings while leveraging existing core competencies. Their hosting and application experience also makes them an ideal candidate to provide MSP services for those companies that want more specialized application support.

Lateral migration to Web hosting is another possibility that would leverage many of the same skill sets currently used by an ASP. We feel that the ability to develop skill sets that benefit your existing offerings and also prepare you for future migration up the value chain is an enormously valuable component of ASP services. Continuous evolution will be the key to retaining differentiation and higher margins.

## ISP to ASP: The Perfect Fit?

ISPs have to deal with a number of internal and external factors that will provide them with advantages over other entrants into the ASP space. These factors are central to the logic behind adding ASP services to the ISP model. ISPs that add ASP services should have significant advantages over existing competitors with similar offerings.

ISPs are able to leverage core competencies in technology and sales. They have a known brand and a valuable customer base to ease market penetration. ISPs also enjoy differentiation from their competitors based on their inherent end-to-end delivery solution. By incorporating ASP offerings into their existing ISP model, providers are able to enjoy a certain synergy that differentiates them from competitors in either individual segment.

ISPs have an enormous asset in their existing customers; they are a known entity to these companies and already have a path of communication open with them. Most of the companies probably fall in the highest growth area for ASP services, small to medium-sized businesses. It is important that the demographics of your existing customers be understood prior to making application decisions. You would be well advised to take the needs of existing customers into account when choosing preliminary applications. These customers will be the most likely to respond to initial contacts and provide very important references to later customers.

ISPs also have the luxury of existing sales models. Development of a highly successful sales model is one of the primary difficulties that current ASPs are facing. The ability to build on existing and proven methods is an extraordinary advantage for ISPs. You should not overlook this challenge because your staff will require training and further development to sell complex services, but existing resources are a significant asset to have.

It will also be important to bring in more advanced application specialists and develop an application-focused sales engineering group to support the existing resources. Another key advantage for many ISPs will be their established brand and marketing efforts. Penetration into new markets is always easier when you are established in a related field.

ISPs also have a variety of technical capabilities to ease the transition to ASP services. C.E. Unterberg, Towbin define the ASP model as including seven service layers:

- Access network connectivity
- Shared infrastructure—data center and backbone
- Collocation hosting
- Application infrastructure hosting
- Net-hosted application
- Implementation/business process design
- Ongoing application management

The first four layers are the core competencies of services providers, while the last three are specific to the ASP model. The access network connectivity and shared infrastructure layers are truly the domain of the ISP. These first two layers are essentially the same services that ISPs currently provide. The next two layers, collocation hosting and application infrastructure hosting, are services that many

ISPs provide as well. Many ISPs have rolled out hosting services. Those providers are perfectly positioned to leverage those services and move into the higher layers.

While a large percentage of ISPs do not yet provide hosting services, virtually every ISP does host some services currently. Internal server farms are generally hosted by ISPs to provide e-mail and Domain Name Services (DNS). More recently, many ISPs have added Web site hosting, calendaring, and instant-messaging capabilities to these server farms. While ASP services require a more demanding environment, these initial steps provide important understanding of the hosting environment.

The uppermost three layers of Figure 2.7 comprise the value-added services provided by ASPs—Net-hosted software application, implementation/business process design, and ongoing application management. These layers are predominantly new territory for ISPs. ISPs making the transition will have to focus recruiting and retention on these areas to develop their core competencies.

**Figure 2.7** ASP Service Layers

ASP Layers	<b>Continuous Application Management:</b> Application Enhancement and Upgrades; Monitoring and Support
	<b>Business Process Design and Deployment:</b> Application Implementation; Client-Specific Customization
	<b>Web-Hosted Software Application:</b> Development of Web-enabled and Web-Based Applications and Templates
SP Layers	<b>Application Infrastructure Hosting:</b> Server/OS Deployment and Monitoring, Glue-code, Application Monitoring
	<b>Collocation Hosting:</b> Racks, Network Infrastructure, Bandwidth Allocation
	<b>Shared Infrastructure-Data Center and Backbone:</b> 24x7 Support, Security, Networking, Network Management, Bandwidth, Storage
	<b>Access Network Connectivity</b>

While it will not be a simple task to develop these skill sets, ISPs have so many other advantages that they should be well positioned for success. These upper layers also require some skills that ISPs currently have, like an understanding of 24 x 7 service and responsiveness. Existing customer support staff and call centers are major advantages for new entrants.

The net result of all of these factors should provide competitive advantages to those ISPs that decide to enter the ASP market. Their existing customer base,



brand, and sales/support models should be enormous advantages. They will be required to add to their technical capabilities, but should have a solid foundation in at least the first four layers as determined by C.E. Unterberg, Towbin.

The accountability that ISPs can provide for end-to-end service delivery is also another enormous differentiator. Many ASPs might have to outsource their lower-layers capabilities and may have to struggle to gain a solid understanding of the infrastructure. The opportunity to differentiate services based on the access connectivity also exists for ISPs. Because they control the end-to-end solution, they should be far more prepared to role out prioritization services for different applications and demand premium pricing for the service.

## Critical Success Factors

We have spent much of this chapter covering the many reasons why ISPs would want to add ASP services to their product offering. While the value proposition for these services is strong and the synergies between the ISP and ASP model are undeniable, it is always important to carefully review all major strategy changes. The risks related to a transition of this magnitude should be considered at least as significant as the rewards. It is important to mitigate these risks in every way possible.

To that end, providers should clearly understand and address the critical success factors involved in the ASP model. The major success factors involved in application hosting revolve around customer decision criteria, potential business models, technical implementation decisions, and application offerings. We will address technical issues throughout the remainder of the book, but we will focus on the other success factors in this section.

## Business Models

ISPs that are convinced of the opportunities the ASP model provides have a variety of ways to capitalize on the opportunity. We can classify the vast majority of these options in three broad categories:

- Application infrastructure provider (AIP)
- Partnered deployment
- Independent provider

Depending on your core competencies, risk tolerance, and available capital, any of the three options could be right for you. In some instances, an ISP could employ more than one strategy, or develop blended strategies. Each option has

inherent benefits and risks. Generally speaking, the greater the risk involved in the strategy, the greater the potential reward. We will discuss each strategy in order, starting with the least amount of risk.

Most ASPs are not looking to develop their own network or hosting facilities. This fact can be related to time-to-market issues, capital requirements, or core competencies. Often times, ASPs see the application integration as their core competency and they look to offload the underlying services.

*Application infrastructure provider (AIP)* is a term used to describe a provider that offers ASPs wholesale network and data center services. ISPs are in an excellent position to offer this service. They can provide network access and hosting experience. This is a way for ISPs to migrate up the value chain in a very controlled manner while developing additional revenue streams.

Some AIPs are offering value-added services to differentiate their offering and more completely develop their value proposition. Some of these services include:

- System and application monitoring
- Provisioning
- Billing
- Server and operating system management

The second strategy, partnered deployment, involves developing a partnership with a systems integrator or software vendor. This strategy attempts to leverage the synergies of a service provider's network and facilities with the application expertise of another company. Some of the notable examples of partnered deployment include Qwest Cyber.Solutions, a joint venture between Qwest and KPMG, and Sprint's partnership with Deloitte Consulting.

Partnered deployment can mean improved time to market and reduced capital expenditures for each partner. Partnerships can be very tricky and time consuming to keep productive and viable. Many well-conceived partnerships have fallen apart because of friction between parties. Additionally, shared ownership means less of the total return for each partner.

The final option is to deploy an ASP offering as an independent provider. This option will require the greatest capital expenditures, and is likely to require significant hiring and retraining of existing staff. Many new skills will have to be purchased and developed. Support that cannot be found or developed can be outsourced to a variety of professional services organizations.

This strategy also provides the greatest potential return and the most significant differentiation of your existing services. Retaining total control of the

offering and keeping it within the original organization provides the opportunity to market new services as a direct means of differentiating all offerings under your brand. Additionally, this strategy does not require the challenges that are inherent in the integration of two separate companies.

Once an organizational model has been chosen for your entrance into the ASP segment, it is imperative that a solid sales model be developed. You should make every attempt to leverage existing resources and customers. Existing relationships with customers will be one of the most important assets that ISPs bring to the ASP model.

ISPs will also have to develop the skills of their sales forces to enable them to clearly articulate a more complex value proposition. Significant opportunities exist for ISPs through the development of new sales channels. Channel opportunities exist with many of the players in the segment, including ISVs, system integrators, Original Equipment Manufacturers (OEM), and Value Added Resellers (VAR).

Leveraging channel partners with complementary offerings can be very effective, but these channels must be managed differently from direct sales methods. Successful channel sales require a great deal of support and education to keep representatives motivated to offer your product, and must be capable of persuasively presenting your value proposition in addition to their own.

## Determining Your Offerings

The choice of offerings—in this case, applications—is critical to the success of any strategy. Determining which applications to offer, and in what order, will be one of the most critical tasks undertaken during the evolution of your offerings. Most business applications sold today are capable of being hosted for remote access.

The software industry has successfully undertaken the call to Web-enable applications over the past few years. The viable application categories cover a wide range of services, including:

- E-commerce
- Finance
- Human resources
- Procurement
- Logistics
- Supply chain management (SCM)
- CRM

- ERP
- Collaboration
- Storage
- Document management

With multiple ISV vendors in each category, the total number of potential applications is enormous. IDC has developed an effective model for beginning to categorize potential offerings. It delineates potential hosted solutions into seven categories:

- Analytical applications
- Vertical applications
- Enterprise resource management
- Customer relationship management
- E-commerce
- Collaborative applications
- Personal applications

The IDC model ranks the categories based on the level of complexity and required level of business process design. The higher-level applications are more difficult to implement and require significant business process capabilities, but provide the potential for improved margin and additional professional services. It is generally a good idea to focus your search on a limited number of these categories based on internal and external factors.

Determination of these target categories should take into account a number of factors. The most important of these include the core competencies of your organization, the demographics of your customer base, and the current and potential market acceptance of that category. As with most instances, the higher levels provide significant opportunities, but also increased investment and greater potential risks. Probability of success in a given category will be highly dependent on your ability to leverage existing skill sets.

It is important to honestly assess the core competencies and specializations of your organization, as well as the areas that need to be developed. It is advisable to uncover your strengths and weaknesses, and include or exclude categories based on internal capabilities.

The next step is to develop a solid understanding of another significant asset, your existing customers. It is important to leverage every advantage available when entering a new market segment, and existing customers are an excellent differentiator when entering the ASP segment. Important information to track on customers includes industries, employee count/growth, revenue/revenue growth, locations with employee count, and expansion plans.

As you gain an understanding for the demographics of your user base, you will be able to make some assumptions about the potential near-term revenue streams of various categories or applications. As potential categories and applications emerge, it is very important to accumulate information on the short- and long-term opportunity for that offering.

While it is important to determine quantifiable information such as potential market size, it is just as important to qualify the willingness of the market to purchase the offering. The ideal offering will fall at the end of the early adopter phase and the beginning of the early majority growth period. Next, it is important to evaluate the opportunity by comparing the cost to bring such a service or services to market with your qualified and quantified projections on the services potential.

As you determine the categories for which your company is best suited, it is important to begin accumulating information on the various ISVs within that segment. It is important to rank these potential partners on factors such as:

- Market penetration
- Innovation
- Financial viability
- ASP strategy
- Application capabilities

It is also important to find an ISV that shares your vision and is willing to develop a mutually beneficial relationship. The best partners will share a commitment with your organization and develop a revenue-sharing model. ISPs should leverage their existing brand, track record, and customer base in ISV negotiations, as they are significant differentiators and should allow better terms to be negotiated.

As with any partnership, confidence in the ability to work amicably with the people across the table from you over a long period of time should be held in high regard. Intangibles are a major component of successful partnerships.

When assessing an application to offer as a hosted service, it is important to clearly understand the potential market for that application. It is infinitely valuable

to uncover problems facing businesses or a specific vertical and bring a solution to market that addresses those issues.

It is particularly advantageous to do so when similar options are not available or not cost effective for that specific market segment. ERP and CRM applications are excellent examples of applications that address existing business problems, but are price prohibitive for many small to medium businesses.

The opportunity to purchase such applications at an acceptable price will only be available to these businesses in a hosted environment. Due to the efficiencies enabled by such packages, the ASP providers of these products will have an excellent value proposition. There are also financial advantages to bringing specialized applications to market initially.

It is easier to develop expertise when focused on a more specialized application. In some instances, this strategy requires that you limit the number of applications offered or focus solely on solutions for one vertical market segment. By doing so, your company will be better able to cultivate market-leading expertise. It can then leverage that expertise to steadily reduce the cost of adding each customer. Document management is such a service for many of these vertical markets.

Companies can leverage their understanding of their application and a specific market segment to truly understand the needs of their potential customers. The ideal application in this instance would have the capability to be ported to other industries as the company's capabilities grow. It is important to remember that your offering must continue to evolve and expand, or it will ultimately succumb to the pressure of commoditization.

## Customer Issues

As with any industry in their early stages, the addressing of the needs of potential and existing customers is going to be an important critical success factor for the ASP industry. As we discussed previously, choosing the correct offerings and bundling them in a way that is valuable will be critical.

It is imperative that ASPs develop an understanding of the decision criteria used by potential customers and the most important criteria used to judge service levels by existing customers. Successful providers will have to focus on these issues to improve their market penetration and customer retention. Customer-satisfaction levels and reputation will be important success factors in the ASP segment.

Current Analysis published the results of their survey of ASP customers that ranked the major decision criteria they used to choose an ASP provider. Major factors included:

- Support Capabilities
- Hosting and Facility Experience
- Cost and Pricing Structure
- Reputation and Client Reference
- Service Level Agreement
- Past Performance
- Scalability and Completeness of Solution

In order to successfully market to potential customers, it will be important to address these issues.

ISPs should be at an advantage because of their existing support capabilities, track record, and reputation. It would be advisable to develop reference accounts from within the existing customer base in the initial stages of the ASP rollout. A successful value proposition will have to incorporate these criteria and should be supported by marketing materials focusing on the same issues.

Many of the same issues will remain a factor for customers after they have signed on. While customer satisfaction is always an important statistic, it is even more important in a developing industry. In order for providers of hosted applications to succeed, they will have to prove that the technology is ready for primetime. One of the most persuasive ways to prove that fact is with high customer satisfaction. The main components of long-term customer satisfaction should include four main factors:

- **Support** Customer support will be a very important component of every successful ASP offering. Much of the ASP value proposition is based on the reduction of internal IT requirements. To achieve that goal, the provider must assume the responsibility for those services.

ASPs must be prepared to manage the system end to end. They will be confronted with application issues, network issues, and desktop issues. This means a broad spectrum of support services must be offered on a 24 x 7 basis. In some instances, the ASP will be required to work with the ISV to troubleshoot application failures and issues.

ISPs have a significant advantage over other providers because they control and manage the network. This allows them to control the end-to-end offering. Successful ASPs will understand that their offering is truly a service, and strive to differentiate their services based on the support they provided around that service. It should not be overlooked that

the number-one issue in the current analysis survey was support capabilities. Do not expect that ranking to change.

- **Performance** Performance is a major issue facing ASPs. Users are demanding performance on par with internal implementations. In many instances, they are expecting better performance than they would receive from an internal application across a WAN link. Some ASPs have tried to address these issues with thin client solutions or direct connections.

In this instance, the strong pricing pressure on bandwidth is an advantage. Many customers are simply failing to invest in their own infrastructures. Link speeds and LAN implementations are too slow for the increasing demands being placed on them. As 10/100/1000-megabit WAN connections become widely available and are adopted by businesses, many of these problems will disappear.

Software applications are helping ASPs gain an understanding of user issues and determine the exact cause. Mercury Interactive ([www.mercuryinteractive.com](http://www.mercuryinteractive.com)) is one such company whose products allow sessions to be tracked from end to end. Statistics and information are gathered from the server through the network to the user's desktop. Problems can even be determined and reported within the service provider cloud. Tools such as these are extremely effective for documenting the source of user issues and proving the provider's innocence when problems exist within the customer's network.

- **Pricing** Customers are looking for pricing models that are reasonable and predictable. While pricing methods in the industry do vary, most providers charge an upfront integration and installation fee, and then bill services on a flat monthly per-user charge.

Some ASPs are offering bundled pricing that does not require an upfront charge but spreads the cost over higher monthly fees. In specific instances such as e-commerce hosting, ASPs have offered risk/revenue-sharing arrangements.

It is unclear what the long-term pricing model will look like, but it is obvious that pricing will remain an issue. Customers will look for a solid and measurable return on investment (ROI) for ASP implementations. Those offering ASP services will have to struggle with the balance between customization and pricing. An efficient balance that allows cost advantages over traditional implementations while addressing the general needs of most potential customers will be required.



- **Security** Security will remain a major concern for all businesses, and will increasingly become a residential issue. Companies hesitate to allow outside connections into their network. They will also fight the migration of their mission-critical information to the data centers of others without clearly defined security practices.

More importantly, ASPs cannot afford a security breach of their customers' data. Such an event would likely receive media attention and destroy the future of that provider, as well as impact the industry as a whole.

Strong measures must also be in place to stop traffic from one customer from gaining access to other customer's connections. In the most severe instance, an experienced hacker could gain access to the ASP network through one customer's link and then infiltrate the networks of all of your customers.

Security measures will have to be taken to address the concerns of potential customers as well as potential threats to the viability of the ASP model.

## Summary

ISPs must quickly come to understand that *Internet* is no longer the most important word in Internet service provider—*service* is. ISPs must bring services to market that combat the financial realities their current offerings are facing. They must efficiently deploy and market impact services that drive profitable revenue, attract customers, and retain existing customers; services that offer customers improved capabilities, increased efficiency, and definable financial savings. The services must address real problems and demonstrate value for the end user. Services that seem to be ready to handle these demands are voice over IP (VoIP), unified messaging, managed global VPNs, and application hosting, among others.

It is our belief that application hosting provides the greatest opportunity for ISPs to address the market issues facing them. ASP services can be used to address issues as diverse as commoditization of existing offerings, pricing pressure, customer churn, market valuations, and access to capital. ASP offerings also have two other important advantages; they are based on proven technology and provide a highly persuasive value proposition to potential customers.

The ASP model offers companies services and capabilities that would not otherwise be available to them. The ability to deploy productivity-enhancing services quickly will prove to be a requirement for most businesses. The scalability and agility inherent in the ASP model will also be required by enterprises and fast-growing companies to support remote offices, users, and corporate partners. It is our belief that the capability to rapidly deploy and scale applications will drive the growth of ASP services in all sectors at the expense of traditional methods.

As ISPs roll out ASP services, they must carefully address the main concerns and demands of potential customers: support, performance, pricing, and security. All of these issues must be addressed in a successful ASP strategy. Addressing these issues should be the foundation of your offerings, as the success of your company and the widespread acceptance of the industry will likely depend on your handling of these issues. The solutions to these concerns should be addressed in your marketing message and value proposition; they will drive customer acceptance and satisfaction.

As all true entrepreneurs know, the greatest opportunities come when the view of the future is at best, unclear. Traditionally, it has been times such as these when the brave and innovative have reaped the greatest rewards.

# Solutions Fast Track

## ISP Market Conditions

- ☑ Internet access reached 50-percent market penetration in less than eight years of existence. The growth rate in the United States is projected to be anywhere from 40 to 110 percent for at least the next few years.
- ☑ According to *Boardwatch Magazine*, there are currently more than 7700 ISPs (early 2001) that are doing business in the United States alone.
- ☑ The reality of the DSL market is that providers must rely on the Incumbent Local Exchange Carrier (ILEC) for the all-important connection to the customer. That forces ISPs into the position of commodity resellers in direct competition with their suppliers.
- ☑ While broadband connections seem to be following the same economic pattern as their slower counterparts, their significance should not be overlooked. Increasing broadband access speeds will be the foundation for the value-added services that will allow ISPs to differentiate their offerings.

## Service Provider Business Requirements

- ☑ In order to break out of the current cycle, many service providers and ISPs in particular will have to address these factors: commoditized offering, significant pricing pressure, high customer churn, drastically reduced valuations, restricted access to capital.
- ☑ The current demands of the financial community once again include traditional terms such as *differentiation*, *barriers to entry*, and *profitability*. The easy money is gone.

## The Evolving ISP

- ☑ The evolving ISP must overcome the issues that are facing its core business, the demands of its customers, and the demands of the investor community.
- ☑ Among the first required steps to migrate to value-added offerings is to develop a highly reliable service model.

- ☑ Current implementations of hosted applications and Web sites are accessed across existing connections, sometimes with significant delay, but saturated links and latency will not be tolerated in the future.

## The Service Provider of the Future

- ☑ ISPs must ask themselves what type of services they will need to have available in two, three, and five years in order to remain competitive and profitable.
- ☑ Businesses and consumers will not purchase services from a provider that cannot include all required data, voice, and entertainment offerings. Over the next few years, providers who have not embraced new offerings and developed methods for continuously developing new offerings will not be facing commoditization, but extinction.

## The Case for Application Service Provider Conversion

- ☑ The ASP offering is a revolutionary response to the inefficiencies in our current distributed computing environment.
- ☑ Application hosting presents enormous potential for ISPs. It addresses many of the market realities that are currently plaguing the segment. Application hosting provides the opportunity to differentiate Internet connections and create additional high-margin revenue streams.
- ☑ International Data Corporation (IDC) placed worldwide ASP spending at \$300 million for 1999 and estimated spending of \$7.8 billion by 2003 based on 92-percent compound annual growth. Many other companies have projected much higher figures.

## Critical Success Factors

- ☑ *Application infrastructure provider (AIP)* is a term used to describe a provider that offers ASPs wholesale network and data center services.

- ☑ Leveraging channel partners with complementary offerings can be very effective, but these channels must be managed differently from direct sales methods.
- ☑ Current Analysis published the results of their survey of ASP customers that ranked the major decision criteria they used to choose an ASP provider. Major factors included support, expertise, price, and reputation.

## Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the “Ask the Author” form.

**Q:** What resources are available for those interested in the ASP model?

**A:** A variety of industry groups have sprung up over the past year. Information can be found at these Web sites among others: [www.atlentis.com](http://www.atlentis.com), [www.aspisland.com](http://www.aspisland.com), [www.itaa.org](http://www.itaa.org), [www.aspnews.com](http://www.aspnews.com), [www.aspindustry.org](http://www.aspindustry.org), [www.aspwire.com](http://www.aspwire.com).

**Q:** What applications should we choose?

**A:** Choosing applications is one of the most critical events in your migration strategy. It is important to assess the current skill sets of your organization, your existing client base, and a number of ISVs. The best applications will leverage existing talents and knowledge. It is also important to offer applications that are attractive to your existing customer base. Finally, it is highly desirable to form a strong partnership with the ISV that produced the application.

**Q:** How will we market and sell our product?

**A:** This is one of the most urgent questions facing the ASP industry. Start by leveraging your existing sales organization. In general, services should be marketed to existing customers first, because they have an existing buying relationship with your company. Various sales models and reseller channel relationships are possible. Your organization will have to develop a customized

strategy specifically for your suite of offerings. As you successfully add additional services to your offerings, it is imperative that you transition your sales strategy to focus on the services instead of the connectivity.

**Q:** How can we determine if we are ready to add ASP services?

**A:** It is our belief that two main long-term strategies exist for ISPs: become the low-cost producer and compete on price, or add a variety of additional services to fend off commoditization. Historically, these have been the only potential strategies for companies in a market segment facing commoditization. Only a few competitors will be able to transition to become low-cost producers, and they will likely be among the largest players in the market. The chance of succeeding with additional services is far higher, and we believe ASP services to be an excellent solution.



## Server Level Considerations

### Solutions in this chapter:

- Implementation, Where to Begin
- Software Solutions for Your ASP
- Application Software Types
- Network Service Considerations
- Data Backups and How They Can Affect You
- Virus Scanning Suggestions
- Thin Client Solutions
- Maintenance and Support Issues
  
- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions



## Introduction

In the past, Internet service providers (ISPs) were not concerned with the OSI reference model beyond layer 4. Up until that time, their only job was to deliver data and connectivity to their clients. This meant that they needed to provide the wiring, switching, and routing infrastructure, and possibly advanced services in this range such as Quality of Service (QoS) when they wanted to move to the next level of service. ISPs were originally only concerned with the streamlining of their data and network services, so there was very little understanding of their customers' complex application systems and requirements. What was necessary to these ISPs was the creation of a successful business model when the Internet was first beginning to grow, since many companies wanted to have a presence on the Internet, and wanted to turn to their local ISP to provide this connectivity. Traditionally, these companies usually provided their own application services, and rarely looked at the Internet as a way of cutting costs and providing new features to their clients. In fact, many of these companies did not even understand why they needed an Internet connection, except that everyone else seemed to be getting one.

As more of these companies and individuals signed up for Internet services, there was a quantifiable need for more ISPs to supply the growing demand. As the boom began to stagnate, the ISPs looked for ways to differentiate themselves from their competitors in order to maintain and increase their customer base and raise their market share. This led to an all-out battle among ISPs that resulted in many of these startup ISP businesses closing their doors as their customers were lured away to fancier ISPs with new product offerings and promises. In the end, customers became far more finicky and looked for a wide range of service offerings as well as a bargain. This led to a new, unexplored realm of application offerings.

Today, there is no limit as to what types of service application service providers (ASPs) can offer. There are market-ready databases, e-mail, rich multimedia content, storage, and data backup services, just to name a few. Whatever the application is that is being offered, they all have one thing in common: they all require their servers to actually compute and alter data in order to provide the services. Servers provide the groundwork for the mapping to the upper layers of the OSI model, mainly layers 5 through 7. These servers will host the applications that provide the services that can be deployed to support and benefit the internal resources and external customer base.

ASPs need to support content, database information, storage area networking (SAN), and file servers to truly provide well-rounded application service offerings.

These “server farms” must be maintained and supported with the priority that the application demands. Remember that the design of a system and implementation of an infrastructure that will work for your company and satisfy your customers will require exceptionally careful planning and forethought.

This chapter is written to help you better understand how servers work in your environment and how the services that they offer can assist you in the creation of a successful ASP. It is designed to take you from the ground up, and discusses different server hardware and operating systems, as well as intricate considerations such as server fault tolerance, connectivity options, data backup, virus scanning, maintenance, redundancy, and system recovery.

## Implementation, Where to Begin

At first glance, it may seem remarkably easy to just throw some servers onto your existing network, install a software package, and begin selling a particular service and raking in the dough. If only it were that easy. Numerous considerations and a good deal of planning for the future are required to provide a reliable, cost-effective solution that is scalable and requires a minimal amount of maintenance to satisfy your customers’ requirements. All of that needs to be handled, as well as your own business needs, to keep the system robust and resource efficient.

At the heart of an ISP/ASP are the server base and the application software packages. If they do not function efficiently, the ASP will not run effectively. On the other hand, if the servers are under-utilized, your ASP will be inefficient and it will be difficult to achieve a decent return on investment (ROI). As with most network-based solutions, there will need to be a terrific balancing act in order to make a profitable and reliable business model that is both fine-tuned and predictable. With this said, let’s define what a server actually is, its components, the applications that run on such a system, and other reasons for concern.

## Server Hardware

I thought that I might go a little tangential here. The abacus is credited as the first computer. It is essentially a wooden rack that holds two or more horizontal wires with beads strung along them. With this device, it became possible to solve basic arithmetic problems by sliding the beads around according to a predefined set of rules. This made mathematical problem solving more efficient, and allowed for greater math developments. Many years later, in 1642, Blaise Pascal built what is considered the first ‘digital’ computer. This device was very different from the abacus, as it was capable of adding numbers using dials instead of beads. Many

years later, complex computers began to take form, and instead of using dials or beads, these new machines used vacuum tubes to compute data. This breakthrough quickly led to the solid-state electronic devices that we still see in use today, such as our computers, calculators, and any other electronic device.

The differences between all of these devices are vast, but the most distinct difference has to do with the hardware that each type of device incorporates. It is, after all, the hardware that allows each of these devices to perform its intended function. It is the hardware in a solid-state device that allows it to perform operations and alter data, just as it is the hardware of the abacus that allows for arithmetic computation. Over the years, the hardware has improved significantly, as has the computer's ability to alter data and offer new functionality that had never before been thought of.

The computers we use today share a common bond and are based on the same concepts. Although there are many different manufacturers, makes, and models of computers, they all share the same basic design. All computers in use today rely on processors, memory, and mass storage to perform their functions. Some of these systems might have more of one item than another, or have a design that is very different from another. However, even the most powerful mainframe is designed using the same basic electronic concepts as the lowliest computer manufactured today.

## Central Processing Unit

Many individual pieces comprise a server. At its core is a central processing unit (CPU), or microprocessor that serves as the brain of the computer. This device operates in binary mode, meaning that its components can be in only one of two different states at the same time: on or off. Binary provides the basis for electronic logic, which is capable of computing the instructions of a computer program. Since the microprocessor is involved in every aspect and function that a server provides, and serves to control the other circuitry and devices in the system, it should always be a consideration when purchasing server hardware.

Today, there are only two basic types of microprocessors available for computers:

- Complex Instruction Set Computers (CISC)
- Reduced Instruction Set Computers (RISC)

CISC microprocessors refer to a processor that has a full or complete set of instructions. Each of these instructions could be a single operation, or might perform a series of operations inside the microprocessor. Since one instruction might

perform multiple operations, there is a reduction in the total number of instructions required to implement a given program or sequence of code. This means that the programs or code designed to run on a CISC microprocessor are smaller and contain fewer instructions. The primary manufacturer of CISC microprocessors is Intel Corporation, which makes the Pentium family of processors. These microprocessors are found in almost all personal computers (PCs) and compatible systems. In the past, we typically thought that a PC could or should only be used for personal use. However, the power of these new processors and technology have made them a very viable and inexpensive server solution that is capable of rivaling most mainframes of yesteryear.

RISC microprocessors are a newer design that provides a smaller set of computer instructions. Since each instruction that a computer must perform requires additional components and circuitry, a smaller amount of possible computer instructions will make the microprocessor inherently more simplistic. This translates into an efficient microprocessor that can rival the speed of an equivalent microprocessor with a larger list of instructions. In fact, RISC microprocessors often achieve two to four times the performance of opposing CISC microprocessors made from the same semiconductor materials and using the same clock rates. The drawback to this is that a smaller set of microprocessor instructions requires more complex application code to help the microprocessor perform its function. What this means is that the software code built to run on a RISC microprocessor is typically larger and contains more instructions than the same code running on a CISC microprocessor.

Although there are only two types of microprocessors, there are many different manufacturers of them. The primary manufacturers of RISC microprocessors are Sun Microsystems, which makes the SPARC family of microprocessors; Digital Equipment Corporation (DEC), which makes the ALPHA family of microprocessors; and Motorola, which makes the microprocessors installed in Apple Macintosh systems. These microprocessors are all proprietary and designed for specific server manufacturers. It is impossible to install one manufacturer's microprocessor in a server for which it is not designed. For instance, a SPARC microprocessor cannot be installed in a DEC or Macintosh server; it will not fit due to form factors, nor is it able to function with the other components of the server. Furthermore, each microprocessor has a differing set of computer instructions, making it impossible to use a software program designed for one processor on a different RISC processor; instead, the software will first need to be recompiled for the particular processor. Because of this, you may need a particular server and microprocessor in order to run a specific application or software package. Moreover, if you decide to

purchase different servers in the future, it may not be possible to simply transfer your applications and data to the new systems.

On the other hand, CISC processors are not proprietary in nature, and instead conform to a standard. For instance, Intel's current microprocessor offering is built upon the processor they designed many years ago, and conforms to the standard known as x86. Because CISC processors have a complete set of instructions, a particular manufacturer's microprocessor can be used in other manufacturers' servers. Software tends to be less complicated since it is designed to run on a microprocessor platform that is highly standardized. Moreover, as new CISC microprocessors are developed and instructions are added, the new microprocessor still uses the old set of instructions and therefore provides backward software compatibility. It is more difficult to provide true hardware compatibility since there are several different physical standards, and different or newer processors tend to have different dimensions and connections to the server's circuitry. The major manufacturers of CISC microprocessors are Intel Corporation, which makes the Pentium and Celeron family of microprocessors, and Advanced Micro Devices (AMD), which manufactures the K6, Duron, and Athlon families of microprocessors.

## Symmetric Multiprocessing

In addition to purchasing a fast processor for your server, there is another technology called Symmetric Multiprocessing (SMP). SMP is an architecture that provides better performance by using multiple processors in the same server. Unlike uni-processor systems, where a single processor is installed in a server, symmetric multiprocessing allows multiple processors to work in parallel while using a single operating system along with common memory, storage, and other installed components. SMP relies on the server hardware, operating system, and applications for support, and not every server and operating system is designed to use SMP. However, it has been standardized, and the majority of hardware and software manufacturers have designed this capability into their products, and will provide technical support on its features.

Installing multiple processors in your server can help speed applications running on the system if your current processors are overloaded. However, adding additional processors does not always solve your problems. You may need to consider other factors before adding additional processors.

## Random Access Memory

The memory in a computer is known as random access memory (RAM). Much like a human's short-term memory, a computer uses RAM to temporarily store operating system and application data. RAM provides quick and easy access by a computer's microprocessor in order to use and alter data, and make the system perform a desired function. RAM is volatile, which means that when the system loses electrical power, either by accident or purposefully, all data stored in RAM is lost. This is not a big issue, since permanent data is stored using a different media, such as a hard disk, floppy disk, or CD-ROM.

RAM is very quick when compared to other storage media, as it can be timed in milli- and micro- seconds. The drawback, though, is that it is also far more expensive. It would cost too much money to have a system that used only RAM as its storage media, and it would not be reliable; if the system lost power, the server would lose all of its data and application software. It is important, however, to have a sufficient amount of RAM for your applications to perform efficiently. The exact amount required varies significantly depending on the operating system, software, application, and size of the data being accessed. If too little RAM is installed in a system, the excess data will run over onto the hard disks, which require physical access, and therefore operate at much slower speeds. This will cause a serious degradation in performance, and as we know, that is almost always unacceptable. On the other hand, if too much RAM is installed in a system, and is never or infrequently used, it may be a waste of money. Most software applications will list the system requirements, which include basic memory requirements. Try to keep in mind that these are generalizations, and the factors mentioned previously may change the true memory requirements of a particular software package. Moreover, if multiple software packages are run on the same system at the same time, the memory requirements are semi-additive, meaning that you must add all of the requirements together to arrive at the correct amount of RAM necessary to run the system smoothly and efficiently. As a general rule, it is always a good idea to exceed the memory requirements slightly, and to never allow a computer to begin using the hard disks as RAM.

There are numerous types of RAM in use, ranging in size and function, and physical dimensions. Some server manufacturers require proprietary RAM, while others use standardized versions. In addition, sometimes RAM will not function in a device that has other types or sizes of RAM installed. Whatever the case, always check with your vendor to ensure that a particular manufacturer's RAM will work in your make and model of server.

## Mass Storage

The most typical form of permanent mass-storage media is the hard drive. A hard drive is comprised of a set of metallic disks (hard disks) that have the capability of electro-magnetically storing data on their surface. The disks rotate at a predetermined speed, which is usually between the range of 4500 and 10,000 revolutions per minute (RPM). The device contains multiple heads that “float” above the spinning disk and perform the actual access of data on the disks. A single hard drive will typically provide many billions of bytes of data storage.

There are many different types of hard drives and manufacturers; however, they all operate in the same physical manner previously described. The main differences lie in how they are controlled and interface with the server or mainframe.

The main types of interfaces that attach mass-storage devices to computers and mainframes are:

- Enhanced Integrated Drive Electronics (EIDE)
- Small Computer System Interface (SCSI)
- Fibre Channel
- Enterprise Systems Connection (ESCON)
- Fiber Connectivity (FICON)

Enhanced Integrated Drive Electronics (EIDE) provides a standardized interface between the computer and hard drives. It is based on an older standard known as Integrated Drive Electronics (IDE), but has been “enhanced” to support drives that are in excess of 528 megabytes (MB). The enhancements also provide faster access to the hard drives, although it is still not as fast as SCSI drives. Most personal computers are equipped with EIDE controllers and hard drives because they are inexpensive when compared to SCSI drives. EIDE drives are rarely used or provided for in a server architecture, since they offer less reliability and currently have a maximum bus throughput of 66 Mbps.

Small Computer Systems Interface (SCSI) describes a set of interfaces that interconnect mass-storage devices and computers. SCSI provides speed and flexibility that is unrivaled by EIDE technology. In fact, newer SCSI implementations are capable of providing a throughput in excess of 80 Mbps. SCSI also allows for multiple SCSI-capable devices to be “chained” together while connected to a single SCSI port on a controller. Although there are many different levels of SCSI, the controllers are almost all backward compatible, but usually require the use of an adaptor since the actual connectors can vary. See Table 3.1 for a breakdown of each SCSI technology.

**Table 3.1** SCSI Technologies

SCSI Technology	Maximum Throughput (Mbps)	Maximum Number of Devices	Maximum Cable Length (Meters)
SCSI-1	5	8	6
SCSI-2	5 to 10	8 or 16	6
Fast SCSI-2	10 to 20	8	3
Wide SCSI-2	20	16	3
Fast Wide SCSI-2	20	16	3
Ultra SCSI-3 (8 bit)	20	8	1.5
Ultra SCSI-3 (16 bit)	40	16	1.5
Ultra-2 SCSI	40	8	12
Wide Ultra-2 SCSI	80	16	12
Ultra-3 SCSI	160	16	12

Fibre Channel has been introduced as a replacement for the SCSI architecture. Fibre Channel provides a method for transmitting data between computers at a rate of 100 Mbps, and scales up to 1 Gigabit per second (Gbps). It is used to connect hard drives to storage controllers, and is especially suited for complex data-sharing solutions, such as building storage area networks (SANs). Fibre Channel is especially flexible, and can use fiber-optic, coaxial, and twisted-pair cabling to interconnect almost any amount of storage devices. With the use of fiber-optic cables, switches, and routers, Fibre Channel offers the ability to attach devices that are up to six miles apart. Dell, EMC, IBM, Compaq, StorageTek, and Sun Microsystems are among the many manufacturers who already offer a compelling amount of Fibre Channel capable devices.

Enterprise System Connection (ESCON) describes a set of IBM products that interconnect mainframes and storage devices through the use of switches known as ESCON directors. They accomplish this by using fiber-optic cable, which allows the devices to be separated by more than 60 kilometers. ESCON also allows for a maximum throughput of over 200 Mbps. The main problem with ESCON is that it is a proprietary architecture, which does not allow for a great deal of flexibility between unlike systems.

Fiber Connectivity (FICON) offers a high-speed interface that is replacing ESCON in newer IBM mainframes. FICON also uses fiber-optic cabling, but introduces a new architecture that is approximately eight times faster than older



ESCON technology. FICON allows for multiplexing, which means that smaller data transfers can be combined with larger transfers to improve overall performance. Although this technology is also proprietary to IBM hardware, it does offer some compatibility with devices using the ESCON standard.

## Network Adapters

Network adapters, or network interface cards (NICs), provide a way for a computer to attach to a LAN and share data with other servers, workstations, and end users. Since there are several different types of networks and media, there are several different choices when it comes to network cards. The good news is that the only type of network cards you will need to concern yourself with are those that provide some sort of Ethernet connection. The other types of network cards, such as token ring, are generally outdated and should not be used in your LAN.

Ethernet was initially designed by Xerox Corporation in the 1970s and ran over coaxial cable. It is based on a technology called Carrier Sense Multiple Access Collision Detect (CSMA/CD), which helps it operate efficiently in networks with sporadic and occasionally heavy traffic patterns. Over the years, Ethernet technology has progressed rapidly, and through its simplicity, it has expanded to encompass the majority of local area networks in use today. Although Ethernet was initially designed to use coaxial cable, it has been adapted to also include shielded and unshielded twisted pair copper wiring and fiber-optic cable, thereby increasing its usability and flexibility even further.

There are many different types of Ethernet technologies, all of which have different distance limitations, run at different speeds, and use different media to transmit data. Table 3.2 displays some of the characteristics of different Ethernet technologies.

**Table 3.2** Ethernet Technologies and Characteristics

Technology	Speed (Mbps)	Physical Media	Maximum Cable Length (Meters)
10Base5	10	50-ohm thick coax	500
10Base2	10	50-ohm thin coax	185
10BaseFL	10	Fiber optic	2,000
10BaseT	10	Category 3 unshielded twisted pair	100

Continued

**Table 3.2** Continued

<b>Technology</b>	<b>Speed (Mbps)</b>	<b>Physical Media</b>	<b>Maximum Cable Length (Meters)</b>
100BaseTX	100	Category 5 unshielded twisted pair	100
100BaseFX	100	62.5 / 125 micron multimode fiber	400
100BaseT4	100	Category 3 unshielded twisted pair	100
1000BaseT	1000	Category 5 unshielded twisted pair	100
1000BaseLX	1000	5-micron single-mode fiber	> 3,000
1000BaseLX	1000	50-micron multimode fiber	> 550
1000BaseLX	1000	62.5-micron multimode	> 550
1000BaseSX	1000	50-micron multimode fiber	> 550
1000BaseSX	1000	62.5-micron multimode	250
1000BaseCX	1000	Balanced shielded twisted pair	25

100BaseT is the most commonplace form of Ethernet in use in ASP networks today. It operates at 100 Mbps, and usually runs over unshielded twisted pair. Because of its distance limitation of 100 meters, 100BaseFX is sometimes used when longer runs are required. 100 Mbps is faster than the average mass storage device, so it will allow you to deliver and receive data at very high speeds. This is good for systems that are part of a server farm, or cluster of servers. If you require faster speeds or longer distances, Gigabit Ethernet is the way to go.

Gigabit Ethernet provides speeds of 1000 Mbps. It is usually implemented to connect campus locations because it is capable of spanning very long distances without the use of a repeater. However, it is also an excellent technology to use when connecting groups of server farms or backbone equipment. Its speed provides a great deal of overhead for large servers that need to access an immense amount of data, or receive access from hundreds or thousands of users simultaneously. Gigabit Ethernet is quickly gaining acceptance as a means of connecting single devices to a LAN.

Gigabit Ethernet usually uses a laser to direct light over fiber-optic cable. In this method, the laser uses two wavelengths: long wave, or 1000BaseLX, and short wave, or 1000BaseSX. The two key differences between these two technologies

are cost and distance. Short-wave lasers are identical to the lasers found in your average CD player, and are therefore more commonly manufactured and available. This translates into an inexpensive laser design; however, due to its nature, it does not traverse extremely long distances as its counterpart does. Long-wave lasers are not as readily available and are more expensive, but are capable of traversing very long distances over single and multimode fiber.

*Single* and *multimode* refer to the way in which light is transmitted into the fiber-optic cable. With single-mode fiber, the light is transmitted straight into the core of the fiber-optic cable, and the light traverses a straight line through the core. Long-wave lasers are designed to run over single-mode fiber using a 9-micron core and 1300-nanometer laser. It has been optimized to take advantage of these properties, and works great when used for very long runs, up to 10 kilometers in length. Because of this, single-mode fiber has traditionally been used in campus environments, but it is rarely used to connect a single device such as a server to a LAN; it just wouldn't make sense.

Multimode fiber, on the other hand, is designed for long-wave as well as the less expensive short-wave lasers. In this design, light is shot into the fiber-optic cable at many different angles. Multimode fiber comes in two different sizes, 62.5 and 50-millimeter diameter. The 62.5-millimeter has a lower modal bandwidth, which means that short-wave lasers are able to traverse shorter distances than long-wave lasers. Servers are usually connected using 62.5-millimeter multimode fiber because they are typically located close to the terminating layer 2 device. In contrast, 50-millimeter fiber-optic cable has better modal bandwidth, and allows short-wave lasers to traverse longer distances. The 50-millimeter cable is generally used to traverse medium distances, such as connecting two buildings that are in close proximity to each other.

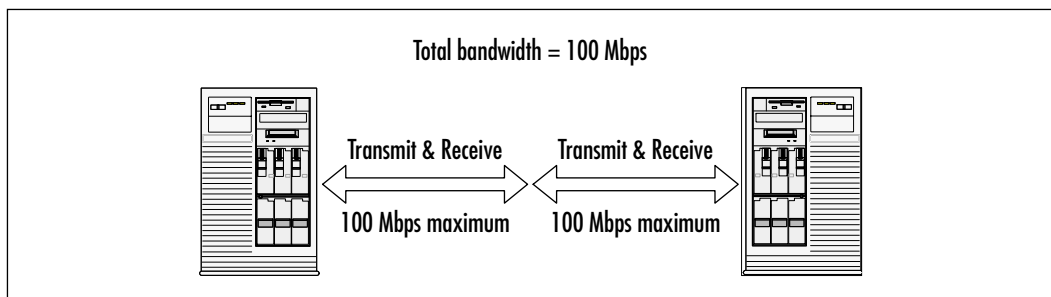
For cable runs of less than 25 meters, it is also possible to use a balanced 125-ohm cable known as 1000BaseCX. This technology does not use a laser, or fiber-optic strands, and instead uses electrical energy to pass information across the wire. Because of the serious distance limitation, it is typically used to make very short-haul data center connections, or to aggregate multiple racks together.

### *Half versus Full Duplex*

*Duplex* refers to a link's ability to send and receive data. When a device speaks over the network, it simply spits its information onto the wire. Since Ethernet is a shared media, this means that no two devices that are on the same network can speak at the same time; otherwise, it creates what is known as a *collision*, and the conversation is dropped. When this occurs, the devices need to retransmit their

data. This means that if too many collisions exist on the line, it is possible for all traffic to come to a standstill. Imagine a room filled with people all trying to speak and listen at the same time. If more than one person is talking, it is impossible to understand the conversations. If others are speaking while you are speaking, it is also impossible to understand what they are saying. What's worse is that it might be hundreds of people speaking at the same time. In order to understand anyone, you would have to do exactly what Ethernet does: ignore all the speakers, and wait for one person to repeat what he or she was saying without a disturbance from anyone else, including you. This means that the people speaking will need to wait their turn and listen before they speak to ensure that their words do not collide and cancel out someone else's. When this is done in an Ethernet environment, it is known as *half duplex*. As you can see, this could present serious ramifications, especially if one person just loves to hear his or her own voice and refuses to be silent. With Ethernet, there will be too many errors in the transmissions, and the retransmissions could stop all the conversations from progressing. For these reasons, "normal" Ethernet operates in half duplex, which hampers its overall throughput. Figure 3.1 shows two devices that are using a 100-Mbps half-duplex connection to speak.

**Figure 3.1** Two Devices in Half-Duplex Conversation

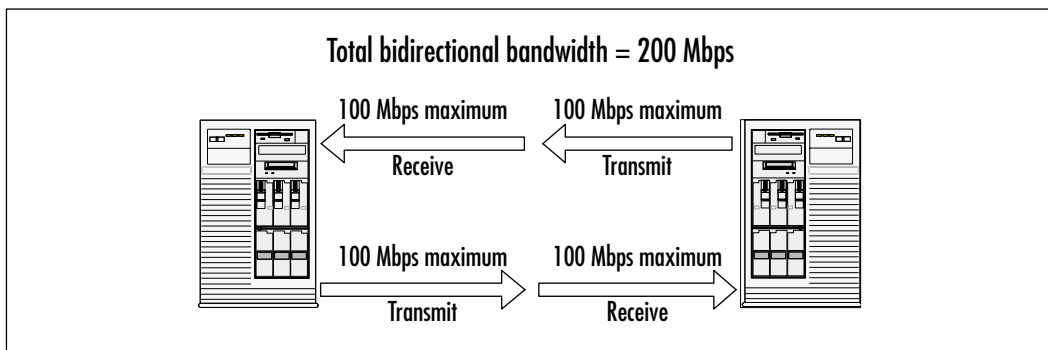


Now imagine that same room filled with people, except this time, everyone has broken into groups of two. Since each person is able to hold his or her own conversation with one other person, the two are generally not speaking over each other. For the most part, one person is listening while the other speaks, and vice versa. There might still be hundreds of conversations occurring at the same time, but nobody is attempting to listen to more than one conversation at once. This is akin to a full-duplex connection, except that network adapters use more than one wire, which means they actually have the ability to speak and listen at the same time as long as they are segregated from all the other conversations. The

advent of switches, and especially their reduction in price over the past 10 years, has made full duplex possible.

In essence, a switch “switches” what is spoken by one device and directs it to the destination listener, instead of allowing every device to hear the conversation. This reduces the amount of network congestion and collisions, making full-duplex speech a reality. Since this allows a single network adapter to send and receive data at the same time, it can, in theory, double the amount of available bandwidth. For instance, a 100-Mbps adapter may be able to transmit 200 Mbps assuming it can send data at 100 Mbps and receive data simultaneously at 100 Mbps. Figure 3.2 depicts two servers engaged in a full-duplex conversation.

**Figure 3.2** Two Devices in Full-Duplex Conversation



Obviously, whenever possible, it is better to use full duplex to transmit data. Most new switches and network adapters will support this operation. There is even a built-in method for negotiating speed and duplex, and will always attempt to make the fastest connection without errors. The problem with using these automated features is that once an error is received, they will usually drop to half duplex or a slower speed. The process of renegotiating the link can be slow, making the server or device inaccessible for that period of time. Moreover, it is not good to have interfaces constantly flapping between full and half duplex, since this alone can cause framing errors on the line.

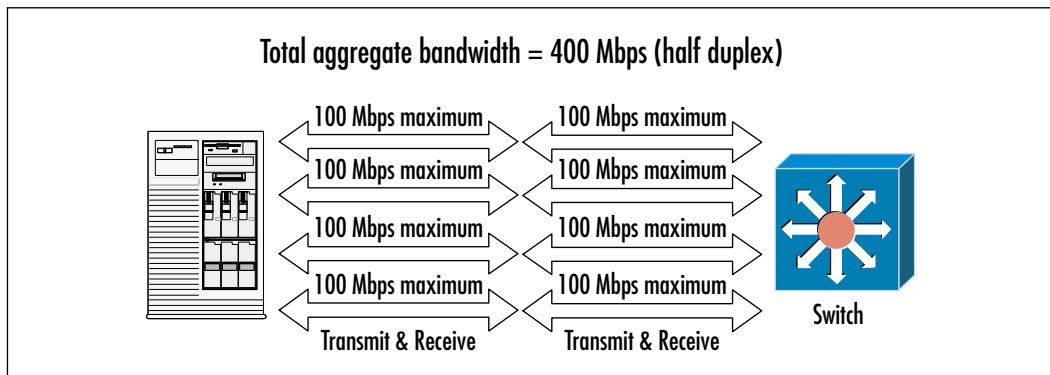
For these reasons, we recommend manually configuring both the speed and duplex on the servers and on the switches to which they attach. This will help prevent this flapping, and ensure full-duplex operation at all times. Most switches and network adapters have diagnostic utilities that can help inform you of transmission errors. If these errors are exceptionally high, it might prove helpful to fall back to half-duplex operation. However, the cause might be related to a different problem altogether, such as a faulty cable. In many instances, it may be helpful to

use a device to sniff the traffic on the network in order to get a better idea of the root of the problem.

### *Link Aggregation*

Link aggregation allows a single server to use two or more installed network interface cards (NICs) to aggregate bandwidth across several links; hence, the name. This is accomplished by assigning a single IP address to a group of network adapters. Since these adapters are all installed in the same server, it does not create an IP conflict. In fact, when a request is made to the server, only one device will answer the request—the server. The request and the reply could have been transmitted using any one or a combination of NICs in the aggregate group, but since the server and connecting layer 2 device is controlling the flow of data, there should not be any issues. It is similar to the way a load balancer “owns” a virtual IP address and load balances multiple servers, except in this case, the server has the IP address and balances its load across multiple Ethernet adapters. For an example, see Figure 3.3.


**Figure 3.3** Link Aggregations



A server might have four Fast Ethernet network cards installed. Instead of having only 100 Mbps of usable bandwidth, the cards are combined for a total aggregate bandwidth of 400 Mbps. If full duplex were used instead, the total usable bandwidth would come close to 800 Mbps. The network interfaces do not have to be Fast Ethernet; in fact, if Gigabit Ethernet were used instead, we would have usable bandwidth of 4 Gbps at half duplex, and 8 Gbps for full duplex. The benefit of this is that it allows you to add extra bandwidth to each server, and it can be done in increments on an as-needed basis. If you find yourself running out of bandwidth on a particular server, you can solve the problem simply by

adding an additional network card. In general, the amount of bandwidth and network throughput the server has is directly related to how many interface cards can be installed in a given server. Beyond providing additional bandwidth, this type of solution will also allow you to add additional fault tolerance to the server. Should a link in the aggregated grouping fail because of a non-operational NIC, a faulty cable, or a dead switch port, the load will automatically shift to the remaining NICs, thereby leaving the server operational and connected to the network. Once the link is fixed, it will be automatically added back into the aggregate grouping, allowing for full-speed operation to resume. Some vendors' implementations will even allow multiple links to be attached to multiple switches, improving fault tolerance and preventing a server from losing its network connection due to a total switch failure (see Figure 3.4).

**Figure 3.4** Link Aggregation Using Multiple Switches



## Configuring & Implementing...

### Link Aggregation Standardization

A link aggregation task force has been formed to launch the IEEE 802.3ad standard. This consortium will help vendors build a compatible form of link aggregation into all their network interface cards, which will allow link aggregation to be performed using different vendors' cards in the same server.

Although there are no current standards for this type of link aggregation, most vendors support proprietary technologies that will reliably perform link aggregation. The drawback to this method is that the same brand network card must be installed in the same server. The good news is that most of these technologies will interoperate with multivendor networking equipment such as switches, and there is usually little or no configuration required on the switch.

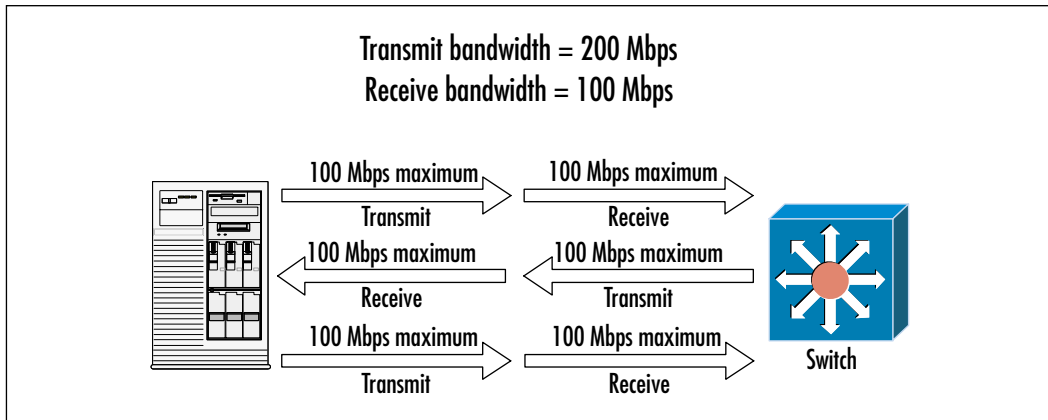
Since there are so many NIC manufacturers, be sure to research a particular card before you purchase a thousand NICs expecting a certain level of functionality. The following are some of the most popular forms of link aggregation from a few of the top manufacturers:

- Intel Adaptive Load Balancing (ALB)
- 3Com's Dynamic Access
- Cisco's Fast EtherChannel and Gigabit EtherChannel

Intel *Adaptive Load Balancing* allows up to four Intel Server network adapters to be installed in the same server. This allows for a total transmission aggregate bandwidth up to 400 Mbps with Fast Ethernet, and 4000 Mbps with Gigabit Ethernet adapters using full duplex. All outgoing data is evenly balanced with a proprietary algorithm, and each link serves as a transparent backup link, making this an excellent fault-tolerant solution.

Since the load balancing takes place only with outgoing packets, this technology assumes that the majority of data is coming *from* the server rather than to it. For incoming data, only a single adapter is used, so the effective downstream bandwidth is still only 100 or 1000 Mbps depending on the adapter. As an example see Figure 3.5.



**Figure 3.5** Intel Adaptive Load Balancing

In our example, we have a single server with two network adapters installed using half duplex. Two of these adapters will be able to transmit data allowing up to 200 Mbps upstream, while only one adapter can receive data. In a server farm environment, this is usually an acceptable solution, since the servers do not typically receive a lot of data. However, if you have a server that needs more bandwidth downstream, you may want to look into a different solution.

In order to use ALB, all the network cards must support ALB, and must be configured as part of the same load-balancing group.

3Com's *dynamic access* uses various proprietary algorithms to maximize throughput across a grouping of up to eight Fast Ethernet or Gigabit Ethernet network adapters. This technology is capable of transmitting and receiving on all network adapters that are configured as members of the aggregate group. Each network adapter maintains its own MAC addresses, but the group shares a single IP address. When transmitting from the server, all the data is evenly balanced across all of the adapters within the group based on the destination IP address and the interface on which the request was initiated. Conversely, when an Address Resolution Protocol (ARP) request is made to the server, the server uses one of the Media Access Control (MAC) addresses from one of the network adapters and answers the request in a round-robin fashion.

If the same client has multiple connections to the server, inbound traffic will rotate between all of the network adapters. This type of bidirectional communication allows for a total aggregate bandwidth of up to 1600 Mbps using Fast Ethernet, and a whopping 16000 Mbps for Gigabit Ethernet adapters assuming full-duplex operation.

Much like Intel's approach, this solution can be implemented across several switches to maximize your fault-tolerance capabilities. All the adapters must support *dynamic access* and must be configured as part of the same aggregate group to function properly.

EtherChannel was developed by Cisco Systems and allows for multiple link aggregation anywhere in the network. EtherChannel can be leveraged to alleviate bottlenecks between routers, switches, and servers. Depending on the hardware used, it is possible to combine up to eight links into a single transmission path supplying up to 1600 Mbps for Fast Ethernet, and 16000 Mbps for Gigabit Ethernet adapters.

The EtherChannel transmission path is bidirectional, and works in full-duplex operation. The links can be either Fast Ethernet or Gigabit Ethernet but cannot be mixed and matched, meaning that all network adapters within a group need to operate at the same speed. It is not possible to split a grouping across multiple switches, but some switches do allow an EtherChannel group to be spread across different modules within the same switch.

Traffic is balanced between the network adapters with an algorithm that remembers and uses each adapter's MAC address. If a port within the EtherChannel group fails, the remaining ports remain in operation, and the port aggregation database is updated within one second. Although it has not been adopted by all, numerous vendors have incorporated this technology in their design, including Adaptec, Compaq, Hewlett-Packard, Intel, Network Appliance, Silicon Graphics, and Sun Microsystems. Of course, to use EtherChannel, all of the network adapters must support Cisco's implementation, and they must be configured correctly.

Many switches are set up to automatically detect and activate EtherChannel connections, and begin functioning without any configuration changes required on the switch. Even if the switch does not come with this feature turned on by default, it is generally a very simple task to perform. Figure 3.6 shows how to set up EtherChannel on a Catalyst series switch using blade 3 and ports 1 through 4. This figure also shows how to verify the configuration.

### Figure 3.6 Configuring and Verifying EtherChannel

---

```
Console> (enable) set port channel 3/1-4 on
Port(s) 3/1-4 are assigned to admin group 52.
Port(s) 3/1-4 channel mode set to on.
Console> (enable) show port channel
```

---

Continued

**Figure 3.6 Continued**


---

```

Port Status   Channel      Admin Ch
      Mode           Group Id
-----
3/1 connected on           52  835
3/2 connected on           52  835
3/3 connected on           52  835
3/4 connected on           52  835
-----

Port Device-ID      Port-ID      Platform
-----
3/1 090031611(6500) 3/1          WS-C6500
3/2 090031611(6500) 3/2          WS-C6500
3/3 090031611(6500) 3/3          WS-C6500
3/4 090031611(6500) 3/4          WS-C6500
-----

Console> (enable) exit

```

---

In Figure 3.6, the Admin group was automatically assigned. This provides an identifier for the particular EtherChannel group and can be used to configure other features within the group such as the Spanning Tree Protocol (STP). If additional groups were configured on the switch, they would need to be assigned a different administrative group. Figure 3.7 shows how to manually configure an Admin group for a second EtherChannel grouping.

**Figure 3.7 Manually Configuring and Verifying an Admin Group**


---

```

Console> (enable) set port channel 3/5-6 99

Port(s) 3/5-6 are assigned to admin group 99.

Console> (enable) show channel group 99

```

---

**Continued**

**Figure 3.7 Continued**

Admin	Port	Status	Channel	Channel
group			Mode	
99	3/5	connected	auto silent	0
99	3/6	connected	auto silent	0

Admin	Port	Device-ID	Port-ID	Platform
group				
99	3/5			
99	3/6			

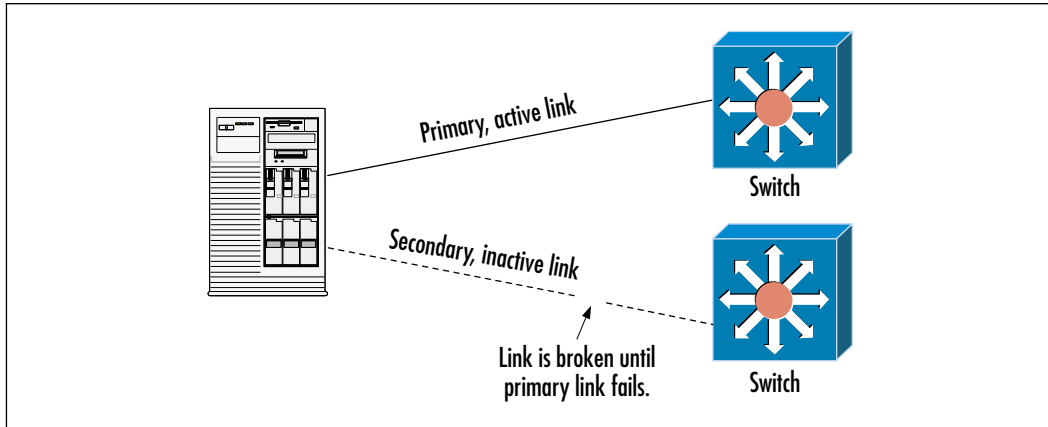
Console> (enable) exit

### *Adapter Fault Tolerance*

Building fault tolerance into a mission-critical server is a must. Although this was discussed in our link aggregation section, it is not always necessary to aggregate links in order to provide a backup or fault-tolerant link to the network. It is possible to use Adapter Fault Tolerance (AFT) to perform this task.

Instead of aggregating bandwidth between multiple links, most network adapters will include drivers that allow a second installed network adapter to be configured as a secondary, or redundant, link, especially those based on the Intel specification. In this type of configuration, the secondary link remains idle until the primary link becomes faulty. When this occurs, the secondary link is automatically configured, and becomes the primary adapter maintaining the server's network connectivity. The new network adapter uses the same IP address as the previous, allowing access to resume once the link is back online. In addition, if the second link is plugged into a separate switch, it could provide a fault-tolerant solution that allows the server to remain connected even if the primary switch were to fail altogether. Figure 3.8 shows an example of a server connected to two switches using adapter fault tolerance.

With this solution, it is possible for a connecting client to receive timeout problems when a link fails. This is due to the time it takes the second link to become operational. This is usually only a matter of seconds, but could increase significantly if STP is enabled on the associated switch port. STP is usually

**Figure 3.8** Server Connected Using Adapter Fault Tolerance

enabled by default on all switch ports to ensure that there are not any loops in the network. To accomplish this, STP does not allow a switch port to transmit data until it has been verified, and does not create a loop in the network. This can cause over a 60-second connection delay depending on the switch and configuration. To avoid this connection delay, it might be a good idea to configure PortFast on the switch ports that connect to the server. This will allow the switch port to transmit data almost immediately upon connection. When configuring this, be sure as to which port you are configuring, and never enable this on a port that connects to another switch. Figure 3.9 shows how to configure PortFast on a Cisco Catalyst series router that is in enable mode.

**Figure 3.9** Configuring Spanning Tree PortFast

```
switch> (enable) set spantree portfast 3/1 enable
```

Warning: Spantree port fast start should only be enabled on ports connected to a single host. Connecting hubs, concentrators, switches, bridges, etc. to a fast start port can cause temporary spanning tree loops. Use with caution.

```
Spantree ports 3/1 fast start enabled.
```

```
switch> (enable)
```

---

Continued

## Figure 3.9 Continued

```
switch> (enable) set spantree portfast 4/1-24 enable
```

Warning: Spantree port fast start should only be enabled on ports connected to a single host. Connecting hubs, concentrators, switches, bridges, etc. to a fast start port can cause temporary spanning tree loops. Use with caution.

```
Spantree ports 4/1-24 fast start enabled.
```

```
switch> (enable) quit
```

In Figure 3.9, the number “3” in the statement *set spantree portfast 3/1 enable* designates the third module or blade installed in the switch. The number “1” represents the actual port number. It is possible to give an entire range of ports in a single command, such as with the command: *set spantree portfast 4/1-24*. This command configures the switch and enables spanning tree PortFast on ports 1 through 24 in the fourth module.

### Configuring & Implementing...

#### Spanning Tree Issues


With some switches, it may be necessary to disable spanning tree on the ports that connect to the server. When doing so, be careful to not disable spanning tree entirely on the switch or ports that connect to other switches. If this is done, and loops do exist in the network, they could severely effect network traffic and possibly cause data delivery problems or a temporary network outage.

### *Dual or Multihomed Servers*

I do not want you to be confused about link aggregation or adapter fault tolerance when we speak about dual-homing, or multihoming a server. *Multihomed* does not simply mean to connect two or more network adapters to different switches. Instead, when we speak of dual or multihoming, we are referring to a

server that is connected to two or more separate LANs. This is usually done to help segregate traffic, and increase the usable bandwidth on a particular network. It is also possible to enforce multiple sets of access lists on each network, which can help to increase security. See Figure 3.10 for an example of dual-homed servers.

**Figure 3.10** Dual-Homed Servers



In this example, there are two servers connected to Network #1. Network #2 is identical, except that it contains a third server with only a single connection. The dual-homed servers could be Web servers that receive requests from the Internet, while the third, single-homed server might be an application server that only communicates with the two Web servers. In this instance, all traffic from the Internet will only traverse Network #1. However, when one of the Web servers receives a request from an Internet host, it can retrieve the necessary data from the application server through a separate network connection, Network #2, and not congest the first network, Network #1, with internal traffic. Additionally, the two Web servers might need to exchange information with each other from time to time or with other “internal” servers behind a router. These servers could be configured to use Network #2 for this communication so as to not congest Network #1. This will help keep servers in Network #1 available for requests from the Internet hosts.

It is possible to multihome servers using aggregated links to ensure available bandwidth for data-hungry servers. It is also possible to use adapter fault tolerance

in conjunction with multihoming techniques to ensure a high level of availability with your server. In fact, when these features are combined, it can make for a highly available, scalable, fault-tolerant solution that will keep you sleeping at night, with confidence that your servers will remain connected to the network with more than enough bandwidth available.

## Software Solutions for Your ASP

*Software* is a term for the various types of programs that can be installed on a computer system. There are many different types of software available for all types of computers and mainframes, but they can all be broken down into three distinct variants:

- System
- Application
- Middleware

### System Software

*System software* describes software packages that provide the basis for all other applications that are run on a computer. System software directs the different components that comprise a computer. Without system software, other applications would be unable to accomplish any task, and there would not be a platform on which to run them. System software includes specific device drivers that help the components contained in the system interface with one another.

The most common form of system software is the operating system, or O/S. The operating system provides the basis for all other software packages, and choreographs the interaction of device drivers and components. There are several different types of operating systems on the market today, but because of your business model, you should only concern yourself with those that offer network services. These are known as network operating systems (NOS).

### *Unix*

*Unix* is a server network operating system developed by Ken Thompson and Dennis Ritchie at Bell Laboratories in 1969. It was originally designed as an interactive system that allowed users to share CPU time on a single system. In 1974, Unix became the first operating system to be written in the programming language known as C.



Unix was not a proprietary operating system, and the source code has been available to the public since its inception. This has allowed many companies and individuals the opportunity to expand upon Unix and reengineer it many times over. This has probably been due in part to the fact that its user interface, known as “shell,” has been standardized by the IEEE as the Portable Operating System Interface (POSIX). Since there are many versions of Unix-type environments, the term now describes a large group of operating systems that may seem different from each other in operation and functionality. Due to its versatility and reliability, Unix has become one of the most popular operating systems in use today, and has helped shape the Internet and other computer networks to the way we know them.

As stated earlier, there are many different versions of Unix. Some of these versions have evolved into proprietary operating systems that only run on particular types of servers, while others remain in the public domain and offer support for many computer architectures. All of these operating systems have advantages and disadvantages, so it is important to understand their design and function and how they relate to your application.

Currently, the leading Unix environment is Solaris from Sun Microsystems. This is a proprietary operating system, and is licensed through Sun without the source code. What this means is that it cannot be adapted or recompiled by any user or company other than Sun Microsystems. It has been designed around IP networking for more than 15 years, and offers an extremely stable and scalable operating system. Typically, it is not considered as efficient or as “snappy” as other versions of Unix, but it offers excellent support for most enterprise-type applications. The servers and software licensing can be a little on the pricey side, but Sun offers excellent technical support for both their hardware and software. Recently, Solaris has adapted a version of Solaris that will operate on an Intel x86 platform. However, this is not recommended, since it does not offer the same level of reliability as with Sun’s own hardware, and other operating systems can outperform it when running on the same platform.

There are numerous other proprietary Unix operating systems, each of which is designed for a particular server or mainframe. Much like Sun Microsystems’ design, Hewlett-Packard (HP) also offers a form of proprietary Unix called HP-UX for their servers. HP-UX runs on high-performance Hewlett-Packard Unix servers, and offers a good amount of scalability and reliability throughout. IBM also makes a different proprietary Unix O/S for a wide range of their servers known as AIX. Their operating systems tend to be designed around a particular model of server, so there is very little flexibility built into the code. Quite a few

applications are designed to run on all of these different operating systems, but they are usually proprietary in nature. Although these systems may not offer the same level of versatility as other flavors of Unix, they are capable of delivering a serious amount of horsepower with an extreme amount of reliability built into them. However, before rushing out with your million-dollar check, it might be a good idea to reexamine the application support for a given operating system.

Berkeley Software Distribution (BSD) is another Unix-based network operating system that was developed and distributed by the University of California at Berkeley. BSD is a popular and reliable operating system that has prompted many other Unix distributors to base their features and functionality on it. There are commercially available releases of BSD, such as BSDi that delivers an Internet-based server solution that is used by many ISPs and ASPs and has been adapted by some hardware manufacturers to run on their Internet-based appliances. There are also free versions of BSD available such as freeBSD and openBSD that provide a reliable network operating system that can rival other commercially distributed operating systems. The main advantage to BSD is its high level of security. Directly out of the box, it is generally more secure than other operating systems. It's designed for and targeted at the Internet market, and is definitely worth looking into.

Linux is a network operating system that has been gaining a lot of ground, especially in the last five years. It was initially designed to provide personal computer users with a free or very low cost operating system that was comparable to more expensive Unix operating systems.

Linux was named after its inventor Linus Torvald who wrote the first Linux kernel, which is the central core of the operating system. The operating system makes use of software components that were written for the Free Software Foundation, or GNU. This allows the kernel and associated software and modules to be copied, modified, and redistributed at no cost. Linux is considered to be in the public domain, equipping anyone and everyone with the ability to alter the operating system to his or her liking. This has led to many different flavors of Linux, and many companies that distribute variations on the operating system.

Linux is a complete operating system, and offers a great amount of multi-vendor support. In fact, device drivers can sometimes be found for devices that have yet to be officially released. There are versions of Linux for the x86 platform, PowerPC, SPARC, Alpha, and even the IBM S/390. Linux also conforms to the POSIX standard, which allows software to be easily ported to other operating systems. Linux is based on the IP protocol, and offers many useful tools for the network environment. This can sometimes be a bad thing, though, since out-of-the-box security is sometimes questionable.

The real problem with Linux is the lack of support. Although many companies distribute their own flavors, the modules and applications included are in the public domain, which makes it nearly impossible to find anyone who is compelled to provide technical support or big fixes. Some companies, however, do provide comprehensive support for an additional fee.

### *Microsoft Windows*

Microsoft was first founded in 1975, and quickly came into the mainstream early in the 1980s. One of its first operating systems was the Disk Operating System (DOS), which was created as an interface for the IBM personal computer. This operating system quickly progressed to incorporate a graphical user interface (GUI), which led to the Windows revolution. Their first release of Windows NT 3.51 brought Microsoft heavily into the network server operating system market. Today, Microsoft makes several network server operating systems that are in widespread use.

Most applications that run under these operating systems are typically graphical in nature, and the use of a keyboard to select items or perform system functions has diminished. With Microsoft's innovations and its widespread acceptance, it has quickly become one of the top operating systems on the market today.

Windows 2000 Server is the latest entry-level version of Windows. It expands upon Microsoft's previous NT technology to create a powerful solution for next-generation business computing. Win2k incorporates Web and file services, as well as a good amount of security and performance. It has been designed to scale from one server with a few clients, all the way to hundreds of servers with thousands of clients. Although all of the Windows 2000 interfaces may seem the same, this base version has been designed with the small business in mind, and does not include all the bells and whistles that other, more expensive versions offer.

Windows 2000 Advanced Server offers all of the features available in the standard version, but includes more reliability and scalability, as well as additional features for applications that require a higher level of scalability. Advanced Server is typically used for essential business and e-commerce applications that handle heavier workloads. It is commonly considered the popular favorite of many ASPs because of its ease of use and price tag. However, if you require more features, there is another version of this operating system available.

Windows 2000 DataCenter Server offers the highest level of availability and scalability that Microsoft has ever presented. It was designed for enterprise and mission-critical applications that require an extremely high amount of server reliability. With this release, Microsoft has teamed up with other industry leaders to

provide an alternative to other proprietary software packages that usually run on powerful Unix platforms. Combining this operating system with the extensive support infrastructure that Microsoft has built provides a sensible solution that can be deployed in a large data-center environment.

### *Novell Netware*

Novell offers a powerful network operating system called NetWare. This operating system was originally designed for use in small to enterprise businesses and networks, and typically used a protocol stack called Internet Packet eXchange (IPX). With the growth of the Internet, and IP networks, Novell has started to design its operating systems around the IP protocol. In fact, its latest operating system, Novell 5.1, is capable of running a plethora of IP-based applications, and comes standard with an enterprise Web and Web-based search server, Network News Transfer Protocol (NNTP) server, FTP server, and a multimedia server capable of delivering audio and video feeds across an IP network. The latest version of NetWare is designed to operate on a PC with an Intel Pentium class or higher processor. The support offered by Novell is excellent, and the operating system can be easily integrated into existing Windows and Unix environments.

## Application Software Types

*Applications* is the term used to describe a group of programs or code designed to perform a specific function directly for users or other application packages. Some examples of application software are word processors, databases, Web browsers, development tools, image-editing programs, and communication programs such as e-mail. Although these programs provide different functionality, they are all designed to use the services of the underlying operating system to provide a high level of user interaction with data. Without application programs, a computer or server would be a useless heap of technology.

## Web Applications

As you are probably already aware, a Web server is an application that uses the client/server model along with the HyperText Transfer Protocol (HTTP) to serve data to client systems and their users across a network. This data forms what we commonly refer to as a Web page, and is the most popular function that the Internet provides. This type of server allows companies to make an individualized presence on the Internet, while serving data, product offerings, and valuable information to their customers across the Internet.

Web servers are not always publicly accessible and are sometimes reserved for internal business use only. Most ASPs use Web servers to deliver secure content to their customers, directly over the Internet, through virtual private networks (VPNs) and also by way of private network connections. Most Web servers offer a method of authenticating users in order to dictate access privileges. Regardless of the usage, the Web server offers a popular method of delivering content to users.

In addition to serving Web pages, many Web server packages bundle numerous network-related products that are capable of serving e-mail, delivering files using the File Transfer Protocol (FTP), and providing other network-centric content delivery techniques. They also include other application development services that are required to build integrated, component-based applications for the Internet, extranet, or private intranets.

Many Web software packages include an application server in addition to a Web server. The application server bridges the gap between applications, and the Web server delivers content to users. It could be considered “middleware,” since the application server itself does not usually deliver content directly to the users. Instead, it is the application server’s job to gather static or dynamic data from applications such as a database, package this data into a usable format, and pass it along to the Web server for delivery to client systems.

The application server is at the core of an ASP, and is therefore a very important aspect to consider. The application server provides an ASP with the ability to deliver nearly any type of application to your end users, and allows for advanced functionality that will help set your company apart from other ASPs.

Web servers and application servers are very broad topics, and it would be possible to write a book just on these alone. However, because they are so important to the ASP model, we will explain some of the most popular programs available today, and describe their purposes and functionality. To get a better understanding of how these applications interoperate, and especially how to configure them, it would be wise to purchase a book on the individual software package you are thinking about using.

With that said, the following are some of the most popular brands of Web-based application suites:

- Microsoft Internet Information Server (IIS) and components
- Apache HTTP Server and additional modules
- Netscape Fast Track Server
- Allaire family of products

- Novell Web & Application Services
- Domino family of Server Software

Internet Information Server (IIS) is a scalable Web server offering from Microsoft Corporation that runs under the Windows family of operating systems. IIS includes Web development tools, e-mail, and FTP services. It also offers the ability to use multiple IP addresses on the same server, allowing it to host thousands of Web sites on a single server. When used in conjunction with Windows 2000's built-in Terminal Services, it is possible to remotely manage the Web server from any authenticated machine running a Terminal Services client.

Active Server Pages are typically referred to as ASPs, but we will not use this term so as not to confuse it with application service provider. Active Server Pages is a programming environment that provides the ability to combine HTML code with scripting and other components to create Internet-based applications that can be accessed through a Web server. In effect, Active Server Pages provide the glue for combining and offering applications over the Internet.

Basically, when users access an Active Server Page on your Web site, their browser requests the Active Server Page directly from your Web server. The Web server uses server-side scripting to process the requested file sequentially, and executes any script commands that may be contained in the file. The result is a dynamic Web page that is sent to the initial client requesting the information. It is not a very complicated idea, but it can provide some very complex solutions to content delivery.

Since the script runs on your Web servers, your servers do all the processing and deliver the content to your customers in plain HTML format. One benefit this offers is that your Web pages are limited only by what your Web server supports. Another added benefit is that the user will be unable to use the "view source" function on his or her Web browser to steal content and features. If this is attempted, the user will only see the generated HTML document, and will be unable to see the actual scripts that were run to create the content.

Apache HTTP Server is an open-source software package that is organized by the Apache Software Foundation. This is a membership-based, not-for-profit corporation that ensures that Apache software projects exist beyond the participation of individual volunteers. Because Apache offers its source code to the public, it is easy to find versions that will run on most Unix-based platforms. Some versions of the software have also been ported to run under Microsoft Windows operating systems, but due to the lack of support and interest, it is rarely seen on these systems. Apache comes as a stand-alone Web server that can be compiled

with different modules, allowing for efficient software that can scale according to the needs and requirements of a company. With these compiled and other plug-in modules, it is possible for Apache to support almost any Web language, including Java, SGML, XML, CGI binaries, PERL, and Tool Command Language (Tcl) scripting. Apache offers a wealth of versatility and product support, and since the software runs on a Unix platform, there are numerous tools that will allow for remote administration.

FastTrack Server from Netscape is an entry-level Web software package that provides a suite of applications that allow for Web development and publishing. It supports HTTP 1.1, Java, and SSL3, as well as support for some Web and Open Database Connectivity (ODBC) enabled database applications. It is simple to install and easy to use. FastTrack can be remotely administrated using a Web-based interface that includes Lightweight Directory Access Protocol (LDAP) support for user authentication. For additional features, Netscape offers an upgraded version called Netscape Enterprise Server that includes additional content management and additional application support. These servers are rarely seen in enterprise environments, possibly due to a concern over the reliability and scalability of the products.

The Domino Advanced Enterprise Server integrates a messaging and Web application software platform that supports many enterprise operating systems, including Microsoft Windows 2000, AIX 4.3.1, HP-UX, Solaris, Linux, OS/2 Warp, OS/400, and OS/390. The software package includes the ability to cluster. This means that if a server or application stops functioning, users are immediately redirected to another server in the cluster. Under this scheme, every transaction is immediately replicated across all of the servers in a cluster, and then it is cached so that it can be replicated to the faulty server when it is functional again. In addition, server clustering allows for dynamic load balancing between the servers. More than nine servers could form a single cluster, and multiple clusters can span different LAN segments to create one single large Domino domain. This software package also allows the ability to partition the software, thereby allowing multiple instances of the program to be run on the same server.

Allaire makes several products designed to deliver application services over the Internet, the most popular of which is known as ColdFusion. This software package is an intuitive visual programming tool that allows for the delivery of high-performance Web applications that incorporate some of the most advanced features available. ColdFusion will work with almost any Web server package, and will run under Windows, Linux, HP-UX, and Solaris operating systems. Allaire offers many other visual development packages from which to choose, including

Spectra, which allows for the creation of customer-centric Internet business applications uniting content and commerce, and Jrun, which allows you to create advanced Java applications that can be accessed through a Web server.

## Database Applications

A database can be defined as a collection of data that is organized for management and access. There are two main types of databases, object oriented and relational. Object-oriented databases are typically used for programming languages such as Java and C++, and therefore do not require your attention unless you are planning to develop your own proprietary applications or middleware.

The second and most prevalent database is the relational database. In this type of database, data is stored and organized in tables that can be reorganized and accessed in a variety of ways. This can be accomplished without the need to reorganize the actual tables. The tables contain one or more data categories in columns. Data is stored in these predefined categories, and each row contains a unique instance of data for the categories defined by the columns. It might not sound like it, but a relational database is typically simple to create and access.

Imagine an order entry database. The database might include a table that describes a single customer. The columns of this table would contain the customer's name, address, telephone number, and other necessary information. There might be another table that stores a particular order, and its columns might contain the product type, customer, and date ordered. Although these two entries contain different information, together they comprise a very useful relational database. Using this database, it would be simple to have one user query the database to tabulate the amount of orders for a single day, or month, while another user might make a query for customers who reside within a certain area code, and yet another user might use the database to generate a report on accounts that have outstanding balances. In all of these situations, the same database and data is used for different purposes, yet the data was "related" to in different ways. Other than their simplicity, other features make relational databases very functional and desirable.

It is very easy to "grow" a relational database, since new tables can be appended to an already existing database. It is also simple to add a new data category without having to modify your existing applications.

A database can be thought of as a large filing cabinet, and in general, databases contain aggregations of data records. These aggregations could be sales transactions, product catalogs, inventories, customer profiles, accounting information, or any other type of data that needs to be stored and queried efficiently.



The most common interface for a relational database is the Structured Query Language (SQL), which is used to perform interactive queries and for gathering data for reports. SQL is both an International Organization for Standardization (ISO) and American National Standards Institute (ANSI) standard, and most database products will support SQL. Queries can be a collection of SQL commands that allow one to select, insert, remove, update, and locate data. In fact, SQL is considered a programming language in its own right, since commands are strung together to form complex data manipulation.

## Middleware Software

*Middleware* can be considered the “glue” that holds applications together. It is a general term for any computer application whose purpose is to combine or mediate between two applications in order to allow them to share data between them. For instance, an application might need middleware to connect to a database application, or a particular application might need a middleware application in order to communicate effectively with the operating system. Whatever the case may be, middleware serves to alleviate communication problems that exist between different applications running on the system.

There is no need to spell out the product offerings from middleware providers, except to say that the offerings can be proprietary, and have sometimes been written with a single purpose in mind. Many times, middleware may have been designed to provide a particular function. This could be a middleware application that allows a company to integrate its legacy systems and applications with more modern systems, or a proprietary application that needs connectivity to a standard database application such as Oracle.

Other types of middleware provide connectivity for a particular function through the use of many different software applications. An example of this would be messaging services. For instance, it is possible to use different mail clients to access your e-mail because of the middleware components that allows access to this resource from other applications that conform to a specific standard or rule set. Whatever the case, middleware can be designed allow two applications to interface with each other, regardless of their design and functionality.

In the ASP space, middleware can be a valuable component that allows your applications to interact with your customers’ proprietary systems and applications. Depending on your particular market, this type of software could be of enormous assistance to your success. If you will be writing numerous types of middleware applications for your system, it may be a good idea to check the references of the person or companies who are doing the design. The end product will more than

likely be a proprietary application, and if you require support or find a need to upgrade the application, you will probably need to contact the same person who initially wrote the application, or else find someone who can rewrite the code for you.

## Server Redundancy

*Server redundancy* is another must for any ASP. One of the last things you want is a server that causes a failure in your network and a loss of functionality to your customers. When designing your network, always try to avoid any single points of failure throughout the entire network. Especially watch for a single point of failure that might have the capability of bringing the entire network down.

Unfortunately, some ASPs who build massive redundancy into all of their network devices sometimes forget about some of their servers, on which all or many of their network and application services depend. When, because Murphy's Law likes "when" rather than "if," a server fails, it could cause many if not all of the applications and functionality in the network to grind to a halt until the server can be fixed or a temporary work-around can be found. This is a very serious situation, and I am sure that your customers will not love you for it or even understand the failure. Instead, many will probably end up knocking on your door expecting a refund of some kind. For these reasons, it is important to consider a solution that allows for the maximum amount of redundancy, with the minimum amount of cost and maintenance. Probably the best alternative is server clustering.

Server clustering is an important fault-tolerance solution that should be considered by any ASP, regardless of its exact service offering. The fact is that an ASP offers applications to its clients. These applications run on servers, and the servers should therefore always be accessible by your clients. If these applications are not available because of a server or component failure, then you are bound to have many unhappy customers who may want to cancel their service with you.

One way to ensure a high level of server availability and reliability is to use server clusters. A cluster is essentially a group of independent computers that work together as a single system to ensure that mission-critical applications and resources are always available, even if one or more servers within the cluster were to fail. The entire groups of servers are managed as if they were one system. They usually share a common namespace, and are specifically designed to tolerate component or whole server failures while still providing high server availability. Each server within the cluster will need to run the same software applications, and have access to the most recent version of the data. There are two different cluster models in use today:

- Shared device
- Shared nothing

These two models dictate how a given server will access resources on the network. It is important for each server within a cluster to have access to the most recent version of the data; otherwise, transactions could be lost forever.

## Shared Device

The *shared-device* model allows applications running within a cluster to access any hardware resource connected to any node within a cluster. For instance, there might be shared storage devices within a cluster that allow all the servers within the cluster to access the same exact data. As a result, access to the shared data and device must be synchronized in order to avoid devices colliding with each other.

Imagine for a moment, two devices both trying to alter the same data; to ensure that the data is altered in the correct order, there is a need to control access to the shared device. In many implementations of shared device clusters is a special component called a Distributed Lock Manager (DLM). A DLM is a service that manages access to devices within a cluster, and resolves any access conflicts that might arise as a result of multiple accesses. Although the DLM will ensure that the data is always synchronized, it can also cause a severe amount of overhead to the cluster. Since the DLM needs to keep in constant communication, and polices all resource activity, the process can generate additional traffic between nodes. In addition to this, there can be a serious performance hit since there is not serialized access to resources within the cluster.

There are several packages that allow for shared-device clustering, such as:

- Digital's Cluster (Windows NT/2000)
- Beowulf Cluster (Linux)
- CustomSystems' Clusters (Tru64 Unix, OpenVMS, and Linux)
- NCR's LifeKeeper (Windows NT, Linux, and Solaris)
- Veritas' Cluster Server (Windows NT/2000, Solaris, and HP/UX)

## Shared Nothing

The *shared-nothing* model does not use DLM, and therefore provides better throughput and less overhead when compared to the shared-device model. When using the share-nothing model, only one node can own and access a single hardware resource at any given time. If a failure occurs, a different server within the

node can take ownership of the failed nodes services and resume normal operation, keeping resources available to the users.

This solution will still allow multiple servers to share resources on a network, so it is still possible to use shared storage devices and the like; however, only a single server will have access to the “live” data. This might seem wasteful, since servers might be sitting idle while they wait for a server to fail. In reality, it is possible to set up multiple cluster groups on the same servers. This would allow each server to be a primary for a particular function, while all the other servers are secondary for each other. For instance, one server might be primary for database 1 and secondary for databases 2 and 3, while the second server is primary for database 2 and secondary for databases 1 and 3. The list could go on to include numerous servers and processes, each having sole access to the resources it “owns” while waiting to provide services for any failed servers.

There are several clustering software packages available for a wide range of operating systems. When looking at different packages, it is important to pick one that will support your existing hardware and operating system. You will also want to pick a solution that has been tested, is known to be reliable, is simplistic in nature, and whose manufacturer provides good technical support. Some of the most popular shared-nothing clustering packages available and in use today include:

- Microsoft Cluster Server (Windows NT/2000)
- Legato’s Cluster Enterprise (Windows NT/2000 and Linux)
- Network Specialists’ Double-Take (Windows NT/2000)

## Under- and Over-Subscription

Under- and over-subscription of a server means exactly that. Under-subscribing a server, or group of servers, means that you have more resources than customers require. Over-subscription is the opposite, meaning that you have too few resources compared to customer requirements. At some point, you will find yourself in one of these situations, either under- or over-subscribed. Because of the growth rate of services over the Internet, it is rare to find a server that isn’t used to its full potential, not a drop less, and not a drop more.

The advantages and disadvantages to under- and over-subscription are fairly obvious. When a system is under-subscribed, you are probably spending more money on hardware than you are receiving in revenue from your customers—this could be serious. On the other hand, your customers are receiving fast and efficient service, and might be willing to pay more if you guarantee a high level of

availability, and deliver it. In contrast, when a service is over-subscribed, you are probably pulling in much more money than you are spending on hardware. At the same time, your customers may have timeout issues, and your service might appear unreliable. In this situation, you will probably lose customers, and will not be able to charge as much as other ASPs with have better reliability.

With that said, probably the only way to ensure a high revenue stream is to over-sell your services, and over-subscribe your servers to some degree. If not, it may be very difficult to have a positive balance at the end of your fiscal year. Over-subscription does not always mean that you will be turning customers away, and offering them unreliable service. In fact, if implemented correctly, your customers should never have a problem and never know that there might have been issues.

The fact is, your customers will more than likely possess very different usage patterns. You are probably already very familiar with the usage patterns of your clients, especially if you are providing Internet services for them. It does not take much extrapolation to figure out their requirements and usage patterns based on the information you already possess. One customer might always use the system in the daytime, while another might always use it at night, meaning that only one of these customers will use the system at any given time.

This can be expanded upon further. You might have 10,000 customers. Maybe half of them are using the system at the same time. That means that to provide an adequate service offering, you only need a group of servers that is capable of serving around 5000 clients at any given time. If you made your capacity 10,000 users, your servers could handle everyone connected at the same time, but the servers would be under-subscribed all of the time. On the other hand, if you built only 5000-user capacity into your system, you would use your servers to their full potential almost all of the time. There might be some instances when there are more users than expected, and some will be turned away or receive timeouts.

This may be acceptable, or else you could consider building a little overhead into the system to allow for these situations. Either way, you should always attempt to over-subscribe your systems without hindering your customer satisfaction.

You should, however, use caution when oversubscribing your services. Depending on the type of service offered, there might be a time when all or most of your users access the system at the same time, regardless of their “normal” patterns of usage. For instance, if you offer time-card services, all of your customers might access the system on Friday afternoon or Monday morning in order to submit their timesheets for the week. This could be a bad

situation if you do not have the capacity for all of these users, because you were not expecting such a trend.

Regardless of the service offered, it is always wise to research your clientele usage patterns, and attempt to predict the maximum number of concurrent users at any given time. This will help you make a rough estimate for the number of servers required for a particular service offering. After all, you would not want to make the mistake of offering inferior service in the beginning, and upgrading your systems to handle the capacity, while at the same time losing customers only to find your servers severely under-subscribed. That could put you in a situation far worse than the one you were in to begin with.

## Network Service Considerations

Since your server will be attached to a network, there are a variety of network-related service offerings that can provide storage solutions, network-based data backup and recovery, and error monitoring, just to name a few. These solutions could provide for a better overall design, and efficiency of the network, as well as provide a method for early detection of failed systems.

To better understand some of these offerings, we'll now discuss several key areas that should be important to you.

## Network Storage

*Network storage* defines the ability to store information on a remote system connected over a network. This kind of storage can help save money, time, and resources by centralizing your data. Typically, we think of servers that possess their own internal storage. In this situation, if a system runs low on storage capacity, we will need to install a larger or additional hard drive to alleviate the problem. This might work in small environments; however, if many of your servers require more storage space, I doubt that you will want to install a new hard drive in each of those devices.

Instead, wouldn't it be great to have one or a couple of servers that the rest of the servers use for data storage? In this situation, you could easily upgrade a few servers, while passing the benefit on to many of your servers. In addition to this scenario, there may be times when you need several servers to have access to the same data. If this is the case, I doubt that you will want to replicate this data many times, and develop methods for synchronizing it between all of the servers. If you decide to use a centralized network storage device instead, it may be possible for all of these servers to access the exact same data without the need to replicate it to many servers.

When looking into network storage solutions, you will find that there are many solutions offered. Some of these solutions are very simplistic, while others are very complex and include many bells and whistles. Among the more complex solutions are Network Attached Storage (NAS) and Storage Area Networks (SAN).

A NAS is essentially a single-purpose device designed solely to attach and serve data across a network. They can be very nifty devices that support most data delivery protocols and offer many management features. SANs, on the other hand, are basically networks in their own right, and can incorporate routers, switches, and other network devices. A SAN effectively connects storage devices and hosts together using a network that is completely separate from the data network.

Both these technologies can help solve very complex storage problems and requirements, and should be explored in some detail in order to fully understand the benefits and implications offered by both. SAN and NAS is explained in detail in upcoming chapters of this book. However, you do not always need to use very complex solutions such as these in order to connect to a centralized storage system.

If you have a smaller environment, or want a more simplistic solution, you might get all the performance and features you require from centralized servers that are using the Network File System Protocol (NFS), Server Message Block (SMB), or Common Internet File System (CIFS) to connect to shared resources.

## Network File System Protocol

NFS was first released in 1984 by Sun Microsystems Corporation. This was around the same time that microprocessors had become available and the old mentality of mainframes and minicomputers slowly started to dissipate. It was originally designed to be a stateless file access protocol to be used across potentially unreliable networks, meaning that if the server were to crash or become unavailable, the client machine would wait for the server to come back up instead of producing an immediate error.

For instance, the second version of NFS would not return with a result until all the data being written had been stored on a permanent storage media. If the attempt never returned to the client with a positive result, it would be up to the client computer to make the decision whether to continue to wait for the write to finish, or return an error result and cancel the write command altogether. Eventually, as time went on, it became apparent that new features needed to be added to NFS since it had become a very useful protocol. This led to the development of version 3 in 1995.

One of the main reasons or necessity for creating the NFS protocol was the cost of file storage. Even though x86-based computers were making their way into environments where only mainframe computers had been used in the past, each of these systems contained its own storage media, the internal hard drive. In many cases, one server would not have enough hard drive space and would require additional storage space, while another computer was under-utilizing its internal storage. NFS was one of the first protocols to allow a centralized file server that could be used to store data for multiple systems. This meant that you could easily add additional hard drives to one central device, while increasing the storage capacity of numerous host systems and servers. This was a new concept, and began to gain acceptance and use very quickly.

Today, many systems use NFS to connect servers to centralized storage. Since NFS was designed on the Unix platform, it has remained a Unix tool, for the most part. It is possible to find NFS servers and clients that run under other operating systems, such as Windows, but they are not very desirable since they are not native to the particular operating system. In fact, in a production environment, you should attempt to use file transfer utilities and services that are native to the operating system.

For instance, if you are using Windows operating systems throughout your network, it is possible to access shared resources on another Windows machine using the built-in NetBIOS protocol, instead of adding an NFS client to a computer. This is because the built-in protocol should outperform any application that is added to the system. The same would be true for Unix environments; NFS would more than likely outperform a NetBIOS-based client, simply because more specific development has gone into the built-in application.

If you are using many different types of servers within your network, and need to share files between them, you may want to consider breaking the rule stated earlier and use a standard method throughout your server environment. In such a case, installing an NFS client on your Windows server might be a wise choice. It could definitely help with the administration of shares, since there would only be a single file transfer method with which to be concerned.

## Server Message Block Protocol

There are some other methods in addition to NFS that are definitely worth mentioning. Server Message Block (SMB), for instance, is a protocol developed by Microsoft that also provides a method for client applications to access data over the network. SMB is designed to run on top of TCP/IP, IPX, and NetBEUI protocols.



The entire range of Windows products provides support for the SMB protocol, and with an additional shareware program called Samba, it is possible to configure both client and server support on most Unix systems as well. The drawback to using this protocol is the limited support available for SMB on the Unix platform.

Since the software available for Unix servers is in the public domain, no one company “owns” it. If you are experiencing difficulties or bugs with the Samba application, you may have to wait for another release of the software to fix the errors, and even then, you may get a version with even more errors than the original you replaced. This sounds like a huge issue, but in reality, it is a rather simple protocol and Samba is fairly simple to configure for a variety of uses, as is shown in Figure 3.11.

### Figure 3.11 Samba Configuration

---

```
# This is the main Samba configuration file. You should read the
# smb.conf(5) manual page in order to understand the options listed
# here. Samba has a huge number of configurable options (perhaps too
# many!) most of which are not shown in this example

# Any line which starts with a ; (semi-colon) or a (#)
# is a comment and is ignored. In this example we will use a
# for commentary and a ; for parts of the config file that you
# may wish to enable

# NOTE: Whenever you modify this file you should run the command
#"testparm" to check that you do not have any syntax errors.

#===== Global Settings=====
[global]
# workgroup = NT-Domain-Name or Workgroup-Name workgroup = TEST
# server string is the equivalent of the NT Description field
# server string = TEST-SERVER
# This option is important for security. It allows you to
# restrict connections to machines which are on your local
# network. The following example restricts access to two C class
# networks and the "loopback" interface. For more examples of
# the syntax see the smb.conf man page
; hosts allow =192.168.1. 10.1.1. 127.
```

---

Continued

### Figure 3.11 Continued

---

```
# if you want to automatically load your printer list rather
# than setting them up individually then you'll need this
printcap name = /etc/printcap
load printers = yes

# It should not be necessary to spell out the print system type
# unless yours is non-standard. Currently supported print
# systems include:
# BSD, sysv, plp, lprng, aix, hpux, qnx
printing = lprng

# Uncomment this if you want a guest account, you must add this
# to/etc/passwd otherwise the user "nobody" is used
; guest account = guest

# this tells Samba to use a separate log file for each machine
# that connects
log file = /var/log/samba/%m.log

# Put a capping on the size of the log files (in Kb).
max log size = 0

# Security mode. Most people will want user level security. See
# security_level.txt for details.
security = user

# Use password server option only with security = server or
# security = domain
password server = TEST-DOMAIN-SERVER

# Password Level allows matching of n characters of the password for
# all combinations of upper and lower case.
password level = 8
username level = 8
```

---

Continued

**Figure 3.11 Continued**


---

```

# You may wish to use password encryption. Please read
# ENCRYPTION.txt, Win95.txt and WinNT.txt in the Samba
# documentation.
# Do not enable this option unless you have read those documents
encrypt passwords = yes
smb passwd file = /etc/samba/smbpasswd

# The following are needed to allow password changing from
# Windows to update the Linux sytsem password also.
# NOTE: Use these with 'encrypt passwords' and 'smb passwd file'
# above.

# NOTE2: You do NOT need these to allow workstations to change
# only the encrypted SMB passwords. They allow the Unix password
# to be kept in sync with the SMB password.
unix password sync = Yes
passwd program = /usr/bin/passwd %u
passwd chat = *New*UNIX*password* %n\n
*ReType*new*UNIX*password* %n\n
*passwd:*all*authentication*tokens*updated*successfully*

# Unix users can map to different SMB User names
username map = /etc/samba/smbusers

# Using the following line enables you to customise your
# configuration on a per machine basis. The %m gets replaced
# with the netbios name
# of the machine that is connecting
include = /etc/samba/smb.conf.%m

# Most people will find that this option gives better performance.
# See speed.txt and the manual pages for details
socket options = TCP_NODELAY SO_RCVBUF=8192 SO_SNDBUF=8192

# Configure Samba to use multiple interfaces
# If you have multiple network interfaces then you must list
# them here. See the man page for details.
interfaces = 192.168.1.1/24 10.1.1.1/24

```

---

Continued

## Figure 3.11 Continued

---

```
# Configure remote browse list synchronisation here
# request announcement to, or browse list sync from:

# a specific host or from / to a whole subnet (see below)
  remote browse sync = 192.168.1.2 10.1.1.2

# Cause this host to announce itself to local subnets here
  remote announce = 192.168.1.255 10.1.1.255

# Browser Control Options:
# set local master to no if you don't want Samba to become a
# master browser on your network. Otherwise the normal election
# rules apply
  local master = no

# OS Level determines the precedence of this server in master
# browser elections. The default value should be reasonable
  os level = 33

# Domain Master specifies Samba to be the Domain Master Browser.
# This allows Samba to collate browse lists between subnets.
# Don't use this if you already have a Windows NT domain
# controller doing this job
  domain master = no

# Preferred Master causes Samba to force a local browser
# election on startup and gives it a slightly higher chance of
# winning the election
  preferred master = no

# Enable this if you want Samba to be a domain logon server for
# Windows95 workstations.
  domain logons = no

# if you enable domain logons then you may want a per-machine or
# per user logon script run a specific logon batch file per
```

**Figure 3.11 Continued**

---

```

# workstation (machine)
; logon script = %m.bat

# run a specific logon batch file per username
; logon script = %U.bat

# All NetBIOS names must be resolved to IP Addresses
# 'Name Resolve Order' allows the named resolution mechanism to
# be specified. The default order is "host lmhosts wins bcast".
# "host" means use the unix system gethostbyname() function call
# that will use either /etc/hosts OR DNS or NIS depending on the
# settings of /etc/host.config, /etc/nsswitch.conf and the
# /etc/resolv.conf file. "host" therefore is
# system configuration dependant. This parameter is most often
# of use to prevent DNS lookups in order to resolve NetBIOS
# names to IP Addresses.

# Use with care! The example below excludes use of name
# resolution for
# machines that are NOT on the local network segment - OR - are
# not deliberately to be known via lmhosts or via WINS.
name resolve order = wins lmhosts bcast

# Windows Internet Name Serving Support Section:
# WINS Support - Tells the NMBD component of Samba to enable
# it's WINS Server
wins support = yes

# WINS Server - Tells the NMBD components of Samba to be a WINS
# Client Note: Samba can be either a WINS Server, or a WINS
# Client, but NOT both.
wins server = 192.168.1.2

# WINS Proxy - Tells Samba to answer name resolution queries on
# behalf of a non WINS capable client, for this to work there
# must be at least one WINS Server on the network. The default
# is NO.
wins proxy = yes

```

---

**Continued**

## Figure 3.11 Continued

---

```
# DNS Proxy - tells Samba whether or not to try to resolve
# NetBIOS names via DNS nslookups. The built-in default for
# versions 1.9.17 is yes, this has been changed in version
# 1.9.18 to no.
    dns proxy = yes

# Case Preservation can be handy - system default is _no_
# NOTE: These can be set on a per share basis
; preserve case = no
; short preserve case = no

# Default case is normally upper case for all DOS files
; default case = lower

# Be very careful with case sensitivity - it can break things!
; case sensitive = no

#===== Share Definitions =====
[homes]
    comment = Home Directories
    browseable = no
    writable = yes

# Un-comment the following and create the netlogon directory for
# Domain Logons
; [netlogon]
;   comment = Network Logon Service
;   path = /home/netlogon
;   guest ok = yes
;   writable = no
;   share modes = no

# Un-comment the following to provide a specific roving profile
# share the default is to use the user's home directory
;[Profiles]
;   path = /home/profiles
```

---

Continued

**Figure 3.11** Continued

---

```

; browseable = no
; guest ok = yes

# NOTE: If you have a BSD-style print system there is no need to
# specifically define each individual printer
[printers]
    comment = All Printers
    path = /var/spool/samba
    browseable = no

# Set public = yes to allow user 'guest account' to print
    guest ok = no
    printable = yes

# This one is useful for people to share files
;[tmp]
;    comment = Temporary file space
;    path = /tmp
;    read only = no
;    public = yes

# A publicly accessible directory, but read only, except for
# people in the "staff" group
;[public]
;    comment = Public Stuff
;    path = /home/samba
;    public = yes
;    writable = yes
;    printable = no
;    write list = @staff

#
# A private printer, usable only by fred. Spool data will be placed in
# fred's home directory. Note that fred must have write access to the
# spool directory,
# wherever it is.
;[fredsprn]
;    comment = Fred's Printer
;    valid users = fred

```

---

Continued

**Figure 3.11 Continued**

---

```
; path = /homes/fred
; printer = freds_printer
; public = no
; printable = yes

# A private directory, usable only by fred. Note that fred
# requires write access to the directory.
;[fredsdir]
; comment = Fred's Service
; path = /usr/somewhere/private
; valid users = fred
; public = no
; writable = yes
; printable = no

# A service which has a different directory for each machine
# that connects this allows you to tailor configurations to
# incoming machines. You could also use the %u option to tailor
# it by user name.
# The %m gets replaced with the machine name that is connecting.
;[pchome]
; comment = PC Directories
; path = /usr/pc/%m
; public = no
; writable = yes

# A publicly accessible directory, read/write to all users. Note
# that all files created in the directory by users will be owned
# by the default
# user, so any user with access can delete any other user's
# files. Obviously this directory must be writable by the
# default user. Another user could of course be specified, in
# which case all files would be
# owned by that user instead.
;[public]
; path = /usr/somewhere/else/public
; public = yes
; only guest = yes
```

---

Continued



**Figure 3.11** Continued

---

```

; writable = yes
; printable = no

# The following two entries demonstrate how to share a directory
# so that two users can place files there that will be owned by
# the specific users. In this setup, the directory should be
# writable by both users and should have the sticky bit set on
# it to prevent abuse. Obviously this could be extended to as
# many users as required.
;[myshare]
; comment = Mary's and Fred's stuff
; path = /usr/somewhere/shared
; valid users = mary fred
; public = no
; writable = yes
; printable = no
; create mask = 0765

```

---

Once Samba has been installed on your Unix system, the first step you will need to take is to set up your configuration file. As was shown in Figure 3.11, this file is fairly straightforward, and usually comes with many comments to help explain each variable. The bulk of the configuration pertains to the actual shared resources and permissions required to access these resources. In our sample configuration file, these are rather self-explanatory. Probably the most important aspect of this configuration file is the Windows NT domain settings. These settings can vary depending on your particular situation, but in all cases, it is very important to configure the same NT domain or workgroup in which your existing Windows clients reside.

If this is incorrect, the devices may not see each other on the network. Once the configuration is complete, the server is easy enough to start. There are two daemons required to run the Samba server on your system.

The first is the NetBIOS Message Block Daemon (NMBD), which provides NetBIOS name support for your client nodes. The second is SMBD, which is the actual Samba program that allows users to access the configured resources. Once that is complete, and your client nodes are configured correctly, you are done. One last note: it would probably be a good idea to make these two daemons start automatically at boot time.

## Common Internet File System

Recently, the SMB protocol has been further expanded upon, and has developed into what is now called the Common Internet File System (CIFS) protocol. CIFS is a public version of the SMB protocol. Although this protocol has yet to be formally standardized, the current version is recognized among many software and hardware designers, and has started to gain widespread use in LAN environments. CIFS is viewed as a complement to Internet-related protocols such as FTP and HTTP. In fact, over the next few years, we may begin to see use of this protocol in the public Internet sector.

Today, there are many alternatives and numerous ways to attach devices to a centralized storage system. Of the more simplistic forms are the Network Attached Storage (NAS) devices that have been designed specifically to attach to and serve data over the network. Most of these devices support multiple protocols for serving data to and from the client nodes. In fact, NAS devices can help simplify a multiple operating system environment. With these devices, it is possible to connect to your Unix hosts using NFS, while connecting to your Windows hosts using CIFS. This can help improve the efficiency, since you are always using the native protocol for the operating system. In addition, management is still centralized, which allows for the best of both worlds.

## Data Backups and How They Can Affect You

Data backups as you probably already know, are one of the most important, if not *the* most important, duties that you will perform for your customer. Many disasters can result in loss of data, but at the top of the list would probably be hardware failure. Although hardware platforms have become more reliable over the years, the fact still remains that your data is stored on what is essentially a mechanical device; a disk that rotates at very high speeds with another bit of metal called a head that floats left and right across the surface of the disk many times a second. What this means is that one day, one of these rapidly moving bits of metal will wear out! When that day comes, and it will, all of the data stored on that disk will be lost forever unless you have a good backup.

In addition to a hardware failure, there is a high possibility that your operating system could crash, thereby forcing you to rebuild the entire system from scratch. With a good and complete backup, it is possible to restore the entire system without having to resort to other drastic measures, such as kicking the

server several times in a fit of rage. Other unforeseen disasters could also result in serious loss of data, such as a compromised system, accidental deletion of files, or an incorrectly configured software application. These reasons alone should help serve as a reminder to design and maintain a backup solution for, at a minimum, your essential systems.

When deciding on the type of backup solution to implement for your network, you should ask yourself a couple of simple questions, such as:

- Could your customers survive if you lose all the data on their databases? Could you?
- What would be the impact of losing months of your e-mail or a customer's financial information in the event of fire or some other catastrophic incident?
- Do you have a large group of lawyers whom you trust to handle all the lawsuits resulting from a loss of customer data?

Unfortunately, backups tend to be low on the priority list for many companies, and for some reason when something bad happens, it tends to crop up in several places. The reason for this is probably due in part to your customers who will expect that you recover their data immediately. In this section, we will attempt to cover different alternatives, and discuss factors that might help influence your decisions when planning a data backup strategy.

You have a vast amount of choices when considering data backup solutions. In just the past few years, vendors have increased the size and the amount of hard disks that they install in their servers to help house the increasingly bulky applications and the data that has become so commonplace. This means that hard drive manufacturers have had the equally difficult task of attempting to keep up with the industry by producing devices and media that allow room for more storage. If you plan to regularly back up between 20GB and 500GB of data, you will need a flexible, reliable, and scalable solution that also fits your budget.

If you require an extremely fast and efficient backup system that includes rapid data restoration speeds, you might consider a high-end system, but it can be pricey. Vendors optimize solutions in this category for high transfer rates, high access speed, and compatibility with robotic equipment. Typical transfer rates are several megabytes per second, and pricing for these drives starts around \$5,000.

Some of the things that you need to take into consideration when you purchase a backup unit include the amount of data you need to back up, whether your backup software supports the unit, and the amount of money you want to

spend. It would also be a good idea to consider the value of the data that you are trying to protect, and even attempt to put a price tag on this data as well as the time required to rebuild the data if it were lost. You should then be willing to spend an equal amount of money on a backup solution, or at least as much as your calculations and budget permit. In this way, you can have some level of understanding as to how much you will spend on backup solutions within a given time.

## Configuring & Implementing...

### Advanced Storage Solutions

If your backup system is vast, requires even more speed, or you are concerned that your data backup will cause problems with the rest of your network, you will probably want to look into providing a server-less backup solution, or a storage area network (SAN). In either case, we recommend reading through Chapter 6 to get a better understanding of storage options.

Many of the solutions given in that chapter offer advanced features that can help reduce the overhead required for your data backup, and some solutions will even allow you to do real-time site replication and mirroring in addition to data backup solutions.

You should determine the maximum backup window possible and choose a unit that provides the technology you need to back up and restore data well within that timeframe. For example, some products use optical disks, which can be a great solution for hierarchical storage management. This will, however, translate into a bad idea when backing up large amounts of data, since the speed of this particular kind of system is fairly slow.

When looking for devices, you need to find a balance between speed, reliability, management, and features, and be certain that your proposed solution fits your requirements in these areas. Even if you are already comfortable with a certain manufacturer and are thinking about using their backup software and devices throughout your network, you will still need to research the solution to ensure it will be a good fit in your network. In addition, always try to foresee your future requirements, and make sure that your solution will not become obsolete in a short period of time.

Tape devices are known for their long-lasting performance with few hardware failures, which is partly attributable to the reliable tape media and robotics that most systems include today. Despite this reliability, you should always research a vendor's service contract to ensure that it includes same-day onsite or overnight cross-shipment service if the unit fails. Due to the strong competition in the tape backup arena, there has been an overwhelming amount of performance claims that are little more than declarations with little backing. When comparing devices, it would be a good idea to determine the type of files the vendor used to test the drive, and the type of test the vendor ran.

Be careful, since vendors have been known to provide inaccurate results, by using compressed files that make their performance look better than it actually is, or some other method of manipulating the results. Many times, these tests do not consider the fact that many file types can't be compressed, so the tests do not simulate real-world scenarios. Your best bet is to stick with the well-known brand names such as HP and Exabyte.

## Software Selection

You will most likely use a third-party backup program as opposed to the generic ones that sometimes come with your operating system, or storage devices. Some of the products that you will run across such as ARCserve, Veritas Backup Exec, UltraBac, or NovaStor, will allow advanced scheduling with various levels of flexibility. These applications are also relatively easy to configure and maintain. Most of the systems available will operate under NetWare, Windows NT/2000, and the various flavors of Unix. However, you should always make sure that the backup software you ultimately choose is compatible with your server platforms.

The backup process can be made to operate more efficiently with the cooperation of the target system. It is common for backup systems to use programs called *agents* that manage the communication between the target system and the backup host. These agents give the backup host access to different servers on the network so that their local drives can participate in the backup system that you implement. Agents are also used to share some of the load on the server that is being backed up, which can help to speed the entire process.

Another use for these agents will be discovered when you need to back up target systems that differ from the platform of your backup host. For example, if your backup host is a Macintosh server, you may need a Unix or NT/2000 agent that understands these specific file systems, and allows for backup of data and applications. The agents can also optimize the performance of the backups by taking advantage of the microprocessor on the target system to pump data to the backup server.

Open File Options (OFO) is another part of the software consideration process that you need to take into account. OFOs are usually purchased separately, and are similar to agents in that they allow you the unique ability to back up specific applications while they are still running, such as a database. As you can imagine, this will help to add significant value to the product that you are evaluating, and depending on the application that you are using, it may even make it easier when it comes to your decision-making process.

In the past, backup solutions introduced a much greater challenge, since many services needed to be completely shut down to allow for a full backup. Fortunately, this is not an issue any longer, as long as you are willing to purchase a solution that incorporates OFO and these types of advanced features.

### *Scheduling Backups and Tape Rotation*

One of the defining factors between backup systems is how tapes are rotated and what files get backed up to which tape. Each rotation method has different advantages that can be applied to systems and provide for different results. Some of the differences between these methods are the amount of time it takes to perform a backup or restore, the number of copies made, and the number of different revisions of a file that gets stored on tape.

Upon planning how you are going to develop your backup strategy and rotation method, it is critical to look at what data is on your system, how critical it would be if some data were lost, and how fast a system would need to be operational if the system did become damaged or inoperative.

Another important task would be to assign one individual the responsibility of administering the backup solution. This will make a single person accountable, and ensure that your backups are restored on a consistent basis. Your plan should include a specific time or times when backup is done, and should allow for testing of backups with a periodic test rotation.

One method of backup requires that a full image or copy of the system be put on tape every day. Each day, a different tape is cycled through the system, ensuring that you are capable of performing a complete restore from a single tape or set of tapes from any day that a backup was performed. The downside to this is the amount of tapes needed to accomplish such a task.

Moreover, if you have a large amount of data, this could make for a very long backup window in order to complete the full backup, and in some cases, your backups may end up taking place during peak hours. This type of backup solution is commonplace for smaller servers with less than 20 gigabytes of data. However, if you have more than this amount of data, we would suggest looking into a different tape rotation solution.

## Designing & Planning...

### Frequent Backups

If your data is altered frequently, and it is critical to have multiple revisions of the file throughout the day, you may need to plan for a system that will allow for multiple backups in a single day. This type of system would obviously require more tapes, and an efficient and swift system will be required in order to complete the backup process several times in a single day.

A different approach to the previous method is to make one backup set that contains a full image of the system, and use every subsequent tape to back up only the files that have been altered or were changed in some way since the last full backup. This type of backup is called a *differential backup*, and allows for the system to be fully restored using two tape sets, one that contains the full backup and a second that contains the newest set of data.

A variation of this method would be to copy only the files that have changed since the last differential to the tape. This type of backup is called an *incremental backup*. This would take less time to back up and is an excellent solution for systems that need to have multiple backups performed in a single day. It would, however, require more time to restore, since you may need several tape sets to perform a restore.

There are some rotation methods that allow for files to be stored multiple times on multiple tapes so that you can have different versions of the same file. This allows a revision history of files to be stored on the tapes in case a past revision of a file should be needed. Many times, this is necessary in order to prove or verify the alleged history of a file, or other times you may require a restore of an older copy of the file because the latest version has become corrupt.

This solution allows great flexibility in single file restoration, but can hinder the time to restore the entire system. It can be especially confusing when you need to restore multiple files that are all from different sets of tapes and revisions. In some cases, you might be busy all day switching between different backup sets.

As discussed previously, there are essentially only three main types of backup solutions possible:


- Full

- Differential
- Incremental

After planning the method you will use to place files on your backup media, you need to choose a rotation method for how tapes are going to be run through the system. There are several tape rotation methods that will incorporate the three backup types listed.

The Grandfather, Father, Son (or to be more politically correct, the Grandparent, Parent, Child) is a simple method that has been used for many years. In this method, tapes are labeled by the day of the week with a different tape for each Friday in the month and a different tape for each month of the year. Using a tape for Saturday and Sunday is optional, depending on whether you have files updated over the weekend. Figure 3.12 depicts the Grandparent, Parent, and Child rotation scheme based on a two-month rotation schedule.

**Figure 3.12** Grandparent, Parent, and Child Tape Rotation Scheme





In the Figure 3.12, a different tape is used for every weekday in a two-month cycle. Since some months have more than four weeks, it will take at least 20 tapes for regular backups to be performed for a single month, and over 40 tapes to back up the system for two months without overwriting any tapes. At the end of each month, one tape should be removed from the set and archived. At the end of the two-month cycle, two more tapes should be added to replace the tapes that were removed for archiving, and the entire cycle should begin again, overwriting the existing data on the tapes.

The Tower of Hanoi solution is named after a game in which you move a number of different-sized rings among three poles. In the game, you start out with all the rings on one pole and must move all the rings to another pole. You can never have a ring on top of one that is smaller than it is. The idea is that you must move them in a certain order to accomplish the task. The correct order of ring movements is:

```
A-B-A-C-A-B-A-D-A-B-A-C-A-B-A-E-A-B-A-C-A-B-A-D-A-B-A-C-A-B-A-F-
A-B-A-C-A-B-A-D-A-B-A-C-A-B-A-E-A-B-A-C-A-B-A-D-A-B-A-C-A-B-A-G-
A-B-A-C-A-B-A-D-A-B-A-C-A-B-A-E-A-B-A-C-A-B-A-D-A-B-A-C-A-B-A-F-
A-B-A-C-A-B-A-D-A-B-A-C-A-B-A-E-A-B-A-C-A-B-A-D-A-B-A-C-A-B-A-H
```

When applied as a backup strategy, we use the same order to rotate tapes through a tape drive, making a complete image of the system each day. Although the theories behind why this works as a good rotation method are beyond our discussion, the benefit of this rotation is that you will always have an older version of a file on one tape. In the case listed previously, you would have used eight tapes. When you consider that you used one tape per day, you could have a copy from as long as 128 days previous.

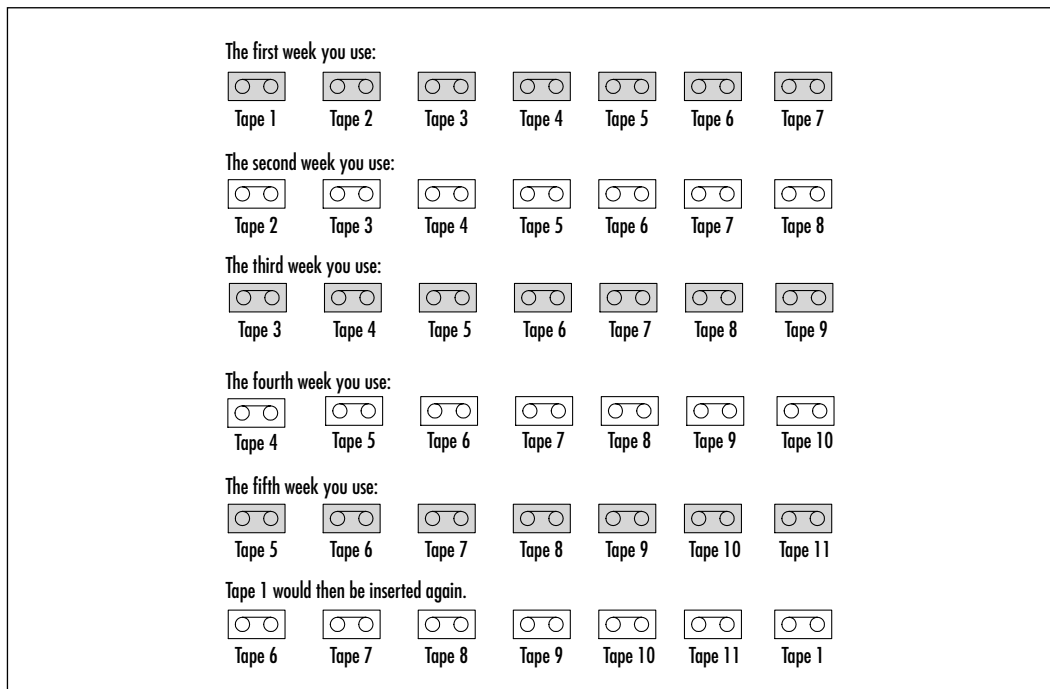
As a further example, tape F will contain a full backup of every file on the system from every 128 days. If your system becomes infected with a virus, you could restore a file without a virus as long as you did not have the virus for more than 128 days. Furthermore, if you require a backup solution that will keep data for longer than 128 days, simply add additional tapes to this particular rotation method. In fact, the number of tapes used with this method depends solely on how far back you would like to be able to go.

There are several other variations on this rotation method. First, using the preceding example, if you performed backups twice per day, you would be able to capture work in progress during the day, but you would also only have versions from as long as 64 days. Again, this limitation could easily be overcome by adding additional tapes to the rotation.

Another possibility would be to perform a full backup on a tape, and to do incremental backups on the same tape for the remainder of the week. By doing this, you could increase the number of versions available, while at the same time decrease the number of tapes required. However, if this is done it is possible that your tape may run out of space or you could risk losing up to a week's worth of data if that tape has a problem or becomes damaged.

The incremental tape method is another rotation scheme that is in widespread use. Although this method goes by a few different names, they are all essentially the same and are fairly simple to implement. This rotation method involves determining how long you wish to maintain a copy of your data and how many tapes you wish to use. It is based on a labeling method in which tapes are given numbers and are incremented by adding and removing one backup set each week. It can be configured to allow for either five- or seven-day backup schedules. Figure 3.13 depicts an incremental tape rotation method.

**Figure 3.13** Incremental Tape Rotation Method



Continue this rotation for as long as you have tapes, and keep one tape from every week that you perform a backup. This tape should be stored for a certain period of time, depending on your requirements and the number of tapes available

to you. This method will evenly distribute the tape usage, and ensure that different revisions of a particular file are stored on every tape.

The disadvantage to this method lies in the fact that you are still doing full backups. This means that your backup window might be large, and the frequency of the backups could become problematic for your users.

One variation of this method would be to perform a full backup on the first day of every week, and then incremental or differential backups every day after that. In this case, you would set the first tape aside after every week in order to keep a full backup.

An advantage of this system is that tapes can be removed or added to the system at any time if additional file history is needed. The key is to keep a log of the tape sequence and the date on which it was last used. This can be calculated months at a time or even for an entire year if necessary.

## Virus Scanning Suggestions

Computer virus programs have a long history. They are considered by some to serve a useful purpose, while the majority of users will tell you they are malicious programs whose authors should be incarcerated. Regardless of the side you are on, a virus is something you do not want in your production network.

A virus can halt your servers, and can even remove data from your hard disks. What's worse is that it can spread to incorporate the computers throughout your entire network and into your client's networks, infecting every server along the way and leaving mass data destruction in its wake.

In the earliest computers, only one application could be run at a time. This meant that to understand the results or changes that a particular application made, it was vital to always know the initial state of the computer, and to wipe out any leftover data from other programs that had already terminated. To perform these tasks, a small program or instruction was created.

This instruction would copy itself to every memory location available, thus filling the memory with a known number and essentially wiping the memory clean. Although this instruction served a very valuable purpose and allowed for the results of an application to be verifiable, this type of program, or instruction, is considered the first computer virus ever created.

As computers progressed, it became possible to run more than one application on a single computer at the same time. To allow this, it became important to partition the applications from each other, so that they did not interfere with one another and they produced reliable results.

Soon after, applications were developed that had the capability to break these boundaries and transcend the partitions. These rogue applications would use random patterns to alter data and break applications by pointing them to memory locations where they would read incorrect data or overwrite valuable data. Because the patterns were random, if one were to trace the patterns and plot these on a map, they looked much like the holes found in wood that has been partially consumed by worms. These patterns soon became known as “wormholes,” and with the help of the “Xerox Worm,” which was the first virus to spread to infect other computers, these viruses have become known as “Worms.”

Nearly everyone has heard the story of the great city of Troy, and the Trojan Horse that was given as a gift. In the computer industry, there are not only worms, and other viruses, but there are some extremely malicious programs that disguise themselves as other beneficial programs. These are known as Trojan horses. One of the first Trojan horses created disguised itself as a program that would enable graphics on a monitor. It should have been a dead giveaway, because this system was incapable of supplying graphics. However, when the Trojan horse was run, it presented a message that said “Gotcha” while it proceeded to erase the hard drive completely. After this, Trojan horses began to spread quickly through the use of early Bulletin Board Systems (BBS). These BBSs were a precursor to the public Internet that we know today. Many ideas that initially began on BBSs were copied and expanded upon on the new Internet, and so were the Trojan horses.

In today’s environment, any computer connected to the Internet or accessible by many other systems or individuals is likely to become infected with a virus. Viruses have been improved upon so many times that the average virus scanning software will scan for tens of thousands of known viruses.

Some of these viruses can be transmitted when viewing a Web page, others can be e-mailed to users, and still others can be manually installed on a system. There are many different ways to infect a system, and new ways are being developed and discovered every day.

Another truth is that there are also malicious individuals in the world. Some of these individuals will attack certain groups or businesses, while others are not as choosy and prefer to attack at random. Regardless of the person’s intent or the method of infection, it is very important to guard your systems against viral attack and to use an anti-virus application that is reliable and capable of detecting viruses before they actually cause harm to the system.

Unfortunately, these days there is a tendency to think that viruses and Trojan horses are only a concern on systems running the Microsoft Windows family of

operating systems. This is definitely untrue. It is true that the majority of viruses and Trojan horses designed today are aimed at attacking systems that use Microsoft Windows, mainly because the operating system is in such widespread mainstream use and comprises the majority of work and personal computers.

However, other operating systems have been around for a long time, and many viruses and Trojan horses have been designed specifically for them as well. It is also possible for a system that is immune from a particular virus to unknowingly pass a virus or Trojan horse to a system that is susceptible to the infection. For these many reasons, you should install an anti-virus solution that incorporates each of your computers, regardless of the operating system used.

The most popular anti-virus suites come from McAfee, Symantec, and Network Associates (NAI). These tend to be good solutions because they have multiple products that can be used on most operating systems. In addition, all three vendors update their virus definition files at least twice a month, and usually create a new definition any time there is a large breakout of a new virus. Their services are reliable, and have been integrated to work with many types of application software.

These anti-virus software suites do not usually cause problems on the system, but there is always a possibility that they may conflict with another program. If you suspect this to be the case, it might help to temporarily disable the anti-virus software in order to test the software conflict. If there is truly a software conflict, you should contact the manufacturer of both products immediately to see if there is a fix or a way around the problem. You might even consider using a different application or anti-virus package to alleviate the problem. As a last resort, you can disable virus scanning altogether and rely on other virus-scanning possibilities.

In addition to installing and executing anti-virus software on each computer, there are some other possibilities that allow you to catch viruses as they enter the system. If you use a shared file server, it may serve to distribute viruses throughout your network. If the file server becomes infected, or contains an infected file, it is possible to transmit this among any of the devices that access the particular file server.

Making sure that your file server is protected, and provides constant virus-scanning services while data is accessed, can cut the possibility of a viral infection significantly even if desktop virus protection is not in use. The disadvantage to this solution is that it can impact performance, especially if the file server receives a significant amount of simultaneous connections. The exact performance loss will vary widely, and depend on the software configuration, server hardware, and number of users accessing the system at any given time.

To alleviate this issue, it is sometimes possible to disable constant system scanning, and to instead schedule scanning during a period of inactivity. This can certainly help improve the performance, but it can also defeat the purpose altogether, since a virus may not be detected before it is spread throughout the system.

It is also possible to run anti-virus software that plugs into popular e-mail applications, such as Microsoft Exchange and Lotus Notes. These enterprise e-mail servers provide many features and services of which a virus can easily take advantage. Anti-virus software is capable of neutralizing e-mail viruses before they are delivered to mailboxes.

Since many new viruses are e-mail based or at least transmitted via e-mail, this can be a very wise solution; however, it could result in slower e-mail performance, especially when large attachments are being sent through e-mail.

Also available are anti-virus Internet Gateway products that are capable of intercepting e-mail that originated from the Internet. These products will catch and quarantine the majority of viruses before they even touch your internal mail servers. There is a minimal performance impact when using this type of solution, since mail usually flows in from the Internet at a leisurely pace.

When using an Internet Gateway product, make sure that you have a system that will allow you to queue incoming e-mail messages. If mail is received faster than it can be processed by an Internet gateway, it could start dropping or bouncing messages unless you have software that allows incoming messages to be queued.

## Thin Client Solutions

In 1996, a comparative analysis was performed of the five-year life cycle for cost of ownership of network computers using a thin client server such as WinFrame for Windows Terminals server, versus the five-year lifecycle cost of ownership for multiple personal computers and a Windows NT-based server. When all aspects were considered, such as the cost of hardware, software, administration, support, and upgrades, this research showed that a company could reduce its five-year total cost of ownership by over 50 percent.

One of the primary focuses for an ASP is to ensure the delivery of its products or services to each client's desktop. For example, if an ASP is hosting an application for a company—let's call them Company X—the ASP has to provide the means for all end users to access particular applications. One approach is to deliver an application to the client using the client/server model.

This approach is based on the idea that all processes are handled at the client level, meaning that the actual computing and data alteration is performed on the client device, and is highly dependent on the capabilities that this machine

possesses. The other approach, which is highly suited for an ASP is the thin client model.

Thin client computing allows the delivery of applications from centralized servers to many remote clients. By using this technology, ASPs are able to deliver any application that runs on their centrally managed server or server farms to remote client desktops. When this is accomplished, the actual computing is taking place on the servers, and the client systems are only receiving graphical updates.

The client devices are essentially acting as terminals, and only serve as an interface to the server. This means that a very powerful computer or group of computers can be installed at the ASP, making it easier to guarantee a certain level of performance to the customer. There are many thin client technology manufacturers in the marketplace today; however, our discussion will focus primarily on Citrix Systems' approach to thin client computing. Citrix is the current industry leader, and uses a proprietary protocol called the Independent Computing Architecture (ICA).

## ICA Protocol

Independent Computing Architecture (ICA) allows the delivery of an application from a centralized server to any end-user desktop, regardless of the operating system or platform. ICA clients are available for every major operating system in the market, including Windows 2000/NT/98/CE, Solaris, SCO Unix, Linux, MacO/S, OS/2, and to provide connectivity to other devices, they have recently added support for most Web browsers. In addition, the ICA protocol only consumes around 10 to 20 KB of bandwidth, which is very little when compared with the bandwidth consumption of today's applications. Low requirement in bandwidth is achieved because only screen refresh, mouse clicks, and keystrokes are sent across the pipe; execution and processing of the application is all done on the server.

When considering application delivery, ASPs should be concerned with two critical issues:

- Heterogeneous operating systems
- Bandwidth requirements

### *Heterogeneous Systems*

The reality is probably that many of your clients are running multiple operating systems in their enterprise. In order to effectively provide services to these

customers, you will need to make sure your client's end users are able to access and use your applications regardless of the operating system installed on their desktops.

In addition, you will need to provide them with a performance guarantee, and will want to reduce your customer support costs. In this type of environment, thin client architecture can definitely save the day. If a client is using an unsupported operating system, it is easy to have him or her access network resources using a Web browser that connects to the thin client server. This is key, since every operating system you encounter should incorporate the ability to use a Web browser.

### *Bandwidth Requirements*

These days, applications are very bandwidth intensive, and as time goes on, more bandwidth is required and will be consumed. End-user satisfaction depends greatly on an application's response time. If your clients receive slow response times, they will tend to be unhappy, and think you deliver an inferior service. If, on the other hand, your service is fast and responsive, it will improve your customers' productivity and make for a much better environment for both them and you.

To alleviate these bandwidth concerns, you could always allocate more bandwidth to satisfy your clients. This could be done by building more or larger pipes in your network, or increasing the amount of bandwidth available to a particular client. It is also possible to do Quality of Service (QoS) within your network, and give certain applications a higher priority over other network functions. Although this might work, without the proper amount of bandwidth, it will cause some other function to perform slowly, and rob other systems of bandwidth.

All of these solutions are not very cost effective for you or your clients. Instead, a thin client solution can provide a drastic reduction in client/server overhead, deliver quick and reliable service to your customers, and allow more headroom for other network services and functions to use the available bandwidth.

Thin client technology addresses these two major concerns and several others. It allows applications to be delivered in a cost-efficient manner and without the restriction of any particular operating system. It will also help you reduce your support costs, which will ultimately translate into a better revenue stream for your company.

All these factors could even allow you to provide a cost reduction to your customers, making your model attractive to other customers and businesses. Since thin client technology can solve so many ASP-related issues, it will prove beneficial to at least look into the services offered, and consider the advantages and disadvantages for your particular company.



## Maintenance and Support Issues

Now that you have planned your server architecture and applications, designed a completely fault-tolerant solution, installed your servers, and began providing services to your customers, you must take on the task of maintaining these systems. You might be thinking that this is the easy part, but in reality, this is probably what will make or break your company.

At this point, you have revenue coming in the door and everything looks on the up and up. However, if you do not maintain your systems effectively, your systems might break, and your customers may desert you. If this happens, you will probably receive a bad reputation in the marketplace, and your competitors, seeing that as an excellent opportunity, will strike fast and may consume your entire revenue stream.

Unless you do not take pride in your work, are planning to get another job, and could care less about the company, this is probably the last thing you will want to see happen. This means that you will need a solid plan in order to effectively perform maintenance on your systems.

## Planned Upgrades

Eventually, every piece of hardware and software operated by your company will need an upgrade of some sort. This might be due to a lack of features that the old system did not possess, a need for further functionality, a way of fixing bugs, or a method of alleviating strain on a server. Whatever the case may be, given enough time, an upgrade will become more of a necessity as opposed to a luxury.

In some rare cases, it may be possible to perform these upgrades in one single session. It might seem tempting to schedule one big system downtime well in advance and notify your customers of an outage for a day or two. Although this might work in a small company with a minimal customer base, chances are that you have a much more complex set of systems and an architecture that requires a small army to maintain it effectively.

When you consider that you might be performing hardware upgrades as well as software upgrades, and that one upgrade might cause another, it just does not make sense to even attempt to upgrade the servers all at once. Besides, I doubt that your customers will be able to adjust to the fact that their service will be interrupted for an entire day or more. Instead, it is going to be far more efficient to schedule routine maintenance windows for your system, and inform your customers well in advance. In this way, your customers can expect a possible downtime or loss of connectivity on a predefined day during a certain time. Between

this timeframe, it will be possible to upgrade some of your servers and applications, leaving the rest to be serviced during one of the next scheduled downtimes.

It makes the most sense that these downtimes should not be scheduled during a time of peak usage, and should instead be performed during the time of most infrequent use. It may be that this is in the middle of the day on every other Wednesday, or it may be at very early hours of the morning on every Sunday—it will all depend on your particular customer usage patterns.

Although your workforce may hate you for it, it would definitely be a very wise choice to monitor your usage patterns and schedule your downtimes accordingly. It may even become necessary to alter these routine maintenance windows as your customer base grows. If anything is certain, it is that what is acceptable today, is not always acceptable tomorrow, and if you would like to keep your customer satisfaction high, it might be a good idea to remain flexible.

When you are performing these server upgrades, it may seem possible to swap out hardware during normal hours, especially if the hardware is hot swappable. I would strongly warrant against this for a couple of reasons.

First, some of the hardware that claims it is hot swappable might not conform exactly to your belief of hot swappable. In fact, I have seen many cases where a supposedly hot-swappable component was replaced, and although the system did not crash, the newly installed component was not recognized until the system was power cycled. You can imagine that this could put you in a bad situation, and you may have to reboot the system just to restore functionality. If this was accomplished during normal business hours, it could prove to be disastrous. I have also seen cases where the removal of a component did in fact stop the system, and caused a complete loss of functionality.

For instance, imagine if you installed a replacement part that was also broken. There is a possibility that this might bring the server down, or even harm a different component installed in the system. I am not saying that it is impossible to hot-swap hardware, though. I have successfully installed components in production servers without a single glitch many times before. However, I do not think that the benefit outweighs the possible consequences in most situations. If, on the other hand, your server has crashed because of a particular component, and you possess the ability to replace the component, by all means, do so. If you do not have functionality, there is very little more that could go wrong!

Performing software upgrades will probably cause the most problems and headaches. For one, it is nearly impossible to replace or upgrade an application on the same server without bringing the system down for at least some period of time. Sometimes it will be possible to install a new copy of a particular program

without removing the first; however, at some point you will need to stop the old program and start the new. During this time period, the application will not be available for use. Most upgrades, however, will not allow this, and will need to be performed after the first program has been stopped and there are no users accessing the application. This can lead to a very long outage, depending on the type of upgrade.

In addition to this problem, there may be times where a software upgrade will first require a hardware upgrade. This can compound the amount of time the server is out of commission, especially if there are problems with either upgrade. As you can see, there is a lot of uncertainty when upgrading systems.

Depending on your size and requirements, it may prove useful to test a particular upgrade prior to performing it in your production network. If time and money exists, test the upgrade and work out any problems ahead of time so that there are minimal surprises and the actual downtime can be better estimated when these upgrades are performed on your production servers.

Whenever performing an upgrade, always incorporate a back-out plan. In some cases, it may even be necessary to provide several back-out plans at every stage of a complicated upgrade. The reason for this is obvious: simply put, you do not want to be stuck with a nonfunctional system if the upgrade fails. Whether you are performing a hardware or software upgrade, make sure that you always have at least one method of restoring the system back to its original state. If this is a hardware upgrade, it might mean to not destroy or lose any components that have been removed from the system, or to be careful when installing new items and not force a component into place because you are pressed for time.

If, on the other hand, this is a software upgrade, probably the best method to restore the server is to perform an entire backup of all the applications and data on the server. In this case, if anything occurs, it is simple enough to restore the applications or entire system from tape.

Whether the system fails, and you have a back-out plan or not, always remain calm and try not to do anything outrageous that has the possibility of damaging the system further. If the contrary occurs, it might be a good time to ask someone for assistance.

## Break/Fix

Regardless of the amount of maintenance and forethought put into a server, some piece will eventually break. It could be a small inconsequential item that breaks, such as a plastic cover on the server, but chances are it will be something much more important that will cause your server to stop functioning.

Hopefully, when this occurs, you will have built a successful fault-tolerant solution, and another server will begin to function in the failed server's place. In the best designs, it may be difficult to even notice the problem without monitoring tools. Unfortunately, this is usually not the case, and the faulty server could even cause a particular application or service to cease functioning altogether.

Whether it is inconsequential or not, it will still be important to replace the faulty part and return the system to a fully functional state. It is never a good idea to let even minor problems remain, since they can easily add up and make for an unreliable server. In fact, even the most trivial problem could turn out to be serious if given enough time.

For instance, if the broken plastic cover mentioned previously is allowed to remain unfixed, an overwhelming amount of dust could build up inside the server, and cause fans and other components to become faulty. There is almost no escape, even from the most insignificant problems.

Since a component in your server is bound to break at some point, we recommend keeping some replacement parts on hand whenever money allows. Depending on your service level agreement, it might even make sense to keep an entire server available solely as a replacement for a defective device.

Sometimes, this may not be possible, and if that is the case, I hope that you have built an excellent fault-tolerant solution, or at least have an excellent support and warranty contract that will allow for replacement in a matter of hours.

Otherwise, you will probably spend hours fielding complaints, and giving your customers refunds. Sometimes it is a good thing to listen to the “what-if monster” and to heed the warnings.

## System Monitoring

In order to catch problems before they arise, you will need to perform some type of system monitoring. This monitoring might be as simple as a small script that will check to see if a server is alive, or it might be more complex and incorporate artificial intelligence that is capable of diagnosing the actual problem and even suggesting a possible solution.

The actual monitoring tool that is used depends on your level of expectation from the system. If you do not mind a server failure, and it will not cause many problems, you might decide to go with a very simplistic solution, maybe even one that only pings devices on the network and sends e-mail when a particular device is unreachable.

This type of solution does not give you many features, and will not inform you of any potential problems. Instead, it will only let you know when a device is unreachable whether it has failed or not.

If, for instance, a switch with 100 devices attached were to stop functioning, your monitoring software might send you over 100 notifications, or at least one for every machine that is unreachable. On the other hand, if you rely heavily on every server in your network, and want to have immediate in-depth notification when even a minor problem occurs, you will need to look into alternatives that offer advanced features and support monitoring of specific applications.

The drawback is that the good packages tend to be expensive, and are sometimes very complex to install and configure. The number of alternatives is endless, though, with packages that range in price from free to systems that will cost hundreds of thousands of dollars to install and implement.

With so many choices, it is important to know exactly what features you are looking for, and exactly what those functions are worth to you in dollars. The good news is that most of the major network monitoring packages such as HP Openview, Veritas NerveCenter, and What's Up Gold, offer many server monitoring tools that are capable of monitoring the vital signs of a system as well as the applications running on them.

The beauty of this is that you may already own a system that is capable of incorporating server monitoring right out of the box or with an additional module. In addition, the same monitoring software can be used throughout your network to monitor all your devices and nodes in the network, which can both simplify and centralize your management and capabilities.

## Summary

Servers should be one of the most important concerns for your ASP, especially since it is these devices on which all of your service offerings depend. In this chapter, we learned about some of the hardware components, such as the CPU, memory, and mass-storage devices that comprise a server and allow it to perform complex instructions and functions. We also discussed network interface cards (NICs), which provide a server with a connection to the network and offer other advancements such as link redundancy, fault tolerance, and aggregation.

After reading through the hardware section of this chapter, you should have come away with a deeper understanding of how a server operates, and some of the pitfalls to look for when planning and ultimately purchasing your servers.

This discussion led us to the topic of operating systems and software applications that can run on particular servers, as well as the many advantages and disadvantages offered by each. We discussed the importance of server and application redundancy, and exactly why you should look at designing these features in all of your devices and throughout every aspect of your network.

We also explained some of the considerations you should have when connecting your servers to the network. This includes some of the services available, such as network storage, data backup and recovery, virus scanning, and thin client. These should help you plan the overall design of your network, and how your servers will interact with one another.

Finally, we discussed maintenance concerns. These will become a very key element to your ASP. Initially, you will be rushing to install new hardware and software, and will be concerned primarily with the design of your network. However, once that is complete, you will need to begin the task of maintaining your network.

Our aim was to give you a better idea and understanding as to how you should maintain your systems without causing huge issues, and customer complaints, as well as a plan to get you started quickly and effectively. All of these topics combined should give you a basic understanding of what it takes to design, purchase, and install a working server from start to finish.

This chapter exposed many topics, and quite a lot of ground was covered. Some of this information is very basic and should be understood by all, while other pieces are more complex and will probably require more research to fully understand the intricacies of the technology. After reading this chapter, you should come away with a plan that will assist you in the design and implementation of servers in your network, and how to make well-planned, thoughtful decisions before purchasing complex solutions.

At times, it may have appeared as if the information was a little repetitive, or unnecessary. However, it is very important to understand the whole picture in order to ensure that your business goals can be met by a particular technology or product offering. It is equally important to look at the fine details, to allow you to build a system that meets your expectations on all levels.

If you take one thing away from this chapter, it should be to plan your future server growth carefully. You should put much thought into your design, since you will probably not have the luxury of scrapping your equipment and starting over. You will want to choose solutions that are both cost effective and scalable.

As we discussed, you should look for redundancy and fault tolerance, or at least understand how they operate and how they can be added to your design at a later date. You should also always be wary and conscious of the pitfalls, and remain flexible and open-minded. There are many options out there, and the sky is really the limit; however, in many cases, the cost can easily outweigh the features, and claims can sometimes hold many caveats.

An ASP relies heavily on its servers. There are so many possibilities with servers that it is very difficult to cover them in a single chapter. If you did not understand something fully, or want to learn about a particular technology a little more in depth, we recommend researching the technology by contacting the vendor, or calling a consulting company that has done these types of installations in the past. If possible, use several sources for your information.

If there is one thing that is certain, it is that a claim is only a claim until it has been proven. Because of this, it may take several sources to ensure that a particular technology or manufacturer's product offering is really capable of its claim.

## Solutions Fast Track

### Implementation, Where to Begin

- ☑ At the heart of an ISP/ASP are the server base and the application software packages. If they do not function efficiently, the ASP will not run effectively.
- ☑ Today, there are only two basic types of microprocessors available for computers: Complex Instruction Set Computers (CISC), and Reduced Instruction Set Computers (RISC).
- ☑ SMP is an architecture that provides better performance by using multiple processors in the same server.

- ☑ Fibre Channel has been introduced as a replacement for the SCSI architecture. Fibre Channel provides a method for transmitting data between computers at a rate of 100 Mbps, and scales up to 1 Gigabit per second (Gbps).
- ☑ Link aggregation allows a single server to use two or more installed network interface cards (NICs) to aggregate bandwidth across several links.

## Software Solutions for Your ASP

- ☑ *System software* describes software packages that provide the basis for all other applications that are run on a computer.
- ☑ Unix is not a proprietary operating system, and the source code has been available to the public since its inception. Currently, the leading Unix environment is Solaris from Sun Microsystems.
- ☑ Windows 2000 Advanced Server offers all of the features available in the standard version, but includes more reliability and scalability, as well as additional features for applications that require a higher level of scalability.
- ☑ Novell offers a powerful network operating system called NetWare. This operating system was originally designed for use in small to enterprise businesses and networks, and typically used a protocol stack called Internet Packet eXchange (IPX).

## Application Software Types

- ☑ *Applications* is the term used to describe a group of programs or code designed to perform a specific function directly for users or other application packages.
- ☑ *Internet Information Server (IIS)* is a scalable Web server offering from Microsoft Corporation that runs under the Windows family of operating systems.
- ☑ Apache HTTP Server is an open-source software package that is organized by the Apache Software Foundation.
- ☑ A database can be defined as a collection of data that is organized for management and access.



- ☑ *Middleware* can be considered the “glue” that holds applications together. It is a general term for any computer application whose purpose is to combine or mediate between two applications in order to allow them to share data between them.

## Network Service Considerations

- ☑ *Network storage* defines the ability to store information on a remote system connected over a network.
- ☑ NFS was first released in 1984 by Sun Microsystems Corporation.
- ☑ Today, many systems use NFS to connect servers to centralized storage. Since NFS was designed on the Unix platform, it has remained a Unix tool, for the most part. It is possible to find NFS servers and clients that run under other operating systems, such as Windows, but they are not very desirable since they are not native to the particular operating system.

## Data Backups and How They Can Affect You

- ☑ Although hardware platforms have become more reliable over the years, the fact still remains that your data is stored on what is essentially a mechanical device; a disk that rotates at very high speeds with another bit of metal called a head that floats left and right across the surface of the disk many times a second.
- ☑ You will most likely use a third-party backup program as opposed to the generic ones that sometimes come with your operating system, or storage devices. Some of the products that you will run across such as ARCserve, Veritas Backup Exec, UltraBac, or NovaStor, will allow advanced scheduling with various levels of flexibility.
- ☑ One of the defining factors between backup systems is how tapes are rotated and what files get backed up to which tape. Each rotation method has different advantages that can be applied to systems and provide for different results.

## Virus Scanning Suggestions

- ☑ A virus can halt your servers, and can even remove data from your hard disks. What's worse is that it can spread to incorporate the computers throughout your entire network and into your client's networks, infecting every server along the way and leaving mass data destruction in its wake.
- ☑ When using an Internet Gateway product, make sure that you have a system that will allow you to queue incoming e-mail messages. If mail is received faster than it can be processed by an Internet gateway, it could start dropping or bouncing messages unless you have software that allows incoming messages to be queued.

## Thin Client Solutions

- ☑ One of the primary focuses for an ASP is to ensure the delivery of its products or services to each client's desktop.
- ☑ Independent Computing Architecture (ICA) allows the delivery of an application from a centralized server to any end-user desktop, regardless of the operating system or platform.

## Maintenance and Support Issues

- ☑ Eventually, every piece of hardware and software operated by your company will need an upgrade of some sort.
- ☑ When you consider that you might be performing hardware upgrades as well as software upgrades, and that one upgrade might cause another, it just does not make sense to even attempt to upgrade the servers all at once.
- ☑ Whenever performing an upgrade, always incorporate a back-out plan. In some cases, it may even be necessary to provide several back-out plans at every stage of a complicated upgrade.
- ☑ In order to catch problems before they arise, you will need to perform some type of system monitoring.

## Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the “Ask the Author” form.

**Q:** What is Symmetric Multiprocessing, and how does it benefit a server?

**A:** Symmetric Multiprocessing is an architecture that provides better server performance by allowing multiple processors that are installed in the same server to work in parallel.

**Q:** If my server came with EIDE hard drives installed, can I install a SCSI hard drive instead?

**A:** It depends. EIDE and SCSI are two separate technologies that require different controllers installed in a computer in order to function. If your system came with an EIDE hard drive, it also has an EIDE controller installed. It may be possible to add a SCSI controller to the system to allow you to install SCSI hard drives, but it will be an additional cost. This is the case for all of the different mass-storage technologies, such as Fibre Channel, and ESCON.

**Q:** What is the difference between single-mode and multimode fiber?

**A:** The difference between these two technologies lies in how light is transmitted over the cable. In single mode, light is transmitted straight through the core of the cable, while multimode transmits light into the core at different angles.

**Q:** What is the advantage of single-mode fiber?

**A:** Single mode fiber will allow for longer distances between segments.

**Q:** What is link aggregation?

**A:** Link aggregation allows a device to use multiple network interfaces on the same network in order to provide additional bandwidth. For instance, if you have four 100-Mbps network interface cards installed in a server, it is possible to aggregate these for a total of 100 Mbps (half duplex).

**Q:** What are the three types of software required to run a computer?

**A:** The three main types of software are system applications, which comprise software such as operating systems and system drivers; application software, which consists of programs such as databases, Web browsers, and e-mail; and middleware, which helps to tie applications together.

**Q:** Is there any way to provide redundancy in my servers?

**A:** Yes. There are numerous ways. First, many servers offer redundant components such as spare power supplies. These can help keep a server remain operational even if a particular component were to become faulty. In addition to this, there are several software packages available that will allow you to cluster servers. This essentially allows multiple machines to serve as backup servers in addition to their normal functionality.

**Q:** What is NFS?

**A:** NFS stands for Network File System, and is a client/server application that allows for file and data sharing. It is usually run in a Unix environment, but has also been used with other operating systems.

**Q:** What is a computer virus?

**A:** A computer virus is a program that is usually intended to do harm to your system. While some viruses are merely pranks, others have been known to cause serious damage and sometimes removal of data, and can be spread among devices. No operating system is immune to viruses.

**Q:** What is thin client?

**A:** Thin client provides a way of accessing data and applications on a server. Instead of using the standard client/server model, it uses an approach that forces the server to perform all of the processing, leaving the client as a terminal that merely acts as a user interface to the actual server.



## Performance Enhancement Technologies

### Solutions in this chapter:

- Web Caching and How It Works
- Deployment Models for Caching
- Load Balancing in Your Infrastructure
- Load Balancing Solutions from F5
- Cisco Systems' LocalDirector
- Foundry Networks' ServerIron
- Content Delivery Networks
- CDN Solutions from Various Vendors
  
- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

# Introduction

The growth of the World Wide Web has greatly stressed the performance of the Internet within the last few years. The Web is now a major center for business transactions, with an increasing proportion of bandwidth taking the form of e-commerce. The projected growth of the Internet economy is enormous. The importance of the Web as a stimulant for this economic blossoming means that it must become more reliable and predictable, so that it can be an acceptable medium for doing business.

The Web is essentially inefficient, as every user is seeking to view specific content and must obtain it directly from a server that is the point of origin for that content. It is not cost effective or feasible to have a dedicated, point-to-point trunk allocated to users—without this, congestion is inevitable.

Problems that contribute to user frustration include:

- Slow connection speeds
- Unpredictable performance
- Limitations in available bandwidth
- Overwhelmed Web sites

The Internet is constantly being built out to handle the capacity of the growing load. In the foreseeable future, the build-out will lag behind demand. Simply increasing bandwidth by building up the network with bigger bandwidth pipes cannot address all of the Quality of Service (QoS) issues that will become involved in the scaling and evolution of the Internet.

For purposes of this discussion, QoS means a high-quality user experience that can be measured in low latency for downloads and generally faster download times. Adding bandwidth may improve speed, but it does not remove the latency or delay that is inherent within all networks. Moreover, adding bandwidth at one point may only change the location of a bottleneck.

As an Internet service provider (ISP) or an application service provider (ASP), your Web site and infrastructure will generally consist of distributed areas that can provide network monitoring, Web content, and application services that will assist in improving response times. Several technologies can be used to enhance the performance of your Web site, such as caching, content routing, and load balancing.

## What Is Web Caching?

As part of the caching solution, there are suites of effective technologies such as Web caching, which moves and saves Web content as close to the end users as possible. With this method, both static and dynamic Web pages can be cached for later usage.

Static Web pages are usually cached in RAM so that end users can access them quickly. Dynamic Web pages can also be cached, but they require the use of predictive algorithms that allow dynamic pages to be generated before end users request them. Caching helps to make more bandwidth available by using existing pipes more efficiently. This not only improves the QoS for the user, but also gives service providers substantial savings, and allows room for growth.

## What Is Load Balancing?

Load balancing is the one of the most commonly used techniques to improve response time of content on the Internet. Several Web servers are configured to share the load of the processes. A side benefit of load balancing is that it offers fault tolerance, due to the nature of using multiple servers.

## What Is Content Routing?

Content routing can be used to handle mission-critical Web sites, by providing fast response times. Web pages for these sites are replicated to diverse data centers at different geographical locations. This permits end users to access these pages quickly from multiple sources. This technology has enabled one of the newest and possibly most powerful technologies for the future of the Internet: Content Delivery Networks (CDN). This combines traditional routing and switching intelligence with content-aware technology (at a packet level), which is located at service provider distribution areas or enterprise data centers.

## Web Caching and How It Works

Bandwidth shortage is only one of the obstacles that contribute to the slow response time of Web-based content. Building up bandwidth connection will not necessarily solve network latency or slow Web server access. Web caching was created in order to address these problems. The intent of caching is to move Web content as close to the end users as possible for quick access to improve the customers' satisfaction levels, and gives your ASP the competitive advantage.



## What Is Data Caching?

As you have probably seen, *data caching* is a highly efficient technology that is already implemented in many areas of your network as well as in the Enterprise networks.

Data caching is generally used in conjunction with other technologies in order to speed up other applications. These are usually hardware devices that can cache frequently used data and instructions in order to handle bandwidth and resources in a more proficient manner. For example, data that is frequently used by a computer's Central Processing Unit (CPU) will normally be stored in local Random Access Memory (RAM). RAM is very fast memory and is sometimes right on the CPU itself. This high-speed memory helps to reduce the need for the CPU to read data from a disk drive (which is usually much slower as it is mechanical in nature rather than circuitry based, like RAM).

This is not a limited technology, as Web browsers are also designed to cache a limited amount of content locally on a user's machine. What this does is allow for the selection of **Back** or **Previous page** on a browser toolbar which results in near-instantaneous retrieval. But this is not true for Web caching. True Web caching uses a server or some specialized device that is placed close to users in the form of a network cache. This reduces the number of router and switch hops that are necessary to retrieve Web content from remote sites. For instance, an audience doesn't need to travel to Hollywood to see a movie; instead movies are sent to local theaters where people can go to see them. This is intrinsically more efficient and allows for a higher user experience.

Normally, Web caching is separated into two distinct models. There is the "Edge-Services" Model, where a business would subscribe to a third-party service provider to have their content cached and served from. This model has some serious disadvantages for some of the customers:

- The service provider doesn't own or control the infrastructure.
- Many times, the more frequently used sites are not always the ones that are cached. This can lead to poorer performance, which can disappoint the end users.

There is also the "Open" Model which is supported by several of the major caching vendors (Intel and Cisco Systems caching appliances come to mind) in which service providers install their own caching equipment. This allows them the ability to offer data caching as a value-added service to their clients. Some of the advantages of this model include:

- The service provider is able to invest in its own infrastructure.
- There is additional revenue that can be realized by directly offering this at the service provider level.
- The system is able to automatically cache the Web sites that users most often access.

## The Benefits of Data Caching?

Who really benefits from the implementation of Web caching? Everyone, this allows for greater QoS for end users, enterprises, service providers, and content providers. All of these models benefit from the implementation of data caching engines.

The group that benefits the most is the end users. These are the people who drive the Internet economy. Web caching is able to provide diverse benefits for end users that can manifest themselves through an enhanced Internet experience. This creates the perception that customers are getting better value for their monthly service fees.

Data caching also benefits enterprise users, especially in large environments that have comparatively little bandwidth. By providing a local cache for Web content, these larger companies are able to monitor how much bandwidth is necessary to meet employee requirements for their network. This will also help companies initiate policies for access that can limit employee usage of the Web to corporate activities.

For ISPs, data and Web caching have several important advantages:

- Caching can reduce overall bandwidth usage by eliminating redundant requests for popular documents and Web sites.
- In the Enterprise, your client may be able to reduce leased line expenses. A data and Web cache that is able to successfully serve an average percentage of user requests will realize that the amount of outbound bandwidth that is normally required can be reduced by up to 40 percent. As you can see, this can allow for significant savings, or may allow the company to add more users with the current network.
- With the use of caching, you can provide better QoS. This will lead directly to higher customer satisfaction and therefore minimize customer turnover or churn. So there is more money that can be spent in acquiring new customers, while still keeping your current customers happy.

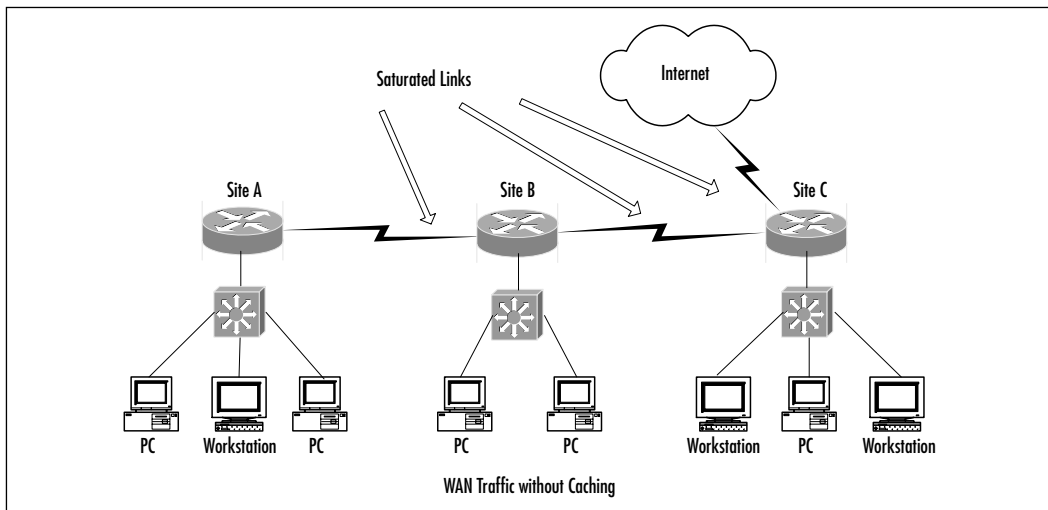
- A Web caching solution provides value-added services that can boost an ISP's profitability.

People that model their business on the Content Providers method can benefit from elevated site availability. This allows for better-perceived user experience that has not only fewer, but shorter delays. This creates a competitive advantage for companies that have these data caches, over those companies that are not cached.

Within the last year, there was a study that was conducted to measure what the average time for Web content to appear before users became antsy. This study indicated that a delay of only five to eight seconds while waiting for a Web page is enough to frustrate the average user into retrying or leaving a site. Within the last year, due to people getting used to faster access, either through Digital Subscriber Lines (DSL), or cable, or through faster connections that they experience within their working environment. By deploying Web caching, this frustration can be minimized or even prevented. As you can see, that from an overall business point of view for service providers and online businesses, you would want your users to be able to visit more sites so that they can do more purchasing of products because content can be delivered faster.

In Figure 4.1, the amount of bandwidth that is required for trips across the backbone is significantly greater in a network that is noncached. With content caching configured, a large portion of the requests can be fulfilled using only local bandwidth.

**Figure 4.1** A Noncached Infrastructure



## What Happens With and Without a Solution in Place

If there isn't a caching solution in place, requests for content delivered from the destination site must repeatedly take the same trip presumably across the Internet or at least through your provisioned bandwidth. The following steps are required to perform a trip from the requesting computer to the destination computer that contains content, and back again to the source machine:

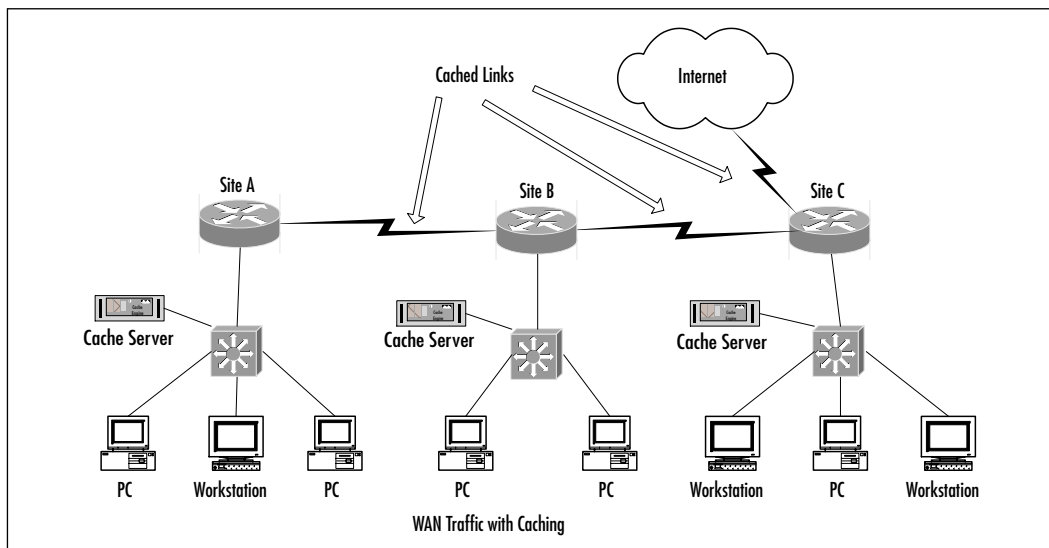
1. A user's Web browser sends a request for a uniform resource locator (URL) that points to a specific Web document that is stored on a unique server on the Internet. Usually this is done from a DNS lookup.
2. The request will go to a DNS and will then be cross-referenced with an IP address. The request is then routed through the TCP/IP network transport.
3. Content requested from these remote servers (also known as a HyperText Transfer Protocol (HTTP) server) may be a static HyperText Mark-up Language (HTML) page with links to additional files, and can include graphics. The content can also be a dynamically created page that is generated from a search engine, a database query, or a Web application.
4. The HTTP server then returns the requested content to the client's Web browser one file at a time. A dynamically created page often has static components that can be combined with the dynamic content to create the final document.
5. If there is no content caching server in place, the next user who requests the same document will need to send a completely new request across the Internet to the Web server, so that it can receive the content by return trip. Thus taking up all of the resources that were used the first time.

The process becomes far more efficient when content caching is enabled, because frequently accessed content does not have to make the long trip from the client to the remote Web server repeatedly (Figure 4.2).

- If the requested document is stored on a cache server that is located within the user's corporate Local Area Network (LAN), at the company's service provider, or some other Network Access Point (NAP) or Point of Presence (POP) that is located closer to the users than the to the remote Web servers, there will be a noticeable savings on bandwidth.

- If the requested document has recently been stored on the cache servers, the servers will check to make sure that they have the most current content (this can also be called fresh). This is done to make sure that a user does not receive an old (stale) or outdated object. There is the ability in some caching devices to set freshness parameters, these can be preconfigured by content providers. Most of the time these are turned on by default when you are configuring and installing these devices.
- If the content is current, then the transaction can be considered a cache “hit.” This allows the request to be immediately fulfilled from the local cache server.
- If the content is old and needs to be refreshed, the cache server can be configured to retrieve updated files from the Internet. This will ensure that the device has the most current information so that it can send them to the client, as well as keeping a fresh copy for itself.
- The more frequently a server can cache user requests, the higher the hit rate and the better the performance for the users will be.

**Figure 4.2** A Cached Infrastructure



The process for caching is similar to the process for File Transfer Protocol (FTP) file transfers. The FTP server will handle each request for a file that is presented from a client’s application. Bottlenecks are a substantial problem with FTP files, because the size of a typical FTP file is larger than a typical Web-based file.

There are many applications such as streaming audio and video that are also examples of Internet applications that can greatly benefit by caching content. Problems with latency through the Internet can cause video that is “jittered” and delayed or distorted audio. By implementing QoS, you are able to better use bandwidth to solve these problems.

## How to Reduce Bandwidth Usage

Data caching reduces the upstream bandwidth that an ISP must provide to meet user content requirements. A cache only needs to pass user requests on to the Internet if it isn’t able to service them locally. The greater the number of requests that can be handled from cache, the less bandwidth that is used to reach distant content servers.

Through this traffic reduction, service providers can achieve significant savings in resources. It has been estimated that 30 percent of an ISP’s operating costs are recurring telecommunications charges. There will always be external traffic, as updates must be performed for freshness. By using caching, though, bandwidth utilization can be much more efficient. Caching is still beneficial when retrieving dynamic documents, because these pages do have some static components that can be served from a cache appliance.

Based on the distribution of traffic and the scalability of the cache, there can be a savings of up to 40 percent (source: Patricia Seybold Group, 1999) of user HTTP requests. This occurs as the traffic is removed from the network and fulfilled from the cache server. This enables networks to be far more efficient, and allows better service at a lower cost.

In order to make your cache truly efficient, you will want to cache as much Web content as possible within the boundaries of an ISP while using small to average amounts of upstream bandwidth so that you can give your clients what they require without creating “black holes” for bandwidth or losing your ROI.

In Figure 4.3, Layer-4 switches and routers can direct requests for data (HTTP, NNTP, etc.) to the cache server while sending other requests to the Internet.

## Key Requirements for a Caching Solution

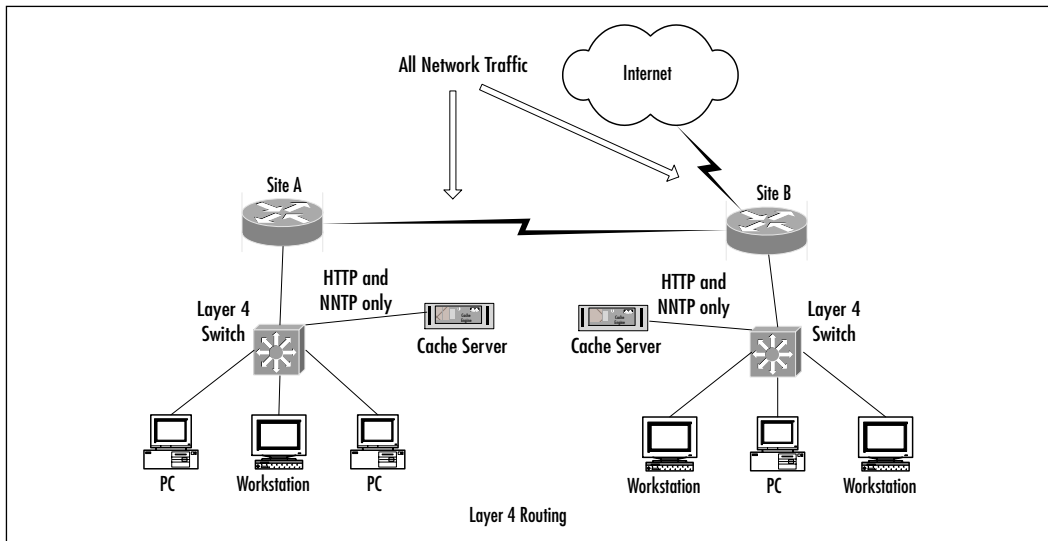
There are several requirements of a caching solution that can allow it to provide optimized performance. Two of the most important sides to cache performance are:

- **Operational capacity** This is handled by the design and deployment of the cache server. In conjunction with raw cache capacity, production

issues include how the server performs with multiple threads and tasks, and how well it executes load balancing with multiple cache servers that are located within the cluster.

- **The ability to be responsive to client requests** This ability can be determined by which technique the cache server uses to maximize its hit rate, including the structure of hierarchies (cache hierarchies are discussed later in the chapter) and the optimization of content. Cache hit rate is a combination of many things, such as cache size and the load on the cache.

**Figure 4.3** Layer-4 Routing



There are many ways that cache servers can be tweaked to improve the capacity and responsiveness in multiple ways. Some of the more common optimizations and improvements include:

- Processing queues for the objects that make up a document
- Determining whether a requested object is cached
- Delivering the requested object to the browser when it is not in cache
- The handling of total throughput based on incoming requests and outgoing data

The performance ultimately depends on how well these potential issues were understood and implemented by those who installed the cache server and created the software.

Scalability is another key prerequisite that a cache must be able to address. Remember that the efficiency of caching increases as the traffic served by the cache increases. This means that the larger the implementation, the more valuable the result. There are also ways to enhance and support very large caches, such as cache server clustering or load balancing when necessary. This helps you to design and implement robust caching solutions that are able to support even the largest customers.

Cache servers support a variety of protocols. Network caching can assist in the content that is delivered over HTTP, Network News Transfer Protocol (NNTP), FTP, and others. All of these protocols are distinguished by having at least some static content.

Manageability is critical for any caching implementation. Cache management must include the ability to easily install and maintain cache servers, and allow access to usage and traffic statistics that the servers provide. You will want to be able to manage groups of cache servers, some of which may be distributed in geographically dispersed areas, from a single point of control.

Graphical user interface (GUI) management interfaces are becoming increasingly common as the typical way to manage these distributed systems. These interfaces provide functionality for configuration, administering security, creating filters, updating the cache, controlling the system, and gathering statistics from system logs.

The solution should provide high reliability and availability. Although the nature of caching has a measure of fault tolerance built in due to replication, the solution must feature first-rate software and a highly dependable platform if it is to be an integral part of the network infrastructure. Configuring fail-over and clustering also contributes to reliability and availability.

Finally, hardware platforms and software packages must be integrated and tweaked to achieve the full benefit of efficiency and caching performance.

## Deployment Models for Data Caching

There are three commonly accepted models for implementing Web cache architecture. The method you choose will depend on where the cache is implemented and the nature of the traffic.



## Forward Proxy

A forward proxy cache is defined by its reactive nature. In the forward proxy cache configuration, a client's requests go through the cache on the way to the destination Web server. If the local cache contains the requested document, it will serve the content directly. If the engine does not have the content, the server will then act as a proxy, by retrieving the content from the external source server on the client's behalf.

## Transparent Caching

Forward proxy caches can also be configured as either transparent or nontransparent. A transparent cache resides in the flow of the network and is invisible to a client's browser. Clients realize the benefits of caching without reconfiguring the browsers. For many service providers and enterprise backbone operations, a transparent configuration is the only way to go, because it helps to minimize administrative and support issues.

The more popular implementation (especially among Enterprises) is to use a switch that is capable of using Layer 4 to connect the cache servers to the Internet (Figure 4.3). These switches can inspect data traffic and make decisions above the IP level. As an example, a switch can direct HTTP (or any other) traffic to the cache servers and send the rest of the traffic to the Internet.

The switch may also send requests to specific nodes within a cache server cluster, for load-balancing purposes. Implementing a pair of Layer-4 enabled switches with multiple cache servers will also allow for redundancy and failover protection.

## Reverse Proxy

There is the ability to create a cache that can also be implemented as a local Web server. This will help to accelerate slower content that needs to be accessed from slower Web servers. How this is handled is through the implementation of a reverse proxy. The documents that are stored locally in cache are able to be served at a very high speed, while documents not locally cached, such as dynamic content or other short-term objects, are requested from the remote Web servers. You will most often see this model deployed to optimize the performance of a Web server farm. The caching system is placed in front of one or more Web servers, capturing client requests and acting as a proxy so that it can fulfill requests in a timelier manner.

The great thing about reverse cache servers is that they can be deployed throughout the network to create a distributed site of hosted content. This model is commonly referred to as site replication. There is also the additional performance enhancement for the clients and providers. These enhancements are created through the inclusion of load balancing, the ability to offer peak-demand availability insurance, and the ability to provide dynamic mirroring for high availability.

## Cache Locations and Placement

The following characteristics will help to identify the most ideal cache deployment points in the network. There are three types of location characteristics to keep in mind:

- **Bottle Necks or Choke Points** These are traffic convergence points where a large majority of network traffic must pass and would be visible to a cache server. This assists the cache in handling more requests and being able to store more content than if it were located in a remote area that is easily bypassed.
- **High Traffic Load or Saturation Areas** These areas are characterized by high traffic throughput that would allow for higher cache utilization. This means that the more the cache is accessed, the greater the benefit will be on bandwidth.
- **Economic potential** There are also points on the network where users will benefit from high cache access rates, while also reducing upstream bandwidth requirements. In implementing cache engines at these points, you will provide both QoS benefits and an economically efficient link for the access provider.

Many of these characteristics are already found throughout major Internet switching locations, dial-up access points, or corporate gateways. Applications for this technology include standard dial-up access, POP, NAPs and network exchanges, “last mile” acceleration, Web hosting, and more. Caching is also used as a more efficient means of updating information stores for online news services and Web sites.

## Cache Hierarchies

Eventually, there will be a time when the information that you request will not be stored in cache (this is considered a cache *miss*). When this occurs, the cache

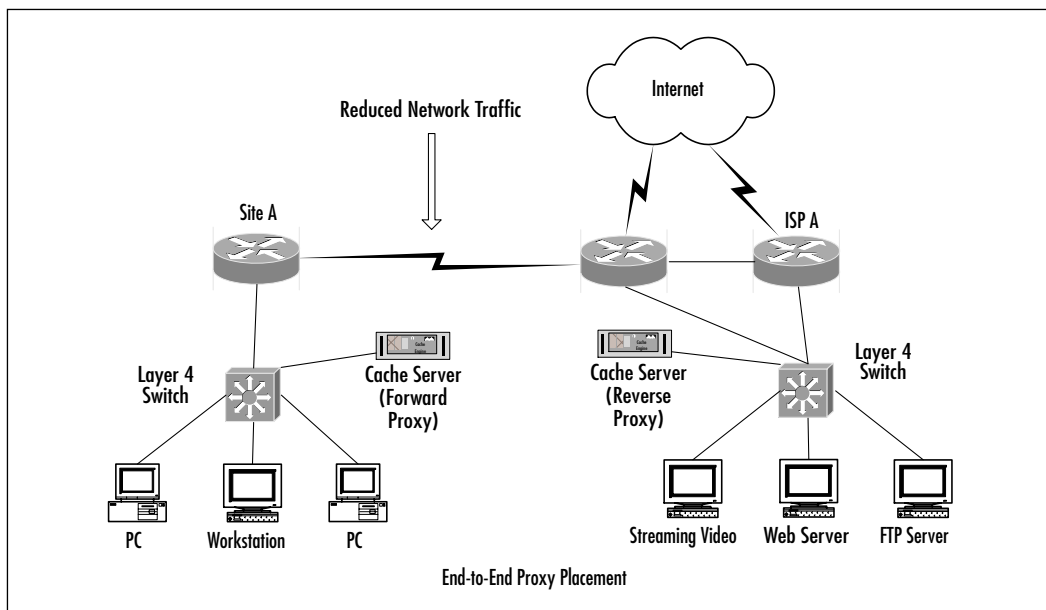
server usually forwards the request to a distant origin server. However, if the cache server was capable of checking with another closer cache server instead, the process could possibly be much faster. This is the concept behind cache hierarchies.

It is feasible to create geographical caches within an organization; for example, a server or cluster could handle certain sectors of a company or limited geographical area, and link them to other larger caches (also known as a *parent* cache) that define larger groups or areas. If a user's local cache does not have the requested content, it will then forward the request for the content to the parent cache. This still provides much faster service than contacting the external destination server, and helps to conserve bandwidth. With multilevel hierarchies configured, a company is able to give cache servers a series of larger and larger cache pools to query if the request attempt misses.

You could also combine capabilities such as site replication and a hierarchical caching structure to create a highly efficient distributed network for Web hosting over a diverse geographical area.

Figure 4.4 shows that cache servers can be placed at a POP to service requests closer to the user. These can be implemented at aggregation points, which are located on the edge of the Internet to reduce bandwidth requirements, or in front of a Web server farm to reduce load on content servers.

**Figure 4.4** An End-to-End Proxy Deployment



## What Are Cache Appliances?

A cache appliance (this can also be called a *thin server*) can be defined as a device that offers a limited number of dedicated functions, and is able to deliver those functions more effectively than a multipurpose device can. By specializing in particular areas, this equipment usually is able to provide features that are more robust, enhanced stability and flexibility, and easier implementation and support.

### Cost Effectiveness

A cache appliance is usually some form of integrated hardware and software that has been designed to provide high-end, carrier-class caching. Some of the capabilities for these appliances include:

- **Ease of installation and management** This is usually accomplished when you are able to configure multiple nodes simultaneously and maintain them from a centralized location.
- **Fault tolerance** Fault tolerance helps to maintain network stability.
- **Scalability and flexibility** Nodes can easily be added to a cache cluster when necessary, and updated services can be implemented as the need arises.
- **Performance and speed** These appliances are capable of handling thousands of simultaneous user connections.

### Ease of Installation and Management

As a “solution in a box,” a cache appliance must contain all of the necessary hardware and software that makes an appliance easy to implement and support. Cache appliances usually have configuration wizards and intuitive software configurations that simplify the setup. This ease of configuration represents a significant cost savings, because there only needs to be minimal resource time allocated to integrate the device into the network and support it.

There are also ancillary benefits provided by the installation of an appliance. Since these devices are relatively compact in size, you are able to provide an increased network capacity in limited rack space.

Cache appliances can also help to minimize the cost of administering, maintaining, and operating a large cache system, by offering several centralized management options that can be customized to suit your needs across a wide range of environments. Some of the common features in cache appliances include:

- **Browser-based interface** The management GUI offers a password-protected, single-point administration for the entire cache cluster.
- **Command-line interface** There is usually a command-line interface that allows the administrator to configure the system's network addresses, and to control, configure, and monitor the cache.
- **Simple Network Management Protocol (SNMP) management** Cache appliances usually support Management Information Bases (MIBs) that allow management through SNMP facilities.

You can access performance statistics from the command-line interface or the GUI, which allows you to tweak performance based on your specific needs. Settings that can be modified include:

- Log file formats
- Content routing
- Site or content blacklist filtering (black hole policy)
- Anonymity
- Never-cache, pin-in-cache, revalidated after X amount of time or action
- Store multiple versions of cached objects for user-defined or browser-defined differences in content
- Domain and host-name expansion

## Fault Tolerance

Because of the mission-critical nature of caching, a cache appliance is designed to provide a highly reliable and available service. Since a cache appliance is designed to implement caches at the highest levels of network traffic, such as NAPs and on the backbone, it is easy to scale.

You can achieve a high degree of scalability with cache appliances in three ways:

- Cache hierarchies
- Clustering
- Symmetric multiprocessing (SMP)

Clustering technology combines the resources of multiple machines to increase capacity and processing power. As nodes are inserted into the cluster,

they assimilate with existing nodes to provide additional disk and processing resources. Clustering also offers failover protection, as node failures can be automatically detected, and traffic can then be redistributed among active nodes.

SMP allows multiple threaded processors to provide the performance that is necessary to accommodate growth, and clustering provides scalability as it is able to spread the workload across several machines.

## Scalability and Flexibility

Since an appliance is usually designed for specialized purposes, it will usually offer a high degree of operational flexibility. Appliances can be used in a range of deployment models, either alone or with other enterprise software, and even integrated other caching products. Here are some of the ways you can implement a normal cache appliance:

- Forward proxy.
- Reverse proxy.
- Transparent caching.
- Nontransparent caching.
- Part of an HTTP cache hierarchy.
- NNTP news cache: The appliance will cache frequently accessed news articles and can receive news feeds for designated newsgroups.

In addition, the cache appliance usually offers a broad range of support for content and interoperability protocols, such as:

- HTTP versions 0.9 through 1.1
- FTP
- NNTP
- Secure Socket Layer (SSL) encryption

## Performance and Speed

Performance is dependent on capacity, and includes how well the server can use multiple threads of execution, and maintain the ability to respond quickly to user requests. Cache appliances are designed to offer top-notch functionality across an extensive range of load conditions.

By implementing multithreading, which allows for the breaking down of large transactions into small efficient tasks, a cache appliance is able to handle thousands of simultaneous connections and maximize CPU utilization. The appliance can respond to multiple requests simultaneously and efficiently even under maximum loads.

## Load Balancing in Your Infrastructure

Load balancing, also called Layer 4–7 switching, occurs when cluster of Web servers are created to handle massive amounts of requests. These server farms also share the workload of processing Web requests through different load-balancing methods.

A load-balancing device is usually implemented to assist in the determination of which server has the least load, so that incoming requests are sent to this server. This can either be accomplished by looking at the transport layer (TCP Layer 4) headers or the application layer (HTTP Layer 7) headers and rewriting them so that the packets can be sent directly to the server.

The load balancers then keep track of the multiple sessions that are associated with these packets so that users can interact with the same server once the request has been processed. When several HTTP sessions are defined in a Web session, it is important to keep the established connections between the same partners.

## Localized Load Balancing

Localized load balancing occurs when the load balancer determines which server should receive new requests. This is usually based on a combination of criteria, such as the history of communications from the client, the load on processor of the server, and the network utilization on the server.

The load balancer will then transfer the traffic between local Web server farms within an autonomous system. Localized load balancing is normally used in a single LAN or a small group of LANs that are directly connected to the load-balancing device.

## Distributed Load Balancing

Distributed load balancing is another method of load balancing for Web servers. This occurs when Web servers route Web traffic across multiple networks. Distributed load balancing sends packets across dispersed networks, which can be located in geographically separate areas from the local server. The purpose of this is for the client request to be processed by the closest (fastest) Web server.

A client request normally tries to access the very first or largest Web server that is identified by its URL hostname (e.g., `www.eXn.com`). The primary server will likely receive all initial client requests, which means that the Web server will have to accept the impact of new incoming traffic.

The site might actually have multiple servers, some which may be geographically closer to the client than the primary server. In order to reduce the latency and improve processing load, the request could be distributed to a closer server.

## Configuring & Implementing...

### What Does Closeness Really Mean to Your Network?

The closeness of a server doesn't necessarily mean that it is the closest geographical server. In the context of this section, closeness is based on several factors, including bandwidth, convergence, and latency that is associated with all networks.

If only it were that easy. With the meshing that is involved with the Internet to keep it stable, available, and redundant, it is difficult to determine the closest server to a client. Because of this, IP packets do not always flow along the same path; the route could even change depending on the network topology of a client's Internet service provider and internal infrastructure.

Distributed load balancing generally uses special protocols between the servers to determine the closeness and the load of remote servers, and allow them to redirect traffic appropriately. As of yet, there is no regulated standard for such a protocol, but several vendors do have their own proprietary standards.

## Comparing Different Load-Balancing Systems

Service interruptions of Web applications can happen in many ways; for example, server and software failure, hardware, operating system and application failure, content failure, error messages, incorrect data, and so on. Heavy traffic loads and network saturation or failure can also limit the availability of your site. Load-balancing systems must be designed to guarantee network availability despite these interruptions.



To evaluate the available products, look for a solution that maintains a balance of quality of service-based availability, can assure continuous operation with little or no downtime, is simple, has consistent management across a wide range of protocols, has robust technical support, and is easy to install.

Having a single product that can offer a good amount of these critical elements can provide tremendous cost savings, and still enhance your users' experience, which will provide significant long-term business value.

Typically, there are varying levels of scalability, availability, and performance found in vendors' products that address this market. In an effort to better define various product offerings, we've put together an overview of the following technology categories:

- Software-only
- Switches
- Routers
- Caching servers
- Clustering
- Hardware-software network appliance

## Software-Only Solutions

Software-only solutions are a category of load balancing in which software is installed directly onto the servers. You can perform granular management of servers, such as analyzing the CPU memory utilization and execute agent-based content management. There will be a cost savings and speedier performance, because traffic doesn't have to traverse an additional device—an important consideration. Some solutions do require more expensive, robust systems to run.

Some software-only load-balancing systems will allow for the synchronization of data between servers within a cluster. This can be useful if servers don't have identical content and are deployed to perform complementary tasks. This requires different servers to work with one another to complete content requests.

While analysis and synchronization of data can be performed and monitored, there will be a degradation of speed and performance within the system. Some vendors extol their ability to do URL parsing. This allows the load balancer to examine the URL that is being requested, and make load-balancing decisions. There is a problem with this, though; performance will be decay when this feature is turned on. This occurs simply because there is a higher drain on the load

balancer to perform load-balancing functions versus content delivery with these more task-intensive features.

OS dependencies are another consideration for software-only solutions; because the software is installed directly on to the servers, businesses are “locked in” to support for specific platforms. There are also security concerns, as this type of solution will expose the server’s real IP addresses directly to the user.

Fault tolerance is compromised, as the software creates a new point of failure on the server that it is intended to protect. There are other issues such as management, system downtime, and cost, as new software must be installed and upgraded on each machine in the network.

## Switches

Switches are able to perform fast load balancing at Layers 2 and 3 in hardware, and are managed by a central processor that can accomplish background tasks such as routing, table, and network management. These solutions create fast balancing of static content, and have high back-plane speed support. In addition, switches have the potential to connect to multiple interface ports simultaneously, thus further optimizing the speed of your links.

This architecture does have several inherent limiting factors. For instance, each packet that requires exception handling at Layer 4 (the transport layer) must be opened and examined to see what port they are destined for. This task uses the switch’s central processor and will create performance decay. In other words, per-frame processing on a central processor will limit the total frame throughput of the device.

You should also keep in mind that these switch solutions often lack functionality such as SSL session ID tracking, user authentication, or application health monitoring, which limits the ability to implement more sophisticated tasks for e-commerce.

Remember that the balancing of packets is only as fast as the uplink. Expandability can only occur through the cascading of network devices, and an additional layer of devices is needed in order to have full redundancy. Additionally, many of today’s switching solutions do not contain practical WAN high availability and load balancing for networks that are in different systems (such as the Internet).

## Routers and Caching Systems

Load-balancing products should be used in a manner that is complementary to routers and caching systems. A load-balancing product, for instance, can offer

additional scalability, availability, and security features well above just the basic routing and caching functionality.

These products can also manage the load balancing between cache servers and routers, which will further enhance system performance. A dynamic configuration, when combined with an appliance-based load-balancing product with routers and cache servers, is one of the best ways to meet user and client demand.

## Clustering

The major manufacturers of clustering systems offer solutions that are excellent for major-scale, mission-critical financial applications—and come with superior performance and a heavy cost. Solutions of this nature offer some expandability, but are of course completely proprietary and represent a long-term commitment and investment.

Applications ideally suited for clustering include those that can be processed on multiple servers, such as banking and other sophisticated host environments. Network appliance-based load-balancing products, for example, differ in that they are optimal for all-inclusive solutions that manage a specific server/application, are plug and play, and maintain an open and flexible architecture to grow over time.

Some lower-end clustering solutions are available (such as Microsoft and Novell), but are limited in the number of servers supported and in functionality. In these scenarios, it's important to look closely at the architectural limitations and tremendous cost of clustering, and determine if it's the right investment for your business. If you do find that these solutions fit your needs, it's important to remember that certain load-balancing products can often complement this approach, offering additional intelligence and reliability to your system.

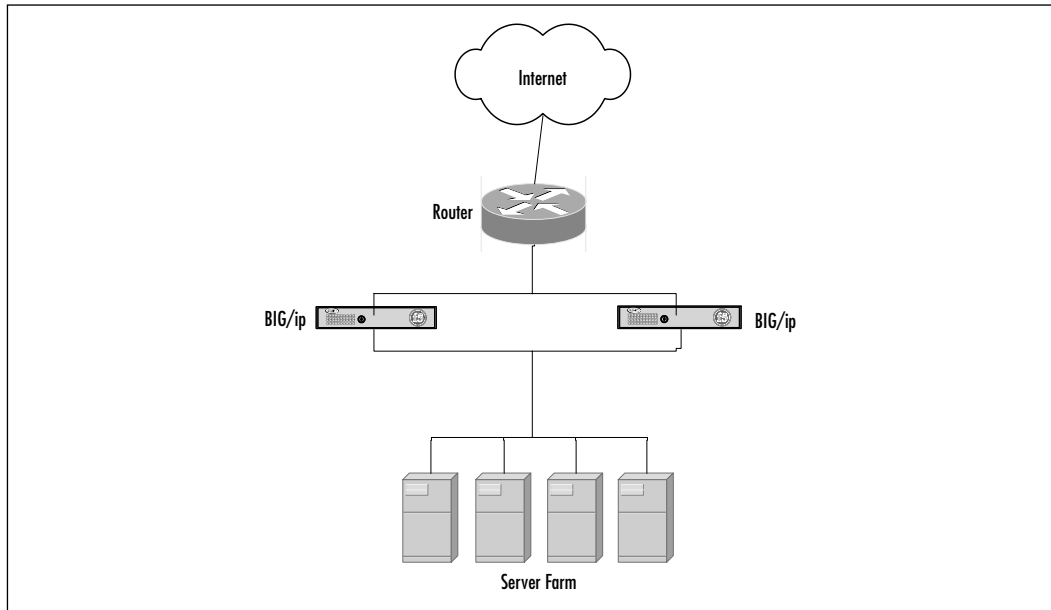
## Network Appliances

Network appliances (e.g., F5 Networks' BIG-IP and 3-DNS controllers) are hardware-software products that are designed to offer full IP support and augment network performance. Their common design will have two redundant load balancers paced between the server array and the network (Figure 4.5). This allows the appliance to operate in cooperation so that they are able to perform parallel and hot-spare load balancing. This redundancy will offer a fail-safe, cost-effective solution that can significantly minimize support. Servers are capable of being upgraded and supported with little to no downtime for your infrastructure.

This high-availability load-balancing solution can provide operating system independence, which allows your organization to implement any type of application or Web server into the network. This design approach can offer functionality,

speed, dependability, and scalability, and still be relatively cost effective. For those service providers that require continual e-commerce and secure connections, as well as user application interaction, these controllers can be valuable assets. The load balancing that is included is primarily subsets of, not a replacement for, applications whose usage is distributed across large clustering systems.

**Figure 4.5** Network High Availability with Network Appliances



## Criteria You Should Look for in a Superior Load-Balancing Solution

The following are some of the criteria that you may consider as requirements in your assessment of the right load-balancing product for your business.

### Dependability

These products need to have failover protection so that you can guarantee availability of your application. The load-balancing product should activate itself immediately once a failure is detected on the device. The malfunctioning unit should then be able to reboot itself automatically so it can come back online as the alternate.

## Quality of Service

When assessing the correct solution for your network, you should consider extensiveness and complexity of load-balancing features within each load-balancing product to see which is best able to meet your networking demands. Some features and options that could be integrated include numerous traffic distribution algorithms (such as round robin, round-trip time, packet rate, etc.), and QoS, which can monitor and manage the application traffic based on current conditions, and should dramatically improve performance.

Load-balancing products that you will consider should be able to detect errors and reroute traffic automatically by actively monitoring content and application performance. By testing performance at the application layer, the user connecting to your application is assured that all the different processes involved are thoroughly checked before service requests are routed to the server. This will diminish or even eliminate the occurrence of error messages that are common with overloaded servers, software failures, and bad or missing content.

Load-balancing products may also be able to provide things called *persistence* features. These features include SSL session ID tracking to ensure that users can stay connected to a single server while using an application. This is especially important to user environments (such as an ISP), where numerous users are assigned the same IP address. This is very confusing to a load-balancing product that causes traffic to be centered on a single server within the server farm. Instead, your load-balancing product should be able to read specific session IDs from an SSL transaction, thus assuring that the user can be uniquely identified and directed to the proper application in a timely manner, until the transaction is complete.

## High Availability

High availability can offer advanced functionality that is sometimes referred to as *traffic prioritization*. This allows the network to vary the user's access service levels based on traffic source, type, or destination, thus guaranteeing access. For instance, rules can be set up to give priority to transactions that are time sensitive, or allow application access from a specific company's domain. This solution can provide some of the most flexible abilities and help to optimize availability for a wide range of business applications.

The product should be able to examine and identify specific types of traffic based on HTTP header information. This gives you greater control over a wider range of traffic, since you can monitor at more granular levels of traffic detail at the application layer.

For the greatest amount of flexibility and control of traffic, the product should be able to recognize and handle high-availability load balancing to any HTTP header; this includes HTTP version, HTTP host field (also known as URL), and the HTTP method being used in the request (get, put, etc.).

Also, ask yourself if you need SSL acceleration, which is usually critical in e-commerce applications. SSL acceleration is used to offload SSL processing from the servers, thus increasing their performance and improving response time. This also helps with traffic management for a customer's transactions. SSL acceleration improves the performance of e-commerce servers, and can provide security, speed, and traffic management for online transactions. All this is from a centralized location without the cost of installing additional hardware or software on each server.

## Can Load Balancing Enhance and Extend Your Network?

Load-balancing products should also enhance your existing infrastructure and enhance security, while offering no limitation as to OS platform or network architecture you can use. Functional requirements include:

- Robust LAN/WAN load balancing capabilities for a scalable, extended network
- Enhanced network security, which includes the following features:
  - Secure socket shell (SSH) and SSL encryption for secure remote access and management either by a Web browser or command-line interface
  - Firewall capability with IP filtering
  - Network Address Translation (NAT)
  - The ability to avert teardrop/land/ping/denial-of-service attacks
  - Prevention of IP spoofing and SYN floods

## Vendor Credibility and Their Support Infrastructure

When choosing the vendor, you should look to see what customers the vendor has, and how well they are able to support them. A proven company will be able to showcase its device with customers that use e-commerce, have high customer

traffic loads at varying times of the day, and are able to distribute clients networks across multiple regions.

One of the most critical evaluation areas is technical support. Vendors should have a strong reputation in this area, by offering onsite installation and training for their systems. Remember that customer references are worth their weight in gold, so look for companies that have implemented these solutions and find out what they have to say about the vendor. Look for specialized technical expertise, as opposed to broad networking support, to ensure that your installed system will be optimized for your application and traffic requirements.

By increasing your application's availability, reliability, and performance, you will be able to offer your clients a high QoS. Today's users are fickle and will move on to the next vendor if they are faced with less than optimal performance and missing or bad content. The key is figuring out how to leverage your existing needs, while protecting your network investment and future growth, so you can create a high-performance service to satisfy your client's needs. Remember that picking a solution that is specifically designed for the task will generally offer better flexibility and performance for your application needs.

## Load-Balancing Solutions from F5

Round-robin DNS was originally implemented to address the issue of scalability. By configuring DNS to return the IP addresses of multiple servers configured to service `www.eXn.com` in a round-robin fashion, the traffic is distributed across the servers in a basic way. This solution does have several inherent drawbacks. First, round-robin DNS has no way or option to verify that the server address that is being returned to a user is actually working properly. Users could be directed to a server that is down or out of service for repair. To that user, `www.eXn.com` is effectively nonexistent, even though other servers for that site may be functioning properly. Second, when the DNS returns the IP address of a specific server to a user, that user's browser will retain the address in its cache. When the user attempts to return to a server for `www.eXn.com` that is no longer in service, it will return a "404 Object Not Found" error.

As more servers are added to the DNS round-robin rotation, traffic will be unevenly distributed. The older servers will tend to receive more traffic than newer servers, as the IP addresses of older servers are usually cached by more users than the addresses of newer servers are. This leaves businesses vulnerable to providing customers with bad or missing content.

## First-Generation Load-Balancing Solutions

The first generation of load-balancing products were able to provide scalability. They were able to fix some of the problems of round robin DNS by presenting a single IP address to end users by mapping requests sent to that address to multiple servers within the network.

These first-generation solutions were incapable of providing true high availability. The devices themselves were passive, so they did not perform active verification of the availability of servers or content that was located on those servers. Instead, they waited for the failure of actual traffic so that they could detect if a server was unable to respond. While this was able to provide adequate load balancing, this solution fell well short of providing the usability that is needed by businesses that demand 24 by 7 uptime.

## What Takes a Site Down?

To fully explain how the network appliance is able to provide both load balancing and high availability, it is important to understand what can cause an application to become unavailable.

There are basically five things that can stop a site from being available:

- **Content failure** The server and application are working properly, but they are responding to requests with “404 Object Not Found” or some other response that does not contain the right content or application.
- **Network unavailable** If a link between the server and the outside world becomes unavailable, the server becomes unreachable. This usually occurs from router failure, a configuration issue with the router, or a cut cable.
- **Server failure** The server becomes unavailable due to a hardware or operating system failure.
- **Software failure** The application hangs or stops responding, even though other applications are working properly.
- **Too much traffic (saturation)** Servers have a response curve in relation to load. As the traffic that they are serving increases, they are able to respond to requests promptly until the server reaches a point at which it stops responding to any request. To think about this in another way, this behavior can be viewed as binary; the server is either on or off.



## Guaranteeing Availability to Your Client

The only way to guarantee that a site is always available and be able provide customers with the quality of service that they expect is to protect against the five possible points of failure that are outlined in the preceding section. You need to protect your sites and applications against all of these points of failure in the following ways:

- **Content failure** As stated earlier, this occurs when the server and application are working properly, but they are responding to requests with “404 Object Not Found” or some other response that does not contain the right content or application. Some network appliances will actively query the servers at the application level in order to defend themselves against this. If an application is not returning the right content or system status, the device should redirect requests and applications to servers that are responding properly. Once the failure is fixed, the devices should automatically detect that the application or content is responding properly, and begin sending requests to it. This functionality will allow you to extend your infrastructure protection to your applications, such as databases.
- **Network unavailable** By using a high-availability solution, businesses can provide redundancy and load balancing to their clients. End users see a single URL `www.eXn.com` and are directed to the geographic site that is best suited to provide the content or application with a high quality of service. This provides protection against network failures, failures that are related to a single data center, Internet slowdowns due to congestion, and overloaded server farms. Remember that businesses want their infrastructure working for them 100 percent of the time.
- **Server failure** Use your network load-balancing appliance in conjunction with two or more servers, so that traffic is automatically routed away from any server that fails or becomes unavailable. Some devices can proactively monitor the servers to detect if there are failures and to keep them transparent to clients using the application. When a server begins responding properly again, it is added back into the server farm.
- **Software failure** When an individual service stops running on a server, proactive monitoring can automatically detect the failure. Requests intended for that service are then sent to another server that has that particular service running as well. For example, if your servers are configured

to support both Enterprise Resource Planning (ERP) and HTTP, and the server's ERP serving process stops, the load-balancing appliance will continue to send HTTP traffic to that server, but will redirect the application traffic to other available servers. Once the ERP process on the server becomes available, the appliance will start sending ERP traffic to it again.

- **Too much traffic (saturation)** QoS is ultimately measured by how long a user must wait for a response. You want your device to protect against clients having to wait too long to receive a response by setting thresholds for acceptable performance. If a server, service, or application is unable to respond within the configured limits, then requests will be redirected to another server until response times return to an acceptable level.

When you implement a network device that is capable of high availability, you want it to guarantee that it can deliver IP-based services, which are always available. To do this, you must remember that it is imperative that both “quality of service” based high availability and load balancing are addressed so that your client has a good usability experience. By deploying a load-balancing solution that is not able to provide high availability, you will not be able to maximize the return on investment for your services.

## Cisco Systems' LocalDirector

When it was first deployed, Cisco Systems' LocalDirector was positioned as a solution for the “round-robin” issues that were encountered in the Internet. The networking and computing trade sheet dubbed them “load balancers,” which is really a misnomer when applied to the LocalDirector. While load balancers are able to equally distribute traffic loads across multiple servers, the LocalDirector is capable of many more things such as scalability, high availability, server connection management, and server security.

## Scaling a Server Farm

There are generally two approaches for scaling a server farm-based system. The first approach is to continuously upgrade the size and processing power of individual servers in the farm. The second approach is to add more servers as you require more capacity.

In many cases, you will need to deploy a plan that allows you to increase your capacity by adding more servers to the system. There are several reasons for adding capacity, including:

- Unanticipated or hyper growth in server traffic
- Internal requirements for redundancy
- Budget constraints that do not allow for wholesale upgrades

When a multiple server environment is created, there are management concerns such as how to best distribute traffic loads so that you get the most utilization of the available resources. The answer might seem simple: just monitor a few key variables such as CPU utilization. In fact, Cisco's models show that three key variables, which are very difficult to monitor, must be considered when optimizing server system capacity. The three variables are network bandwidth, server performance (including application performance), and job size.

One of the main goals that your system needs to ensure is that network bandwidth and server performance are at full optimization. This implementation can help to identify the source of future performance issues so that you can avoid the expense of having to continuously increase server performance when performance bottlenecks are encountered. Because bandwidth is more expensive than server capacity, relatively speaking, it is generally easier to find server systems that have more than adequate resources to handle the total amount of bandwidth that is available in your network.

Keep in mind that you must consider the average jobs or application size that customers are requesting from your servers. Job size is truly important, as the traffic load and the capacity of server systems must be analyzed so you can make sure that the device can handle the capacity. For example, if users typically request small text files, the server should be able to handle more jobs than are normally requested as compared to sophisticated database queries, or downloading large graphics files.

Load-balancing technology does not normally consider variables such as bandwidth, server performance, and job size for optimizing the traffic loads among your server farms. Load balancing can allow you to incrementally scale the capacity of servers in your server farms in a more efficient manner. There is the ability with today's load-balancing devices that will allow you to monitor and manage the number of TCP connections that are allowed to each server.

What this allows you to do is test scalability by gradually adding TCP connections to a server to see what the real capacity is for actual traffic demands. More servers can be added when you encounter saturation of the network, which you can usually tell by slow response times, or lack of communication altogether. What load balancing really offers to you is predictability within the

network, which will help you better plan for server resources and support the growth of server farms.

## High Availability

When connectivity problems prevent access to applications such as ERP, companies could lose an estimated \$5000 to \$20,000 per minute of business downtime. These costs are accrued due to the customer's inability to access your revenue-generating Web site or applications. These costs affect a company's bottom line and its reputation.

Load balancing, and Cisco LocalDirector technology specifically, is a possible solution to implement high availability of critical Web, database, and application connections. A LocalDirector can improve the uptime of the server farm by allowing you to design your networks and server farm build-outs so that you can increase server and application availability by enhancing your redundancy.

LocalDirector is considered a transparent device, as it is able to work with any TCP-based service or application. There is no special software required on the server, as these are external devices. Transparency is one of the main reasons for the product's success (as well as the SmartNet Package, which requires that you purchase the LD), as other high availability devices can require that application software needs to be written for different server architectures.

LocalDirector can provide three components for its availability solution:

- Server availability
- Application availability
- LocalDirector availability

What the LocalDirector does is determine what the server and application availability is by monitoring the TCP connection state. If the server becomes unavailable, the LocalDirector will transparently redirect traffic to another server. When the server becomes available again, the LocalDirector will resume sending traffic to that server.

The LocalDirector is equipped with a hot-standby, failover mechanism. If the LocalDirector unit fails, a failover unit can assume load-balancing duties. For applications that require longer duration, and are connection oriented, you can set up failover between LocalDirector in a stateful mode, so that clients do not need to log in again to access the application or server. You can also add redundancy for the server farms by implementing the Cisco Hot Standby Routing Protocol (HSRP) on the connecting router.

Many companies implement LocalDirector systems to solve the availability puzzle by using redundant server farms, and other networking equipment at one local and self-contained site. A manager can further increase availability by building an identical server system at a different location. In this design, Cisco DistributedDirector balances traffic loads among multiple server sites that are managed internally by Cisco LocalDirector units. If one of the server systems fails for any reason, an identical server system in another location is ready to receive and handle client requests.

## Configuring & Implementing...

### Clustering Technology

Clustering technology is one example of a high-availability solution that has not caught on or become very dominant because it requires server-specific software. This makes implementation and support very time and resource consuming.

## Managing Your Server Connections

The LocalDirector is considered a stateful device, as it is able to monitor and can track all TCP connections that are occurring between clients and servers. This monitoring capability allows you to reduce the inconsistencies that can be associated with load balancing; it can facilitate the identification of failed servers, and allow you to better manage the infrastructure as a whole. Some industry analysts argue that server connection management is the most important reason to implement load-balancing technology.

Cisco's LocalDirector was the first device that allowed the ability to take a server out of the production environment, upgrade an application, and then return the server into production. It also afforded the ability to set a maximum number of TCP connections for each server. This was a way to let you monitor and know that a specific application could handle several hundred simultaneous connections before it crashes, so that there were no surprises in your network.

LocalDirector has another capability, called Real-to-Virtual-to-Real (RVR) communication for managing server connections. RVR can help you enable a real server to access a group of real servers by a virtual IP address. You generally

see this implementation in a server farm of two Web servers and two database servers. When a request to one of the Web servers needs to access the database to do a lookup, LocalDirector will allow for the two databases to be accessed and load balanced by the virtual address.

A LocalDirector device is often deployed so that it can direct traffic to a server based on a source IP address that accesses a virtual IP address. For instance, any traffic that comes in from high-priority customers can be directed to a server that is faster and has more resources than others within the network have. This creates a level of differentiation that is very similar to those for quality of service in networks.

## Security with the LocalDirector

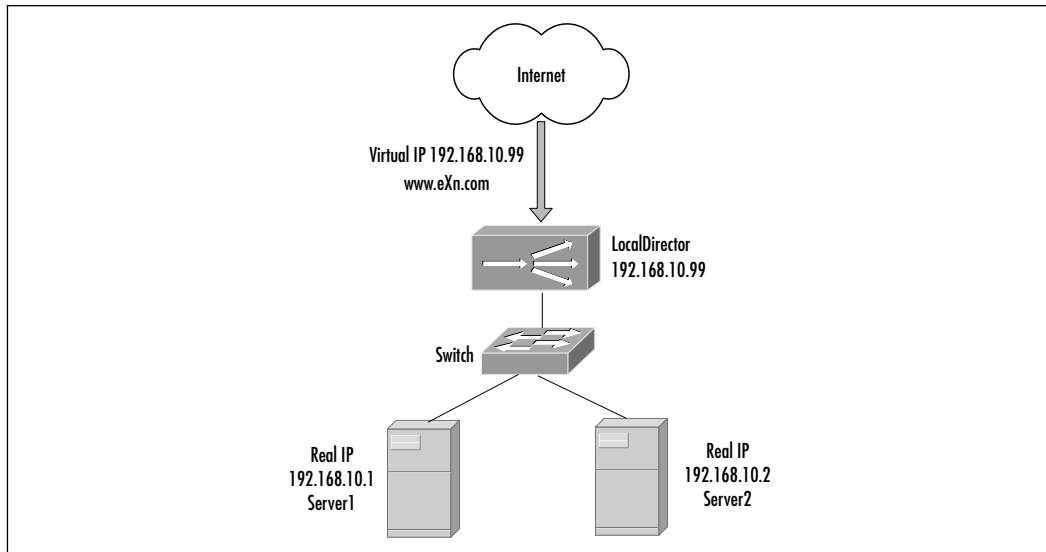
Cisco LocalDirector can also has the ability to help with server security for your system. While a LocalDirector is not a firewall product, the security features in LocalDirector can be an added layer of security for many of your applications. Servers themselves generally do not have security capabilities built into them. A LocalDirector can provide a level of protection against unauthorized access.

The LocalDirector devices can handle the filtering of access traffic based on source IP address and port, within its chassis. A LocalDirector can be set up to block traffic from a specific class of addresses, or by policies; for example, only Web traffic is permitted access to the servers. There is also the built-in capability to secure specific ports, which means that traffic is not bridged through to a real server. By using NAT, a LocalDirector can allow unregistered IP addresses to protect the servers from external attack. LocalDirector can also provide a layer of protection against SYN attacks by allowing you to set the number of unanswered SYNs before it enters into protected mode. This is discussed a bit more in depth in Chapter 6.

## LocalDirector Configuration Samples

The following section contains some examples of how to implement and configure a Local Director within your infrastructure. There are several methods discussed, and these are what we have seen in several of the build-outs in which we have taken part.

**One virtual server and multiple real servers.** In this example, LocalDirector is load balancing all TCP traffic connection over two servers to provide for Web-based services. Figure 4.6 shows the network configuration.

**Figure 4.6** LocalDirector Load Balancing between Two Servers

All traffic that is destined for the virtual IP address of 192.168.10.199 is load balanced across real servers with IP addresses 192.168.10.1 and 192.168.10.2. Only the virtual server appears in the Domain Name System (DNS) tables. Follow this procedure to set up this configuration:

1. Use the *enable* command to enter privileged mode. Type a carriage return at the password prompt if you do not want to assign a privileged password (This is not a long-term option; this is only for the sample configuration. You should always implement a password on production equipment.)

```
LocalDirector# enable
Password:<CR>
```

2. Use the *Configuration Terminal* command to enter configuration mode:

```
LocalDirector# configuration terminal
```

3. Use the *ip address* command to specify LocalDirector IP address 192.168.10.99, and subnet mask 255.255.255.0:

```
ld(config)# ip address 192.168.10.99 255.255.255.0
```

4. Use the *interface ethernet* {interface number} command with the *auto* option (if your interface card supports this option) to automatically set the speed of the Ethernet interface:

```
ld(config)# interface ethernet 0 auto
ld(config)# interface ethernet 1 auto
```

5. Use the *shutdown interface* {interface number} command to disable unused interface ports:

```
ld(config)# shutdown interface ethernet 2
ld(config)# shutdown interface ethernet 3
```

6. Use the *name* command to identify IP address *192.168.10.199* as *domain*, and the *virtual* command to define *domain* as a virtual server:

```
ld(config)# name 192.168.10.199 domain
ld(config)# virtual domain
```

7. Use the *name* command to identify IP address *192.168.10.1* as *server1*, and *192.168.10.2* as *server2*:

```
ld(config)# name 192.168.10.1 server1
ld(config)# name 192.168.10.2 server2
```

8. Use the *real* command to identify *server1* and *server2* as real servers, and the *is* option to enable the real servers to start accepting connections:

```
ld(config)# real server1 is
ld(config)# real server2 is
```

9. Use the *bind* command to associate *domain* with *server1* and *server2*, and establish the load-balancing relationship between the virtual and real servers:

```
ld(config)# bind domain server1 server2
```

10. Use the *is* command to bring the *virtual domain* server into service:

```
ld(config)# is virtual domain
```

11. Use the *write terminal* command to view the running configuration before it is saved. (This will be different for all configurations, so it has been left out of the text.)



12. Use the *write memory* command to save the new settings:

```
ld(config)# write memory
Building configuration...
[OK]
```

13. View the saved configuration with the *show configuration* command:

```
ld(config)# show configuration
: Saved
: LocalDirector 420 Version 3.10.0.106
syslog output 20.3
no syslog console
enable password [Edited] encrypted
hostname localdirector
no shutdown ethernet 0
no shutdown ethernet 1
shutdown ethernet 2
shutdown ethernet 3
interface ethernet 0 100basetx
interface ethernet 1 100basetx
interface ethernet 2 100basetx
interface ethernet 3 100basetx
mtu 0 1500
mtu 1 1500
mtu 2 1500
mtu 3 1500
multiring all
no secure 0
no secure 1
no secure 2
no secure 3
ping-allow 0
ping-allow 1
ping-allow 2
ping-allow 3
ip address 192.168.10.99 255.255.255.0
```

```
route 0.0.0.0 0.0.0.0 172.16.30.1 1
no rip passive
failover ip address 0.0.0.0
no failover
password dft
no snmp-server contact
no snmp-server location
casa service-manager port 1638
virtual 192.168.10.199:0:0:tcp is
real 192.168.10.2:0:0:tcp is
real 192.168.10.1:0:0:tcp is
name 192.168.10.1 server1
name 192.168.10.2 server2
name 192.168.10.99 domain
bind 192.168.10.99:0:0:tcp 192.168.10.2:0:0:tcp
bind 192.168.10.99:0:0:tcp 192.168.10.1:0:0:tcp
localdirector(config)#
```

## Multiple Virtual Servers and One Real Server

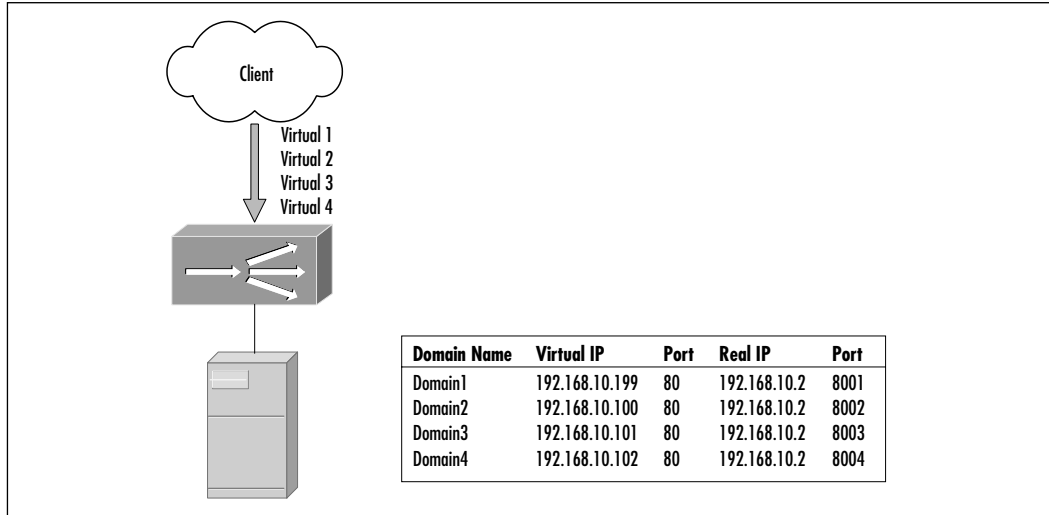
For this example, we have four virtual addresses that are bound to a single Web server, as shown in Figure 4.7. This will allow you to provide multiple DNS entries for one server. To say it another way, there is one real server that is able to support multiple domain names. The virtual IP addresses 192.168.10.199, 192.168.10.100, 192.168.10.101, and 192.168.10.102 are identified as domain1, domain2, domain3, and domain4, respectively. Port 80 traffic for each virtual IP address is bound to different ports on real server IP address 192.168.10.2.

All Web traffic that is destined for domain1 will access information on real server 192.168.10.2 through port 8000. Traffic that is destined for domain2 will access information on real server 192.168.10.2 through port 8001, and so on.

By means of defining a virtual server as an IP address and a port, you will be able to restrict traffic to a specific port. Therefore, port 80 is specified for each of the virtual servers, and ports 8000, 8001, 8002, and 8003 are specified as the ports for the real server. The virtual server ports and real server ports are bound to each other directly using a bind-id on the real server for each port that is bound. The great thing about this is that if the application that is running on port 8000 fails,

LocalDirector does not fail the entire server; the remaining ports will continue to accept connections.

**Figure 4.7** Multiple Virtual Servers with One Real Server



The following is an example of how to configure multiple virtual servers for one real server.

1. Use the *name* command to identify the IP addresses of the virtual and real servers:
 

```
ld(config)# name 192.168.10.99 domain1
ld(config)# name 192.168.10.100 domain2
ld(config)# name 192.168.10.101 domain3
ld(config)# name 192.168.10.102 domain4
ld(config)# name 192.168.10.2 server
```
2. Use the *real* command to identify the IP address named *server* as the real server that is accepting connections on ports *8000*, *8001*, *8002*, and *8003*:

```
ld(config)# real server:8000
ld(config)# real server:8001
ld(config)# real server:8002
ld(config)# real server:8003
```

3. Use the *virtual* command to identify the named IP addresses *domain1*, *domain2*, *domain3*, and *domain4* as virtual servers accepting connections on port 80:

```
ld(config)# virtual domain1:80
ld(config)# virtual domain2:80
ld(config)# virtual domain3:80
ld(config)# virtual domain4:80
```

4. Use the *bind* command to direct traffic that is destined for port 80 to a different port on the real server:

```
ld(config)# bind domain1:80 server:8000
ld(config)# bind domain2:80 server:8001
ld(config)# bind domain3:80 server:8002
ld(config)# bind domain4:80 server:8003
```

5. Use the *is real* command to set the service state for all *real server* ports to in service:

```
ld(config)# is real server:8001
ld(config)# is real server:8002
ld(config)# is real server:8003
ld(config)# is real server:8000
```

6. Use the *is virtual* command to set the service state for all virtual server ports to in-service:

```
ld(config)# is virtual domain1:80
ld(config)# is virtual domain2:80
ld(config)# is virtual domain3:80
ld(config)# is virtual domain4:80
```

7. Use the *show bind* command to display the association between the virtual server ports and real server ports:

```
ld(config)# show bind
Virtual Machine(s)      Real Machines
domain2:80:0:tcp(IS)   server:8001:0:tcp(IS)
domain1:80:0:tcp(IS)   server:8000:0:tcp(IS)
```

```

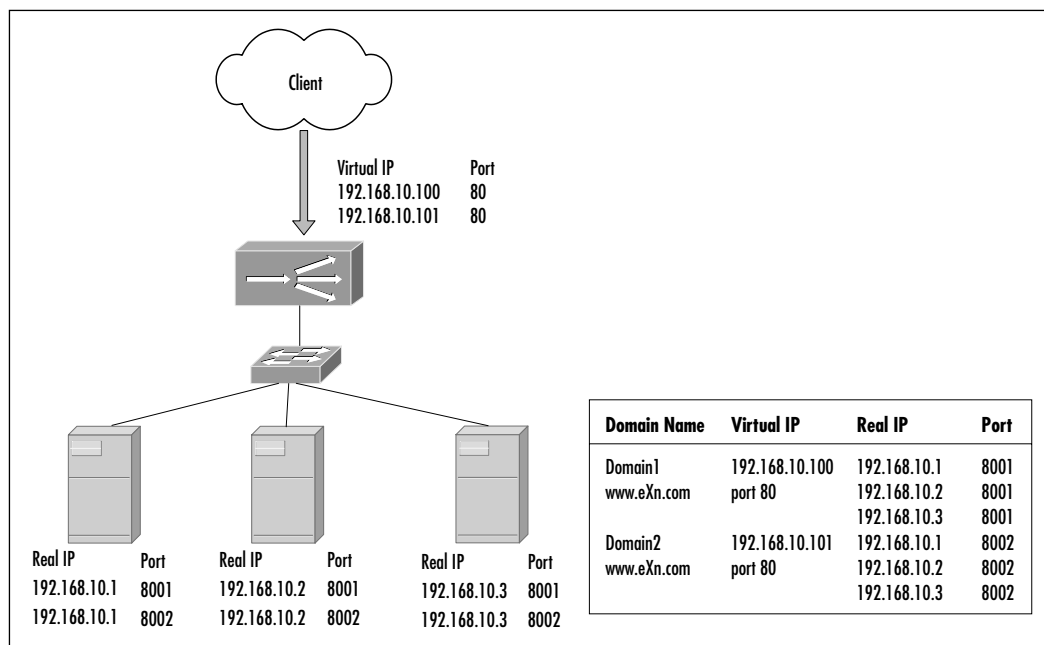
domain4:80:0:tcp(IS)    server:8003:0:tcp(IS)
domain3:80:0:tcp(IS)    server:8002:0:tcp(IS)
ld(config)#

```

## Multiple Virtual Servers and Multiple Real Servers

You can also combine multiple virtual servers and multiple real servers so that each virtual server is able to send network traffic to the same port across all of the real servers. As shown in Figure 4.8, all TCP traffic that is destined for virtual server 192.168.10.100 is load balanced across the three real servers on port 8001. Traffic that is destined for virtual server 192.168.10.101 is also load balanced across the real servers on port 8002.

**Figure 4.8** Multiple Virtual Servers and Multiple Real Servers



With this combination of virtual servers and real servers, you can also load balance traffic across server clusters. Each virtual server can have a different load-balancing option set with *predictor* command. For instance, 192.168.10.100 can be configured to use the *leastconns* option, and 192.168.10.101 can be configured to use the *weighted* option.

The follow is an example of how to configure multiple virtual servers for multiple real servers.

1. Use the *real* command to identify three real servers, each accepting connections on ports *8001* and *8002*. Use the *is* option to put the real servers in service:

```
ld(config)# real 192.168.10.1:8001 is
ld(config)# real 192.168.10.1:8002 is
ld(config)# real 192.168.10.2:8001 is
ld(config)# real 192.168.10.2:8002 is
ld(config)# real 192.168.10.3:8001 is
ld(config)# real 192.168.10.3:8002 is
```

2. Use the *virtual* command to create two virtual servers accepting connections on port *80*:

```
ld(config)# virtual 192.168.10.100
ld(config)# virtual 192.168.10.101
```

3. Use the *bind* command to direct network traffic from port *80* on the two virtual servers to ports *8001* and *8002* on the three real servers:

```
ld(config)# bind 192.168.10.100:80 192.168.10.1:8001
ld(config)# bind 192.168.10.100:80 192.168.10.2:8001
ld(config)# bind 192.168.10.100:80 192.168.10.3:8001
ld(config)# bind 192.168.10.101:80 192.168.10.1:8002
ld(config)# bind 192.168.10.101:80 192.168.10.2:8002
ld(config)# bind 192.168.10.101:80 192.168.10.3:8002
```

4. Use the *is virtual* command to set the service state for all virtual server ports to in-service:

```
ld(config)# is virtual 192.168.10.100:80
ld(config)# is virtual 192.168.10.101:80
```

5. Use the *show bind* command to display the association between the virtual and real servers:

```
ld(config)# show bind
```

```
Virtual Machine(s)      Real Machines
```

```

192.168.10.101:80:0:tcp(IS)
    192.168.10.3:8002:0:tcp(IS)
        192.168.10.2:8002:0:tcp(IS)
            192.168.10.1:8002:0:tcp(IS)

192.168.10.100:80:0:tcp(IS)
    192.168.10.3:8001:0:tcp(IS)
        192.168.10.2:8001:0:tcp(IS)
            192.168.10.1:8001:0:tcp(IS)

ld(config)#

```

## Foundry Networks' ServerIron

Foundry's ServerIron Web switches provide high-performance content and application-aware traffic and server load balancing. ServerIron has the functionality of a traditional Layer 2 and Layer 3 switch built in, and is able to examining the content at Layer 4 and above through the packet header. This ability helps to provide robust content switching.

Foundry has several common features within its ServerIron Line that we have found useful in the past. Some of these features include (this list is subject to change with time, and with model number):

- Several methods of session persistence, including cookie and SSL session ID-based switching.
- The ability to select a server based on host name, prefix, or suffix of the URL, or by matching a pattern expression against the URL.
- The ability to implement up to 256 URL switching policies, with support for nested URL policies so that you can do complex pattern matching.
- The SwitchBack technology that provides wire-speed throughput for server-to-client reply traffic.
- Active/Active and Active/Standby features that support stateful failover to ensure that there is no disruption to client sessions.
- A command-line interface (CLI)
- Web-based GUI
- SNMP-based management
- NAT

- Port Address Translation (PAT)
- Built-in SSH
- Access control lists (ACLs) and extended ACLs.

Foundry Networks' ServerIron server load-balancing switch solutions can provide ISPs with a high-density and high-performance Layer 4 switching. This will improve the performance of existing servers and guarantees an application's availability while increasing network redundancy.

ServerIron solutions do not depend on specialized, vendor-dependent server software, as these devices are protocol based. This ensures that there is compatibility with all major server vendors and operating systems. Therefore, no special server agent software is required for these devices to function properly.

ServerIron switches are deployed in all forms of server farm environments, including Web, FTP, mail, and other application-based settings. ServerIron is able to provide high-speed connectivity to your server farms. These links can be multihomed for greater aggregate bandwidth and more complete redundancy.

ServerIron load-balancing characteristic is based on Layer 4 traffic such as HTTP, FTP, SSL, and email. This creates the ability to transparently distribute data traffic among multiple servers. ServerIron allows you to create logical communities of servers that are represented by a single IP address. Many individual servers can reside in multiple logical communities, thus reducing setup and deployment of servers that support multiple types of end users and clients. By creating a single-server view, you can ease management, and protect your servers from unauthorized access.

Three main options are available in the ServerIron family to optimize application traffic without exceeding the server's capacity:

- **Round robin** This will assign connection in a sequential order among all servers that are located within a logical community. Round robin will treat all servers as equals, in spite of the number of connections or response time that is being experienced.
- **Least connections** This assigns a connection to the server that has the least number of open connections. This is especially useful for Web sites in which there are groups of servers with similar performance and function. Least connection tries to ensure that there is adequate distribution of traffic, in case a server within the community becomes overloaded.



- **Weighted percentage** This option allows you to assign performance weights to each server in a community. These weighted percentages can then be used to calculate which servers are capable of processing connections faster, so as to receive the largest number of connections.

ServerIron increases network resiliency for your mission-critical applications, and includes essential, dependable features that will ensure applications' availability. If there is a server or application outage, the ServerIron can provide millisecond detection and failover to the next server that is located within the same logical community that supports that application. This will guarantee that your data will continue to flow and your applications are always available to the client.

An added feature that will help you meet your customer needs for uptime is a hot standby redundant switch capability that can protect your users against loss of session. The ServerIron creates primary and secondary switches that have identical configuration parameters. The secondary switch (or failover switch) will continuously monitor the traffic that passes through the primary switch. If the primary switch fails, the secondary switch will take over without losing user sessions or connectivity.

## Content Delivery Networks

The Internet has grown to the point where its value transcends IP connectivity for the support of Web pages and email. ASPs, content providers, Web-based publishers, and e-businesses look to the Web for high-performance, reliable transport for bandwidth-intensive, multimedia content such as IP (XoIP), e-commerce transactions, special events, news, and entertainment services.

With this need, there comes the requirement to develop dynamic multimedia content. The networking industry's focus from Layer 3 connectivity issues is shifting to the creation of intelligent, Layer 4–7 networks that can support the rigorous response-time requirements of these new types of content. The emphasis is now turning to content delivery networks (CDN).

Some of the reasons for the movement to CDN include:

- The speed of development and deployment can be much faster at the higher network layers.
- There is a need to grow to improve quality of service and quality of experience for more-demanding clients.
- Content providers must scale to service larger audiences who often consume the same content.

CDNs can also leverage the abilities of strategically placed distributed caching, load balancing, and Web request redirection systems. They ensure that, based on closeness and server resource utilization, content is delivered in the most efficient manner to the user requesting it. This benefits the end user (and therefore, the content provider), as well as connectivity providers, who receive the benefits of a streamlined bandwidth consumption implementation.

Content is normally served from a cache server, which in this model can also be called a *surrogate*, which is located on the edge of a network, close to the user. The surrogate mirrors the content provider's primary servers, which are located in a CDN service provider's data center. This enables CDN service providers to deliver the highest-quality experience to the end users, who are intolerant of response-time interruptions.

Some of the reasons that are driving the CDN's growth are the network design and application requirements of content providers, which are causing increasing numbers of content publishers to consider the economic savings that are offered by CDN service providers. Sites that deal with streaming-media events (such as live conference events and training companies), high-volume e-commerce transactions during holiday seasons, and interactive videoconferencing sessions are just a few of the businesses that are capable of receiving these benefits.

The task of keeping these complex content sites advantageously dispersed and available to a wide base of users can be a costly and time-consuming undertaking for content providers. Web content providers need the following to productively build and maintain their multimedia content:

- Near-100 percent (99.999) server uptime and availability, while still delivering fast response times to users.
- The ability to reach a wide base of customers in a cost-effective, scalable manner.
- Content management and monitoring tools that allow providers to keep their content fresh ready, and track user activity.

## Today's Content Delivery Landscape

Most of the larger content service providers have hosted their own content while monitoring and managing their own Internet connections. However, as Web usage for a business foundation proliferates the market and content distribution demands increase, there has been a mass conversion to CDN-type service providers. The

reasons for this migration include increases to the performance and reliability of their content services, while lowering their total cost of ownership.

Content providers are driving the industry, and their vendors, to develop standardized new technologies. The trends also motivated the industry to construct and implement peering and settlement capabilities among CDN service providers to ensure dependable, high-quality service levels. Two industry groups, the Content Alliance and the Content Bridge Alliance, are establishing these technical and business standards along with the Internet Engineering Task Force (IETF).

## Functional Components of a CDN

The components that are necessary for a CDN to function properly include the following:

- There should be a redirection service that makes sure that a client's Web request is directed to the "closest" cache server.
- Distribution services that are comprised of a distributed set of surrogate servers, so that cache content can be accessed on behalf of a Web owner's point-of-origin server. This enables traffic to bypass heavily congested areas of the Internet. An example of this is when IP multicast might be implemented as a component of the distribution services such as a medium for updating mirrored caches.
- An accounting and billing system that enables the CDN provider to monitor, log, and bill the content provider based on use (the amount of bandwidth consumed by users who access the content provider's site). These systems have also evolved to enable CDN providers to crossbill multiple providers for CDN internetworking services.

## How Do CDNs Work?

CDNs are able to provide QoS to the Internet's IP-based backbone, which helps to eliminate or minimize delay (sometimes referred to as "World Wide Wait").

These latency issues are usually unnoticeable when the application that in use is providing email or static Web page downloads. However, as we move to a world that uses multimedia-rich applications for entertainment services, online gaming, live videoconferences, and streaming broadcasts, all of which are susceptible to response-time delays, extra preparation must be taken to ensure the delivery of a quality experience for the client.

CDNs are capable of addressing these response-time delays by minimizing the number of Internet backbones that an application request and return path that streaming or downloadable content must pass through. Surrogates are one possible way to do this by hosting replicated content in cache servers located on the network edge. This setup enables CDN service providers to deliver content that is stored on these cache servers to be just one hop away from the end user.

User requests to a content provider's Web site are redirected to dispersed data centers that the CDN providers own or lease. Setting up rules, and using the encoding methods that are dictated by the CDN service provider accomplishes redirection. The CDN maintains a service that is able to do lookups, which help steer user requests to the content surrogate that is closest to the client.

CDN service providers also use load-balancing technology to determine if the content server is available and considered the closest. As noted earlier in this chapter, this load balancing can take the form of software or hardware (such as a network appliance) from a third-party vendor.

CDNs are also using content switching or application-layer (Layer 4–7) switching to further enhance QoS abilities. This enables the CDN routers and switches to examine IP address information, and account for the specific response-time requirements of the application or content being requested. These Layer 4–7 switching capabilities can also be delivered in the form of software that overlays the router or switch, or as server software that runs in concurrence with the router or switch.

The surrogates contain software that is able to create logs that track usage and billing information, which is then collected by a central database in the CDN service provider's data center. This information can then be used to determine what to charge for the CDN services provided. These pricing models are generally based on usage, with a fixed rate billed per megabit per second of usage. This collected information by surrogates can also be used to deliver content management capabilities, based on usage trends and performance information, to the customer as a part of the CDN service package.

## Who Needs CDNs?

So, who really benefits from CDNs? The list includes content publishers, CDN service provider specialists, ISPs, CDN infrastructure component makers, local access telecommunications providers, and the content consumers in the public market and in the private sector. Usually, the CDN service providers “own” the content provider clientele, while the ISP or local access provider “own” the end customers (the consumer of the content).

The following are brief looks at each type of provider, why they need CDNs, and what their requirements are.

## Content Providers

Content provider organizations build content for the Web, and are faced with delivering content that has dynamic characteristics to customers who require high levels of service. One of the major issues facing the content provider market is the need for QoS levels in content delivery, as this is what attracts and maintains clientele. The following are examples of companies that are in the content provider category:

- Owners of e-commerce sites, who are concerned about response times for Internet browsers and transaction times for customers
- Retailers who want to broadcast promotional events
- E-learning developers, virtual universities, and traditional sales training companies that are adding Web-based versions of their classes
- News organizations that want to present Web-based video news coverage
- TV stations, radio stations, and entertainment companies that want to use the Web to deliver entertainment services
- Businesses that have mission-critical Internet-based content

Many content providers host and manage their own Internet content sites, and also support mirrored content servers in multiple locations. The reason for this is that CDNs are relatively new and little is known, and there is a perception that there are high costs associated with CDN services.

Currently, some companies are becoming well known in the CDN marketplace. Akamai Technologies Inc., Inktomi, and Digital Island Inc., for example, are priced at a per-megabit per-second of usage. These prices often seem high to content providers when compared with the per-megabit per-second of usage that is charged by most plain Internet hosting and connectivity services such as AboveNet Communications and Exodus Communications Inc. This cost is somewhat dispersed as content providers find themselves needing to run multiple data centers to efficiently serve content based on geographical location, and start totaling up the necessary hardware resources, network connectivity costs, and the human resources that are required to support their sites on a 24 by 7 timeframe.

Calculations by HTRC Group, a networking research firm in San Andreas, California, indicate that as content distribution is outsourced from content

providers, their performance increases and their support costs decrease. The reasons are the same as those found in a typical network outsourcing model; by being able to use the economics of scale, breadth of skill set, and networking expertise, companies that focus on the distribution of network content for many companies are able to priced better, faster service to each customer by sharing their resources.

## What Do Content Publishers Require from CDNs?

Content providers require a combination of hosting and distribution capabilities, so they are pushing the industry to develop newer products and to work together more effectively, not only on technical levels, but also conform to business standards as well. As discussed earlier in the chapter, this resulted in the formation of two industry groups, the Content Alliance and the Content Bridge Alliance. These companies try to foster IETF standards for the interoperability of CDNs on technical and business levels.

- **Multiple provider capabilities** Many of today's largest publishers are taking advantage of CDN services to help them leverage their services. In the past, the coverage offered by a single service was adequate to sustain the relatively limited amount of multimedia content that was commonly used by the Internet. Now, content providers are requiring a broader network reach, but don't want to deal with creating many relationships with multiple CDN service providers. Therefore, these content providers are looking for interoperability among CDNs, where the owners of the CDN could reimburse each other for shared distribution services. This helps to give content providers a more ubiquitous model and more flexibility to choose or change providers.
- **The ability to edit/redirect Web links** Content providers who use CDN services need to be able to edit their content and the links in their Web pages to point to the network of their CDN service provider. In the past, this required that changes be made to their internal naming conventions to that of the CDN operator. One of the drawbacks associated with this change is the fact that it tends to lock the provider into the redirection of content that is confined to a single CDN, which has often discouraged some Web site owners from using CDN services.

There are generally two accepted methods that are used by content providers to redirect their links to a CDN service provider's network.

The first works when CDN service providers tell the content provider to refer its internal Web links in its own DNS server to the DNS name of the CDN. This method makes it easier for content providers to change CDN service providers. The second method requires that the content provider program the name of the CDN service provider into its own DNS, (this is also referred to as *canonical names*, or *C-names*). The tradeoff with this method is that the content provider will gain a broader set of content management and network visibility capabilities, since its server is integrated into the CDN network infrastructure, but it is not as easy to change providers.

There are also efforts to assist content providers through products from companies such as CacheFlow Inc. and Novell Inc. They have created an application that works as front end for a content publisher's Web site and helps to perform URL rewriting on the fly. This protects the publisher from having to rewrite the Web pages and redirect user requests. It also prohibits the publishers from having to obligate themselves to a CDN. Currently, these rewriting solutions will work with multiple service providers' CDNs, but are only able to redirect content to just one provider.

- **Content management/usage visibility** There are now complex, yet user-friendly content management tools that are a critical success factor for CDN services. Content providers must be able to redirect their content to (multiple) CDN sites with minimal configuration and support; and have access tools that assist to refresh their content. Many CDN service systems are designed around a pull model for HTTP objects, which allows the ability to update distributed surrogates throughout the CDN. The surrogates use an algorithm that allows them to detect a mismatch between the surrogate and the point-of-origin server. When a mismatch is detected, the surrogates will perform an update to themselves. CDN service providers also supply their content provider customers with a Web-based front end. This allows the content provider to see what content is being served, and purge and update content as necessary.

Content providers use these tools to track usage histories and trends. This information is required for reporting and promoting their sites to potential advertisers. It is often used to see usage trends to determine how to tweak their content to make it more attractive to customers and provide a better quality of experience. This allows the ability for some content providers to personalize their services and applications at the

edge of a CDN provider's network, and helps to better target users in a given location, similar to a local newspaper or Yellow Pages directory.

The preceding capabilities make CDN migrations more flexible and less invasive on the part of the content provider, and contribute heavily to the acceptance and use of CDN services.

## CDN Service Providers

CDN service providers today include companies such as Akamai, Digital Island, epicRealm Inc., Inktomi, InterNAP Network Services Corp., Mirror-Image Internet Inc., and Speedera Networks Inc. The business model of these companies is to bring management and QoS to what have been mostly best-effort services to date.

They are looking to serve the growing market needs of their content provider customers and help to deliver content that require higher levels of service to end users. This trend gives CDN service providers an opportunity to add levels of control to the Internet, and to build themselves significant revenue streams.

Several of the early adopters of the CDN model have developed proprietary technologies for caching, content management, and load balancing. Some of these companies include:

- **Akamai** Akamai makes exclusive use of its own technology to optimize its ability to deliver better service to its clients.
- **InterNAP** InterNAP has created a software application called the ASsimilator, which works with the Border Gateway Protocol (BGP) moving data traffic from the ISP network that is closest to the user requesting the data, directly to the ISP backbone to which the content provider customer is connected. In order to make this work, InterNAP has created business partnerships with some of the world's largest ISP backbone providers.

When a customer makes a request to a Web site, the InterNAP partner ISP identifies it and transports the message to an InterNAP-owned data center. The traffic is then forwarded to the content provider's Web site. This setup bypasses ISP-to-ISP NAPs for public peering, and avoids much of the congestion on the Internet.

Many CDNs are currently owned and managed by a single body. However, there is a shift in thinking that is driving multivendor CDNs and allowing content providers and clients to benefit from the use of multiple providers. This is a



very scaleable way to add and maintain QoS and connectivity without having to design and implement their own infrastructures.

## What CDN Service Providers Require.

A CDN service provider will need to migrate from proprietary network intelligence to allow for the greater implementation and utilization for a multivendor design to fully realize its capabilities. Software functions that will be required from CDN manufacturers include:

- **The ability to handle accounting and billing** This assists the CDN in the ability to charge customers based on the tracking of usage between multiple CDN service providers that host and deliver content for common customers.
- **Content signaling technologies** These signaling technologies indicate when content should become invalid or when it should be refreshed, and can be extended across multiple CDNs providers.
- **E-commerce capabilities** (such as credit card verification, security, transaction processing) This capability can be used for the delivery of entertainment type services such as pay-per-view and gaming.
- **The ability to provide third-party clearinghouse services** These services are created to assist in the resolution of shared services among CDN service providers.

### Designing & Planning...

#### The Content Bridge Alliance and the Content Alliance

The Content Bridge Alliance is testing a concept for third-party clearinghouse services, using Adero Inc. in the multivendor delivery of content from America Online (AOL). The Content Bridge Alliance was created to proof a multivendor CDN model concept, before defining the technologies to be used in the design of these networks. In this testing model, vendors are participating in real-world multivendor CDN trials.

The Content Alliance alternatively has been working to define technology that supports multiple business models, not just the clearinghouse model, but also joint private peering.

## CDN Deployment Basics and Considerations

Three main architectures are used for deploying CDN services:

- **Facilities-based CDN** The provider owns data centers and provides network services across a wide geographic area, and distributes these services to end users. Usually, these are large ISPs that have built a CDN that works in conjunction with their Internet access and connectivity services.
- **Distributed or multinetwork CDN** In this type of deployment, CDN servers are placed in the PoPs of multiple facilities-based providers (the more the better), which creates an internetwork of CDN servers that spans multiple ISP backbones. This model is also referred to as a meta-CDN and is currently used by Akamai and Speedera.
- **Hybrid CDN** Companies that are deployed with this model maintain and monitor some of their own facilities, but also use the infrastructures of other ISPs or CDN service providers. Digital Island is an example of this model.

## Network Service Providers

ISPs and colocation companies that sell IP connectivity services are looking for ways to differentiate themselves and add new streams of revenue. ISPs that currently comprise the Internet backbone are moving quickly to implement CDN technologies such as load balancing and caching for the benefits they offer. By implementing these solutions, ISPs are more capable of monitoring and managing bandwidth, which is simply good business for network service providers who are looking to keep bandwidth costs low and traffic flow for their customers high.

Due to the nature of deploying these technologies, companies are deciding to become CDN providers on their own by offering content peering and internetworking agreements with other existing CDNs. They could also adopt other business models so that they may “plug in” their networks to multiple CDN infrastructures. They might for example, buy “edge” services from CDN providers to deliver broadband content to their own clients.

This means that local access providers are capable of delivering content for the “last mile,” and using caching and load-balancing capabilities for internal benefits. A company might do this because it may not have the network coverage to catch the attention of large content providers, but could get paid as the source for the ultimate delivery of the content from those providers.

## Satellite-Based Network Service Providers.

These emerging companies also have a place in the CDN value chain with the ability to deliver IP content directly to a local access provider's network edge, or even to a business or consumer site. Hughes Network Systems (their home-based satellite division) has started integrating Inktomi Corp.'s Traffic Server network caches at its network operations centers (NOCs). This enables the company to provide efficient delivery of IP-based applications directly to businesses and consumers.

Hughes and Inktomi are also in development for a new satellite-optimized caching and content distribution software that will then be implanted within Hughes satellite receivers in homes and businesses worldwide. These optimized satellites will operate as remote extensions of the Traffic Server caches that are deployed at the NOCs. Hughes will bundle Inktomi Traffic Server caching software with its DirecPC satellite platform, so that they can provide an integrated caching and content distribution solution.

## What Network Service Providers Require from CDN Service Providers and CDN Component Product Makers

ISPs need to create business relationships with CDN providers and work in tandem for the acceleration of content delivery that is required. Going back to the InterNAP example, direct connections run from InterNAP data centers to global Internet backbones that are managed and maintained by UUNet, Sprint, Cable & Wireless, Genuity, Digex, PSINet, AT&T, Verio, and Earthlink.

These connections are not free, for public or private peering, as InterNAP pays each of these backbones for TCP/IP transport. Because of this, InterNAP is able to provide its customers with differentiated quality of service.

## CDN Product Manufacturers

There is currently a large market for CDNs, and there are significant opportunities for CDN hardware vendors. The customer base of CDN product manufacturers are chiefly CDN service providers, ISPs that want to become CDN service providers, network service providers who are trying to gain more efficient operations in their own networks by using CDN technology, and enterprises that want to build corporate CDNs for internal applications.

CDN product manufacturers need to conform to IETF standards on content distribution internetworking so that there can be industry-standard technologies that allow interoperability between multiple vendors' equipment. Many CDN product manufacturers belong to one or both of the industry content alliances—the Content Alliance or the Content Bridge Alliance.

There are several bonuses to standardizing the CDN hardware; the development standards will accelerate CDN propagation and acceptance as a viable model. This will then allow CDN product manufacturers to sell more equipment and software. The Alliances have joined in supporting the IETF as the forum for the development of standards among CDNs.

## Enterprises

Enterprises will eventually need the same basic CDN infrastructure capabilities as the much larger public content providers. The nature of information that is being distributed by IP technologies within the enterprise is traveling the same path and eventually across the public Internet. Enterprises are developing their own internal e-learning or e-training content that needs to be distributed internally or streamed directly from the Internet. For very dispersed enterprise sites, these companies may want to use the public Internet for videoconferencing and bandwidth-intensive collaboration capabilities while cutting their infrastructure costs.

What do enterprises need? Enterprises that want to run their own internal CDN need to implement caching, redirection, load balancing, and content management tools, similar to that of the service provider CDN.

## Consumers

Consumers generally want the latest and greatest toys at their disposal. They thirst for entertainment, online gaming, distance learning, and videoconferencing capabilities that use their existing Internet connections. This is what will ultimately drive the content providers to offer these Web-based capabilities.

## The CDN Services Landscape

CDN services are fairly new, and many content providers do not know much about them. This is one of the major reasons why most Web site owners have typically hosted their own content. Moreover, using the Web for the delivery of multimedia services requires significantly higher QoS levels than are currently experienced by most consumers. There are also technical and business standards for CDN services that are still evolving.

## Industry Standardization Efforts

The effort to standardize CDNs will need to address the technical requirements for multivendor peering, and encompass areas such as standard methods of billing and settlement. The Content Alliance and the Content Bridge Alliance have made progress in convincing the IETF that a formal IETF group is needed to define and formalize CDN technical specifications. This group could then lead to specifications for multivendor CDN product interoperability.

Many of these technical issues were addressed at the December 2000 IETF meeting in San Diego, California. A CDN “Birds of a Feather” (BOF) meeting was a precursor to the forming of four new IETF Working Groups:

- **Content Delivery Network Peering (CDNP)** This group concerns itself with the specifications on how CDNs run by different operators will be able to share the information that is necessary for their CDNs to interoperate across administrative boundaries. Some of the specifications that are being defined by this working group are those that are designed on how to track usage and exchange billing information across network borders.
- **Open Proxy Extension Services (OPES)** This group defines the standards for how proxy caches execute code and enable special services such as redirection to a server storing foreign-language content. OPES standards specify how to encapsulate content and communicate with diverse servers.
- **Contextualization of Resolution** This group focuses on extending DNS naming capabilities to handle services that are more complicated. This could be a service such as identifying an IP address of a server that stores a foreign-language version of an organization’s content.
- **Web Replication and Caching** This group is tasked with detailing common ways to replicate, distribute, and store content on servers located in multiple geographical locations at the edge of the network. These standards will assist CDN service providers to offer a greater variety of choice among vendors for their CDN components.

### The Content Alliance

The Content Alliance was formed in mid-2000, and was headed by Cisco Systems, Inc. This consortium first defines and standardizes CDNs technologies,

which will then be implemented in CDNs. The Content Alliance has stated that its standards would be designed to support a variety of business models. The Content Alliance went on to form an internal design team called the Content Peering Working Group.

## The Content Bridge Alliance

The Content Bridge Alliance, founded by Adero Inc. and Inktomi Corp. in mid-2000, is comprised of a smaller group of companies whose members have actively participated in building a CDN. These functioning CDNs handle limited commercial applications and availability by using Content Bridge specified technologies. This group is proofing their CDN business model by offering multivendor CDN services.

### *Streaming Media and CDNs*

One of the main challenges facing content delivery providers is the streaming of media content. This issue is focused mainly on incompatible format and bit-rate requirements for the various platforms that need to stream this media. This platform mismatch can be seen between Windows Media Player, whose servers are usually implemented on Windows NT or Windows 2000, and RealPlayer, whose servers may be Unix based. Applications that use Moving Picture Experts Group (MPEG) applications such as MPEG2 through MPEG4 run on every imaginable platform.

In the past, Webcasts have only been able to present a limited selection of bit rates for consumers, and these may not be able to use the bandwidth at the edge optimally. Some of the reasons why the usage is not at optimal levels include the nature of the connection at the edge, time of day, and overall bandwidth congestion.

This is the impetus for the growth of storage products and services being integrated into CDNs. The products and services can provide more efficient service, as they can store and forward multiple formats and bit rates for streaming media; therefore, CDNs can offer these services.

A number of companies now offer streaming media appliances:

- **Midstream** [www.midstream.com](http://www.midstream.com)
- **Network Engines** [www.networkengines.com](http://www.networkengines.com)
- **Vingage** [www.vingage.com](http://www.vingage.com)
- **Vividon** [www.vividon.com](http://www.vividon.com)

Other services that have sprung up will rationalize the delivery of streams from point-of-origin servers to the edges and are capable of providing various media formats and a variety of bit rates.

One software company that enables these services is AnyStream. AnyStream recently announced that its product, called the Agility Edge, will allow for reencoding of streaming content at the edge as a background service so that the consumer demand for varying formats can adjust to changing network congestion conditions.

Streaming media CDNs will face challenges to their business and technology model, such as the “flash” favoring groups. According to NaviSite, a November 2000 Webcast that featured Madonna involved the creation of the largest FastForward media bridge network in the history of streaming media. NaviSite also noted that traffic spikes were much different from those of on-demand usage of content.

## CDN Solutions from Various Vendors

Several vendors in the market today make what I consider exceptional products for CDNs. The following sections discuss some companies whose products I have seen, implemented, and worked with. These are not the only products in this field, but they are probably the most common.

### Inktomi Content Delivery Suite

Inktomi Content Delivery Suite (CDS) is one of, if not the, leading software solutions for content distribution, data delivery, and data traffic management. This suite contains the Inktomi Traffic Server, that works in conjunction with the Content Delivery Suite to assist in the management of complex tasks such as replication, distribution, and tracking content as it traverses some of the largest and most demanding networks.

Inktomi has a good reputation and has proven its product in real-world deployments. The Content Delivery Suite is able to address all of the challenges associated with content distribution, which will make it easier for you to:

- **Move your content and applications** This allows you to easily distribute your content or application to any type of delivery vehicle such as a Web server, cache, or application server. This includes text and graphics, streaming audio and video, and Enterprise applications.

- **Monitor and report on usage and performance** This will assist in the design and assessment that is necessary in content distribution to maintain and grow, as well as give detailed metrics for service level guarantees.
- **Automatic synchronization of content and applications across server communities** This will create caches on servers so that every user has access to the same information at the same time.
- **Rollback** This will assist in the maintenance of your system by being able to correct content, perform error checking, and audit distribution.
- **Update content and applications in real time** This allows you to “refresh” your content or application on live sites and servers without adversely affecting clients that are accessing the “old” content.
- **Monitor preset thresholds** This will alert you to service issues in real time.
- **Integrate content distribution** This will help you implement content distribution with other areas of the content delivery process, such as traffic management, so that you are able to guarantee that clients always have access to the latest, most accurate content.

Large, distributed service providers are faced with many operational issues. For instance, with the mirroring of sites, many of which are often in geographically distributed locations, the problem of updating and synchronizing multiple servers becomes an almost insurmountable task. Homegrown solutions that built around these technologies, such as Remote File Distribution (RDIST) and FTP, do not scale well and are prone to failure. Just as a little kick in the pants, the rapid pace at which technology and personnel change means that the cost of supporting these customized solutions becomes very substantial.

The Inktomi Content Distributor product was designed to provide a scalable content delivery for Web sites and independent software vendors (ISVs). It contains many features and functions that make sure that your clients are consistently served fresh content in a timely manner.

With the integration of the content distribution process, load balancing, and caching, you will be able to, for example, maintain an application or Web site that only sends clients to a server with the “freshest” information and content. Inktomi Content Delivery Suite is a package that consists of Inktomi Content Distributor and Inktomi Content Manager:



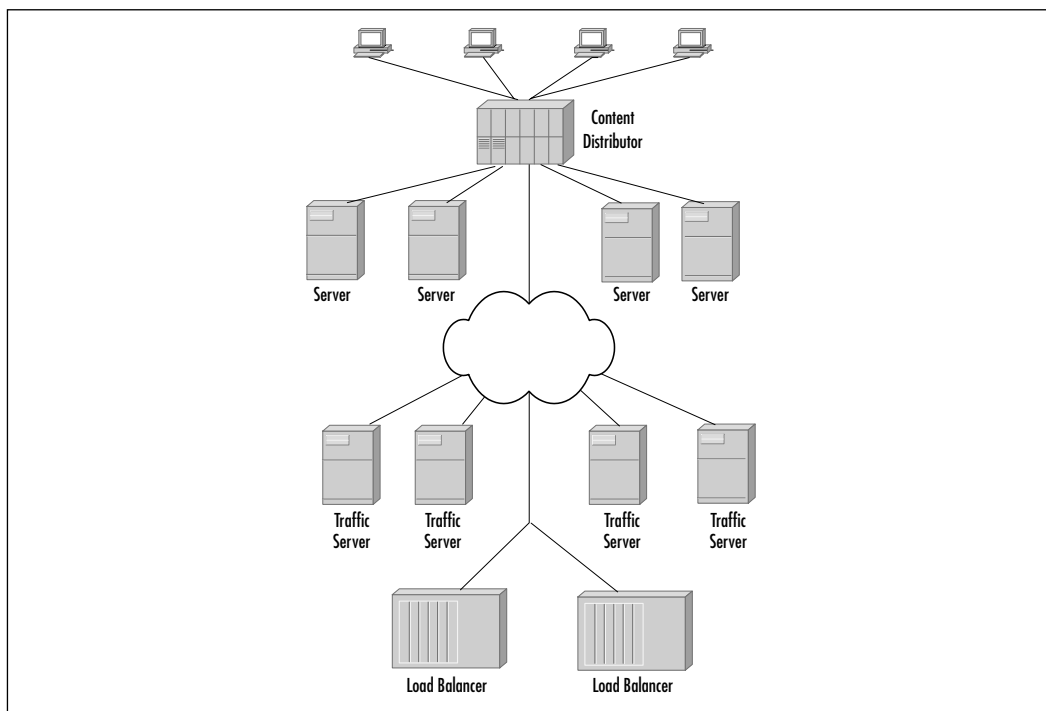
- Content Distributor is designed to replicate and synchronize content and application distribution across multiple network server communities and caches.
- Content Manager monitors, supports, and reports on that content in real time, so you can see how it is performing and being used.

## Inktomi Content Distributor

Content Distributor uses the agent/manager design and a proprietary communications protocol that can replicate content updates to a community of servers over any TCP/IP-based network. It notifies network servers and caches when the content has been changed or updated, and tells them that they need to invalidate old content. Content Distributor can also make available content routing information to load balancers, thus making them aware of servers that did not accept a content update. This will tell the Content Distributor to stop sending traffic to those servers until they are synchronized.

Figure 4.9 illustrates this architecture design. Notice that each server acts as a Content Delivery Suite (CDS) agent.

**Figure 4.9** The Inktomi Content Distributor



Content Distributor also includes a GUI that makes it easier to define and schedule when your content distribution occurs. A CLI is also available, so that there is flexibility and the ability to integrate with the content creation process.

## Inktomi Content Manager

Inktomi's Content Manager uses CDS Agents, which are placed at the distributed servers and caches to capture information on system statistics. It then consolidates this information in real time in a relational database. The Content Manager console can also provide access to the database, and processes the data according to predefined policies. Content Manager also tracks your service level requirements, and allows you to specify what actions should take place if these thresholds are being approached.

## Cisco System's Content Delivery Networks and Next-Generation Content-Based Services

Cisco Systems' Content Delivery Network (CDN) system was developed to help service providers to deploy content delivery services so that they could realize new profit opportunities. With CDN, service providers can augment their users' experience and deliver new services, yet still maintain high availability, add security, and minimize response times.

CDN can help service providers to distribute content "closer" to the client to help overcome issues such as network bandwidth availability, distance, latency, origin server scalability, and saturation issues during peak usage times. CDNs help enterprises speed their deployment of applications such as distance learning and live video and audio streaming.

There are five pieces to Cisco's CDN system:

- Content distribution and management
- Content routing
- Content edge delivery
- Content switching
- Intelligent network services

Cisco's CDN system was designed to handle a wide range of services that include network service providers (ISPs and ASPs), content service providers, Web sites, e-commerce and hosting service providers, and new-model service providers (those that offer applications, voice and video streaming, and storage). With its

structured approach, CDNs can deliver an end-to-end solution or just an individual component that can add value to the existing infrastructure.

Enterprise and Dot.coms can also benefit from CDN by being able to deploy applications including e-commerce, distance learning, online gaming, chat communities, and high-quality streaming media for communication with employees and other businesses.

## Cisco's CDN Group

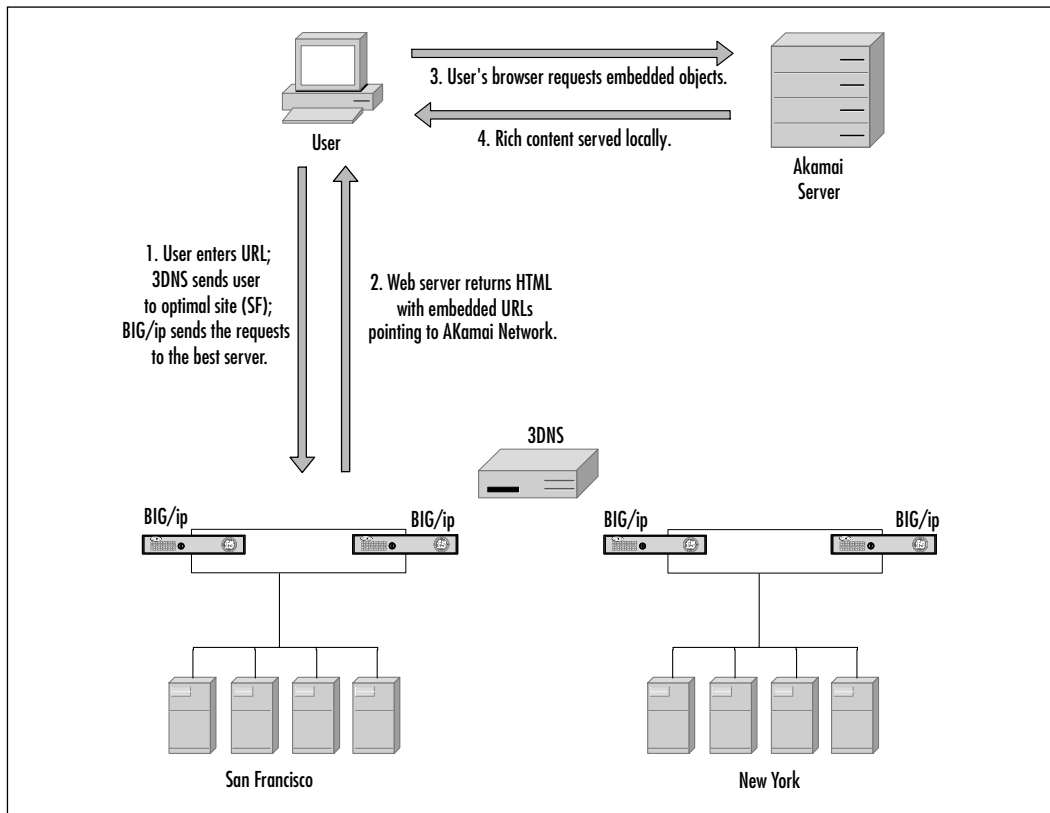
Cisco Systems delivers CDNs that can provide a complete solution that addresses multiple network requirements and situations. Cisco's CDN system is composed of the following:

- **Content distribution and management** This helps to distribute content to nodes that are located at the edge of the network. This allows real-time monitoring, and that will allow you to provide policy settings and centralized provisioning for all delivery nodes within the CDN.
- **Content routing** This redirects client requests to the CDN for maximum scalability and reliability. It is able to handle these requests based on a set of real-time, user-definable metrics that include delay, network topology, current server load, and which policies are implemented.
- **Content switching** This will intelligently load balance traffic across multiple delivery nodes that are located at PoPs (points of presence) or distributed NAPs based on the availability of resources. This level of content switching adds an extra layer of protection against saturation, and maintains transaction connections for e-commerce applications. Content switching also enables QoS-type granularity so that you can prioritize for important content and clients.
- **Edge or access delivery** Allows for the delivery of content from the network to the client. Service providers can define and enlarge the border of their network anywhere from a small number of distributed data centers that are located near the core of the network. This is usually outside of the network edge, and inside the firewall of a client.
- **Smart network services** This area includes network services that are within the IP network, such as security, QoS, VPNs, and multicast traffic. CDN is able to incorporate existing content-aware applications that are required to build scalable and highly available infrastructures.

## Akamai and F5 Networks' Combined Offerings

Akamai and F5 Networks have developed a complementary set of offerings that can provide you with both high speed and reliability for your sites (Figure 4.10). When both solutions are deployed in conjunction, the network is faster for the delivery of content, and there is a guaranteed level of high availability and reliability for that content. How this was accomplished was rather ingenious; Akamai delivers speed, and F5 Networks adds performance and high-availability load balancing. The end product is a leading-edge infrastructure that is able to deliver content quickly and reliably to the client.

**Figure 4.10** How Systems from Akamai and F5 Interact



## Akamai's Solution

Before Akamai's FreeFlow network device was created, the classic state of your everyday Internet site had slow loading time for pages, often had broken images,

and relied heavily on static content. Before FreeFlow, the user would access the Web server, and all available content would then traverse the Internet to arrive at that user's Web browser.

FreeFlow improved this process by instituting global Internet content and application delivery. Akamai enhanced Web performance for its clients by speeding the delivery of content. This was revolutionary, as it did not require administrators to make changes to site layout or browser upgrades.

This vast improvement was created by the placement of copies of the content in caching devices that were much closer to access points. Therefore, a user did not need to pull up the original content that was located somewhere across the Internet; instead, he or she could access the data through a local copy. FreeFlow also provides monitoring of content and applications to keep the local content "fresh" and current with the original server.

FreeFlow also reduced the bandwidth needs of many companies. Internet content delivery services can serve content locally, thereby reducing toll and other long-distance charges.

## F5 Products

Once content delivery was accelerated using the Akamai FreeFlow service, there was a concern that there may not be enough availability for particular sites. Part of this concern stemmed from the fact that if the origin of the content is guaranteed to be available, then content may become stale or inaccessible. Basically, this amounts to users approaching a ramp to a very fast highway, only to find out that the ramp is closed or leads to a dead end.

So, how do you handle the need for high availability with FreeFlow? Many companies need to have a complete high-availability system so that all of their content is readily accessible. Otherwise, this is the weak point, the single point of failure that is the bane of the designer's existence. That's where F5 Networks' products come into the picture.

F5 solutions can be deployed to protect an origin site and guarantee site availability and responsiveness. F5 can also be used to distribute traffic among multiple origin sites, so that there is an ability to have disaster recovery. Remember that you always want to ensure that clients are always able to access your content and applications.

## Summary

In today's complex Web based environments, it is very important to consider end-to-end performance and response time as being the product. The Internet is now carrying increasing loads of mission-critical and bandwidth-intensive multi-media content. There are many factors over which very few Web sites, service providers, or clients have complete (or even partial) control. Service providers need to provide an optimal user experience that can be measured in low latencies and fast download times.

There are various approaches to caching load balancing, and CDNs can be implemented in a variety of ways depending on the specific requirements of the service provider and its clients. When these services are correctly implemented, they can improve the user experience and QoS significantly, while also save service providers significant costs of providing bandwidth that they are then able to pass on to their customers and therefore gain more customers.

The dynamic and delay-sensitive characteristics of mission-critical and real-time content requires much higher levels of Internet QoS than have traditionally been available. In the past, it was acceptable to offer best-effort and no-settlement provisions; however, even network-layer peering is no longer sufficient for the content delivery needs of many Web sites

It is also ideal for enterprises to use caching, load balancing, and CDNs so that they can better manage the usage of network resources. This will also help to provide superior information sharing to employees, and reduce the administrative burdens. These implementations also assist service providers with their advance to managing growth in their infrastructure and connectivity.

The direction that ASPs will be taking in the next few years will likely be toward a multivendor CDN model that will focus on the interoperability and settlement relationships among providers. This will allow for the connection of CDN infrastructure-enabled service providers so that these providers can deploy services to a dispersed set of users in a quicker, scalable, and more economical fashion. CDNs will continue to help ASPs and Web publishers, as they are able to empower businesses and consumers to control and propagate their content in new ways. As we move to this new model, CDNs will evolve to change the Internet into a "pay-for-performance" based environment. This change will then impact the business models of ISPs who, by participating in the CDN value chain, will realize new sources of revenue that are gained from becoming an ISP.

# Solutions Fast Track

## Web Caching and How It Works

- ☑ The intent of caching is to move Web content as close to the end users or the edge of the network as possible for quick access to improve the customers' satisfaction levels, and gives your ASP the competitive advantage.
- ☑ Hardware devices will cache frequently used data and instructions in order to speed tasks.
- ☑ Caching as much Web content as possible within the boundaries of an ISP while using modest amounts of upstream bandwidth is a way to grant clients what they require without creating a “black hole” for bandwidth investment on the part of the service provider.

## Deployment Models for Caching

- ☑ In the forward proxy cache configuration, a client's requests go through the cache on the way to the destination Web server.
- ☑ A transparent cache resides in the flow of the network and is invisible to a client's browser. Clients realize the benefits of caching without reconfiguring the browsers.
- ☑ Reverse cache servers can be deployed throughout the network to create a distributed site of hosted content; this model is commonly referred to as *site replication*.
- ☑ A cache appliance (this can also be called a *thin server*) can be defined as a device that offers a limited number of dedicated functions, and is able to deliver those functions more effectively than a multipurpose device can.

## Load Balancing in Your Infrastructure

- ☑ Load balancing, also called Layer 4–7 switching, occurs when cluster of Web servers are created to handle massive amounts of requests.
- ☑ Localized load balancing occurs when the load balancer determines which server should receive new requests.

- ☑ Distributed load balancing sends packets across dispersed networks, which can be located in geographically separate areas from the local server.

## Load Balancing Solutions from F5

- ☑ As more servers are added to the DNS round-robin rotation, traffic will be unevenly distributed. The older servers will tend to receive more traffic than newer servers, as the IP addresses of older servers are usually cached by more users than the addresses of newer servers are.
- ☑ When you implement a network device that is capable of high availability, you want it to guarantee that it can deliver IP-based services, which are always available. To do this, you must remember that it is imperative that both “quality of service” based high availability and load balancing are addressed so that your client has a good usability experience.

## Cisco Systems’ LocalDirector

- ☑ There are generally two approaches for scaling a server farm-based system. The first approach is to continuously upgrade the size and processing power of individual servers in the farm. The second approach is to add more servers as you require more capacity.
- ☑ Load-balancing technology does not normally consider variables such as bandwidth, server performance, and job size for optimizing the traffic loads among your server farms. Load balancing can allow you to incrementally scale the capacity of servers in your server farms in a more efficient manner.
- ☑ LocalDirector is considered a transparent device, as it is able to work with any TCP-based service or application. There is no special software required on the server, as these are external devices.
- ☑ The LocalDirector is considered a stateful device, as it is able to monitor and can track all TCP connections that are occurring between clients and servers.



## Foundry Networks' ServerIron

- ☑ Foundry's ServerIron Web switches provide high-performance content and application-aware traffic and server load balancing. ServerIron has the functionality of a traditional Layer 2 and Layer 3 switch built in, and is able to examining the content at Layer 4 and above through the packet header.
- ☑ ServerIron load-balancing characteristic is based on Layer 4 traffic such as HTTP, FTP, SSL, and email. This creates the ability to transparently distribute data traffic among multiple servers.

## Content Delivery Networks

- ☑ The networking industry's focus from Layer 3 connectivity issues is shifting to the creation of intelligent, Layer 4–7 networks that can support the rigorous response-time requirements of these new types of content. The emphasis is now turning to content delivery networks (CDN).
- ☑ CDNs are able to provide QoS to the Internet's IP-based backbone, which helps to eliminate or minimize delay.
- ☑ Content provider organizations build content for the Web, and are faced with delivering content that has dynamic characteristics to customers who require high levels of service.

## CDN Solutions from Various Vendors

- ☑ Content Distributor uses the agent/manager design and a proprietary communications protocol that can replicate content updates to a community of servers over any TCP/IP-based network.
- ☑ Cisco Systems' Content Delivery Network (CDN) system was developed to help service providers to deploy content delivery services so that they could realize new profit opportunities.

# Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the “Ask the Author” form.

**Q:** What are some major causes of bottlenecks on the Web?

**A:** Network link congestion, server congestion, network equipment congestion, and distance delay.

**Q:** What are some of the solutions that can accelerate Web performance?

**A:** Geographic distribution, content replication, caching, redirection, and load balancing.

**Q:** What are the limitations of using local server farms to improve performance?

**A:** For simple implementations where DNS round robin is used and a single URL is mapped to multiple Web servers, the status of each server is unknown by all other servers, and users could be sent to a server that is not working or is congested, resulting in an even poorer Web experience. Using local server farms is also susceptible to network connectivity outages, bringing down a Web site. Finally, the solution addresses only congestion associated with the central site; it cannot address delays associated with network link congestion, network equipment congestion, or distance delay.

**Q:** What requirement criteria should I use to select the right load-balancing product?

**A:** Criteria you should consider include dependability, Quality of Service, and availability.

**Q:** What are important things to look for in content-delivery products?

**A:** Important features that need to be considered include a version control feature, site recovery and rollback capabilities, scheduled publishing, logging features, and built-in security features.

**Q:** What types of applications are ideally suited for development as traffic server extensions?

**A:** Content filtering, content transformation, software-on-demand, media delivery, content personalization, analysis/monitoring, and compression.

**Q:** What major benefits can CDNs bring?

**A:** CDNs can improve the end-user experience, respond to network growth and dynamic changes, offer high availability for both network and content access, support a variety of new service provider business models, and enable service providers to build value with CDN offerings or use CDN technology to augment their hosting service offerings.

## Storage Solutions

### Solutions in this chapter:

- Upfront Concerns and Selection Criteria
- Directly Attached Storage in Your Infrastructure
- Network Attached Storage Solutions
- Storage Area Networks
- Scalability and How It Affects Your Business
- Fault Tolerance Features and Issues
- SAN Solutions Offered by Various Vendors
  
- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

# Introduction

Within the last decade, we have seen a complete transformation in computing technology. The Internet has helped shape our current view of business, and with these new businesses, the need for high-tech data storage. These days, applications are requiring more storage space, and simply placing data on a server or workstation's internal hard drive has become a thing of the past—it is almost considered archaic.

According to the research firm International Data Corporation (IDC), in order for a company to keep up with demand, storage capacity will need to double each year for the next couple of years. The global network storage industry is expected to triple its capacity to a whopping 2.2 Exabytes (an Exabyte is about 1 million terabytes) over the next two years.

In today's environment, there is a very real need for quick and reliable access to data from various integrated systems. With the ever-increasing amount of data that is stored on systems, it is becoming progressively more complex to perform routine backups and handle the maintenance of hundreds or possibly thousands of systems without introducing at least a modicum of system downtime and/or network congestion.

Whether you are an Internet service provider (ISP) or an application service provider (ASP), keeping your systems running efficiently around the clock should be one of the most important goals of your organization. Designing your data storage systems properly means that you need to develop reliable, cost-effective solutions. This will provide proof to your customers that your organization truly can offer them value-added services, and is genuinely concerned about their needs.

As I stated earlier, the need for reliable data storage and cost-effective solutions has grown at an exponential rate. This has forced the industry to reassess its approach to data storage. Some organizations have turned to large storage arrays that are centrally maintained, others have opted for clusters of network attached storage servers (NAS), and still others have decided to build the most state-of-the-art storage system available today, the storage area network (SAN).

This chapter is written to help you cut through all the confusion that surrounds SANs and NASs, in order to clearly define and discuss some of the storage issues as they pertain to an ASP. We will start with an explanation of some of the most basic storage methods, such as server-attached Redundant Array of Inexpensive Disks (RAID) arrays, and NASs, and move on to the complex workings of SANs. Since scalability and fault tolerance are also major concerns that an ASP must address, I have provided separate sections that discuss some of the potential issues in some detail.

# Upfront Concerns and Selection Criteria

Currently, there are many differing manufacturers of storage-based equipment, and several methods of delivering storage solutions to your servers and clients. All these pieces of equipment and options range greatly in price, performance, manageability, and features offered. It is very easy to become overwhelmed by the choices available, so I will try to give you some options that will help you make a wise decision that will not cost your organization incredible amounts of money to implement and maintain (although my definition and your definition of “incredible amounts” may differ).

## Concerns for Your Storage Devices

Having to replace a failed implementation with a different solution is a tremendous waste of time and resources, so let’s look at some of the criteria that will assist us in implementing the proper solution the first time. Keeping the issues and concerns that follow in mind will help you make wise, well-planned decisions about your storage solutions.

Six major concerns and criteria should be taken into account before deciding on the storage solution that best fits your requirements, and we’ll discuss them all in some detail. These concerns, in order of importance, are as follows:

- Host independence
- Mixed vendor support
- Security
- Legacy support
- System availability
- Price versus performance

## Host Independence

Some manufacturers design their storage equipment so that it relies on software that is placed on each host to facilitate access to storage devices. While this may work for small implementations, it can quickly become overwhelming in a large service provider environment. Remember that as the number of hosts grows exponentially, so does the complexity of maintaining and managing these systems.

Imagine having to install and maintain software on hundreds or possibly thousands of host systems that reside within your infrastructure and your client’s

environment. The amount of time and money spent on installation alone could grow to astronomical levels, not to mention the associated recurring maintenance costs. Since there are so many types of hardware, software packages, and network operating systems (NOSs) available, you may have a system that different vendors may not support. If this is the case, your storage solution may not work with some of your hosts, and that doesn't sound like a good solution at all.

## Mixed Vendor Support

It is much more advantageous to standardize your environment on a particular vendor's equipment. Doing so can help to minimize the time and resources spent on training your staff to implement, manage, and maintain the infrastructure. There is only a single product to learn, instead of several products that are configured and operate significantly different from each other. In some cases, it may even reduce the amount of time spent troubleshooting equipment, since it could lead to a higher level of familiarity and deeper understanding of a particular manufacturer's product.

Most companies feel that this tendency for standardization is better than diversification in almost every instance. There are, however, instances where this can be a drawback. There are some disadvantages; for example, if a product offering or device is proprietary, it will probably be very difficult, if not impossible, to change or upgrade in the future. In some rare cases, a vendor may not stay in business, leaving us all hanging out to dry and raising numerous support problems and issues.

Because of this, it is important to be concerned about vendor 'lock in' and attempt to plan for future growth and expansion instead of short-term comfort and cost savings. With mass-storage products, some of the major manufacturers may only offer proprietary equipment, while others may standardize their equipment, using a technology such as fiber channel to ensure that their product will work with a similar offering from another manufacturer. For these reasons, it is always important to know what you are purchasing, and whether it will successfully fit into your long-term business model.

## Security

Security should always be a concern, but it is especially important given the high visibility of ISPs and ASPs. Based on the sensitive customer and internal data that is typically stored on their systems, we can not stress enough how important security concerns should be in your storage criteria and design decisions. When thinking of security for your storage solution, there are generally two different methods:

- Host-based security
- Outboard security

### *Host-Based Security*

Host-based security is exactly that; the individual host device will handle the security functions for that equipment. There are some concerns like those we mentioned earlier that are just as applicable when considering security, as there is another level of complexity that could impose even more severe ramifications on your design and implementation. If you plan to use host-based security for your storage network, beware of the host attack.

Since the host device will handle its security exclusively, should it happen to be compromised, there is little to no security preventing it from accessing any slice of data in your storage network. If the host has been compromised by an attacker, or has become a rogue for any number of reasons, it may have full reign and access to the entire storage network. It may also have the ability to read and write to any storage device in the entire pool of systems.

As you can see, this could prove to be a disaster of very large magnitude, especially if it is not caught immediately. For these reasons, we do not recommend using solely host-based security solutions. Instead, it would be prudent to use host-based security in addition to some form of outboard security.

### *Outboard Security*

Outboard security is any type of security feature that is located on the host. It might be an external authentication scheme that is provided by a firewall. A firewall is hardware or a software package that performs real-time security features and monitoring, or even security provided by each individual storage device.

Whatever the case may be, outboard security offers the best level of protection for your storage network by providing a method of centralized access control. In most cases, such a solution will help to reduce maintenance time and the associated costs because the system is also centrally administrated. Even greater may be the ability to audit security trails, adding to the overall sense of security that will allow your staff and customers to sleep well at night. Alone, or in conjunction with a host-based solution, outboard security is really the only way to go for sensitive data that needs real protection.

## Legacy Support

You may already own storage devices that use interfaces other than fiber channel, such as small system computer interface (SCSI) or enhanced integrated drive



electronics (EIDE) for host connections. It can sometimes prove difficult to port older hardware to some newer storage solutions. In particular, it may be difficult to use some devices in a SAN environment, because they may not incorporate the protocols and technology to integrate into the network.

If you would like to retain the ability to continue to use this equipment in your network, you may need to look into particular product offerings to provide this functionality. Fiber channel routers or bridges could be used to allow for functionality between some of these devices, but this may make the overall design quite complex and complicate administration. Instead, it might be more prudent to look for devices that offer a wide range of flexibility and work with your existing hardware. Whether you choose one of these options or decide on brand new equipment that offers more flexibility, it is always important to design a system that is simplistic in nature and always transparent to your end users.

## System Availability

System availability (also known as uptime, redundancy, or high availability) should be a concern whenever purchasing equipment that is mission critical. It is important to look for redundancy in your network as well as in the individual device. High availability might mean having two of every device throughout your network, as well as two possible paths in the network in case one fails. It could simply mean looking for redundant power supplies or network connections for specific devices.

The decision is ultimately yours, and the level of redundancy you require depends on your expectation of server and application uptime, and network connectivity. As an ASP, your services are your business lifeline, so you should always try to identify single points of failure and build solutions to overcome these potential issues. Because of the huge importance of fault-tolerant systems in your network, we'll discuss this topic in further detail later in the chapter. For now, it is important to remember to look for redundancy whenever purchasing equipment and designing your solutions.

## Price versus Performance

Another factor to consider is the cost versus performance aspect of your storage scheme. Obviously, you will want to shop around and compare the prices and features that each device will bring to your infrastructure before you decide to purchase a particular device. However, there is sometimes more than meets the eye.

For instance, some manufacturers may use custom or proprietary hardware and software to provide storage services that may not operate well with other

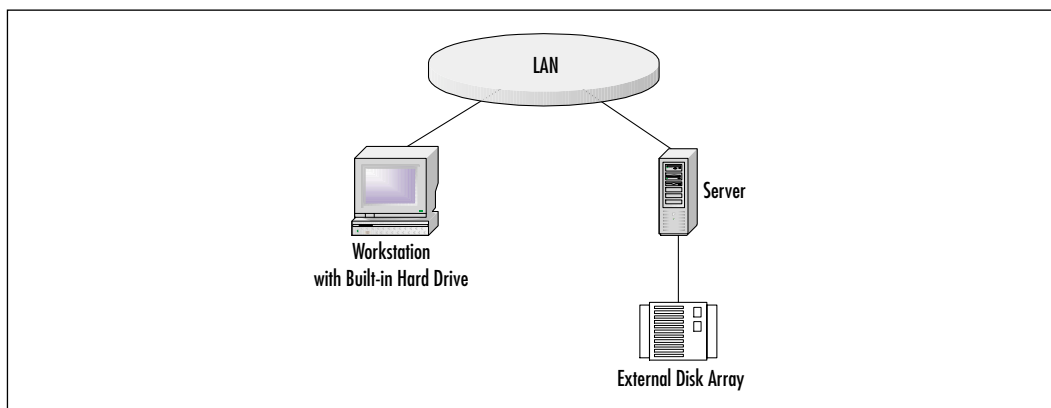
devices or hosts. In order to integrate their platform into your environment, you may need to spend a great deal of money developing and testing their product to ensure that it interoperates correctly with your entire solution.

If it does not integrate well, you may even be forced to purchase other devices or develop new code that allows for better functionality. In some instances, you may even need to rewrite your applications to allow them to access a particular device. At the end of this ordeal, you will probably have spent far more money developing a solution that functions correctly than you would have spent if you had chosen a slightly more expensive but more flexible storage solution. This being the case, you should look for proven solutions that leverage high-performance technologies to provide an upgradeable, extensible, manageable, and cost-effective solution.

## Directly Attached Storage in Your Infrastructure

Server-to-storage access, or directly attached storage, has been in use in much of the history of computing, and still exists in over 90 percent of implementations today. An example of server-to-storage access, as shown in Figure 5.1, could be a workstation that has an internal hard drive, or a networked server that has an external disk array directly attached to it.

**Figure 5.1** Directly Attached Storage



In these network implementations, storage devices are directly connected to a server using either interfaces and/or bus architecture such as EIDE or SCSI. In more recent implementations, it is common to find newer devices that use fiber

channel to directly attach to a server. Regardless of the method used to connect these devices, they are all the same in architecture; a server or host is directly connected to a storage device using a storage bus.

This is not a very flexible model with which to work. Given that some hosts may require more storage space than others may, it is very difficult to move capacity from one server to another. To do so, you would actually need to remove hard drives from one storage array or device and install them in another device when that device needs more space. Even with this solution, you may run out of physical space in a storage array, and need to attach an additional array of disks.

All of this “upgrade” would require the reconfiguration of the storage device and host systems, and would obviously become quite cumbersome and time consuming. In addition to these drawbacks, performance is limited completely by the directly attached server’s abilities and the central processing unit (CPU).

For instance, if a server is too busy doing calculations for other applications, it will have to wait or free up valuable CPU clock cycles in order to read and write from the storage device. This will impair its application and input/output (I/O) performance significantly. This may be acceptable for someone’s personal computer, but in a mission-critical, performance-impacted business environment, it can prove to be a serious problem with severe consequences and limited options.

## Network Attached Storage Solutions

Network attached storage (NAS) is one of the latest solutions to hit the streets. When you hear someone talking about SAN, you usually hear “NAS” in the same sentence. While they both provide methods for accessing data, and resolve many file access issues when compared to traditional methods such as directly attached storage, in practice they differ significantly.


A NAS is a device that provides server-to-server storage. What does this mean? The answer is simple: It means that NAS is basically a massive array of disk storage connected to a server that has been attached to a local area network (LAN) as depicted in Figure 5.2. In fact, it is very simple, and means exactly what it states.

As an example, imagine a host accessing data on a NAS server. The actual data is transmitted between these devices over their LAN interfaces, such as Fast Ethernet, using a communications protocol such as Internet Protocol (IP) or Internet Packet eXchange (IPX).

With the existing network infrastructure, the communications protocol might also allow data to be transmitted between hosts that are extremely far apart. For instance, a personal computer might access data on a file server that is thousands

of miles away by using the existing network infrastructure of the Internet, or a customer computer might mount a drive on a remote server over a private wide area network (WAN) connection such as a T1. In both of these cases, the server being accessed is, for all intents and purposes, acting as NAS.

**Figure 5.2** Network Attached Storage



This can provide a great solution for applications, and will more than likely be the method most of your customers will use to connect to data that resides on your systems. It offers quite a lot of flexibility and requires very few upgrades to your network infrastructure. We already discussed the best benefit of this type of architecture, but it bears repeating it here: you can use your existing network infrastructure for accessing data that resides on NAS servers.

There can be some serious drawbacks that are inherent to this solution, though. Probably the most important is the impact that such an architecture will have on your LAN and WAN. When we talk about sharing data, we might mean terabytes of data. Using a NAS device can easily bottleneck your network and seriously impact some of the other applications within your network.

I do not want to scare you away from this architecture, because it is still a very viable and robust solution. In fact, when connecting hosts or servers to data over very long distances, it is still a very good solution, and sometimes the only option available. Many of your customers will more than likely already have an existing connection into your network, so it becomes easy to add services with

very little impact on your other clients. Some methods can be used to help eliminate the impact that a cluster of SAN devices might impose on your network.

## Quality of Service

You can combat network performance problems by designing Quality of Service (QoS) into your network. In fact, we recommend using QoS throughout your network, even if you decide not to use NAS. QoS has the ability to delegate priority to the packets traversing your network, forcing data with a lower priority to be queued in times of heavy use, and allowing for data with a higher priority to still be transmitted.

A well-designed and implemented QoS schema can definitely help eliminate the impact that large volumes of data may have on other time-sensitive data, but it could still expose your network to a level of latency that is capable of growing exponentially. This is especially true if you do not plan correctly. When designing QoS in your network, it is very important to look at all the data traversing your network, and carefully weigh the advantages and disadvantages of using a particular QoS strategy and its effect on types of data and the network as a whole.

## Location of NAS in Your Network

When designing NAS in your network, probably the most effective solution for latency and saturation issues is the location of your NAS servers in relation to the hosts and systems that access their data. The placement of NAS devices becomes extremely important, and performance can vary significantly depending on your design.

For instance, if you have a single large cluster of NAS devices in the middle of your network, all hosts will need to traverse deep into your network in order to access the servers and data. Consequently, you will have large amounts of data flooding every part of your network that will more than likely create serious bottlenecks and latency issues at every step along the way.

In contrast, if you were to use smaller clusters of SAN devices, and locate these groupings close to the hosts that access them, the hosts will not need to traverse your network to access the NAS servers, thereby keeping network saturation to a minimum.

Unfortunately, there is no clear and concise way to design NAS in your network. Your ultimate design will depend greatly on your current and future growth patterns. As a general rule, remember that NAS devices should always be kept as close as possible to the devices that access them. However, always keep

their purpose in mind, as well as who will be accessing the data, patterns of usage, and the costs associated with distributing these systems.

In some cases, you may have very few clients accessing the data, or saturation may prove to be the downfall of your network or a nonissue. However, when comparing price versus performance issues, try to keep your projected future growth in mind, as it can significantly alter the decision-making process.

## Storage Area Networks

A storage area network (SAN) is a networked storage infrastructure that interconnects storage devices with associated servers. It is currently the most cutting-edge storage technology available, and provides direct and indirect connections to multiple servers and multiple storage devices simultaneously.

With the use of technologies such as Fiber Channel, the SAN actually extends the storage bus, thereby allowing you to place servers far away from the storage devices that they access. In fact, the servers may be housed at locations that are completely separate from the site housing the storage. In this situation, we would be taking advantage of one of the greatest features that SAN technology provides.

A SAN can be thought of as a simple network that builds off the familiar LAN design. Instead of connecting hosts with other hosts and servers, it is designed to connect servers and hosts with a wide range of storage devices. A SAN uses network hardware that is very similar to what can be found in a typical LAN, and even includes the use of hubs (very rarely), switches, and routers. In its most basic form, it could be thought of as a LAN that is dedicated solely to accessing and manipulating data.

## The Need for SAN

There are several scenarios behind the move to storage area networks. The major one is the need to manage the dramatically increasing volume of business data, and to mitigate its effect on network performance. The key factors include:

- **E-business** Securely transforms internal business processes and improves business relationships to facilitate the buying and selling of goods, services, and information through the Internet.
- **Globalization** The extension of information technology (IT) systems across international boundaries.

- **Zero latency** The need to exchange information immediately so you can maintain a competitive advantage.
- **Transformation** The ability to adapt, while maintaining the ability to immediately access and process information that drives successful business decisions.

Distributed computing, client/server applications, and open systems give today's enterprises the power to fully integrate hardware and software from different vendors to create systems tailored to their specific needs. These systems can be fast, efficient, and capable of providing a competitive edge.

Unfortunately, many enterprises have taken a far less proactive approach with their storage systems. Storage, unlike a Web application server or a database system, is rarely viewed as a strategic tool for the enterprise; this view, however, is beginning to change.

With the explosive growth of e-business, IT managers are working very hard to keep pace with managing the significant growth of data (multiple Terabytes, if not Exabytes, per year). They are installing high-performance storage systems to meet the demands for smaller backup windows and greater application availability. However, these systems are sometimes much more complex and expensive to manage. In addition, they are often single platform, restricting access to data across the network. To improve data access and reduce costs, IT managers are now seeking innovative ways to simplify storage management, and SAN is a promising solution.

## Benefits of SAN

SANs remove data traffic—backup processes, for example—from the production network, giving IT managers a strategic way to improve system performance and application availability. Storage area networks improve data access. Using Fiber Channel connections, SANs provide the high-speed network communications and distance needed by remote workstations and servers to easily access shared data storage pools.

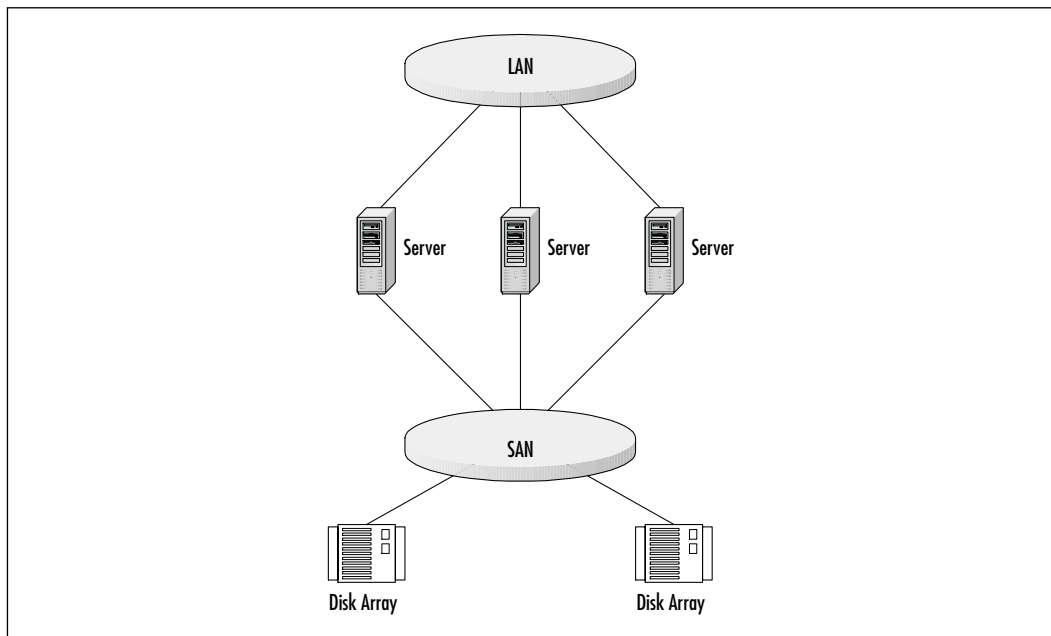
IT managers can more easily centralize management of their storage systems and consolidate backups, increasing overall system efficiency. The increased distances provided by Fiber Channel technology make it easier to deploy remote disaster recovery sites. Fiber Channel and switched fabric technology can help eliminate single points of failure on the network.

With a SAN, virtually unlimited expansion is possible with hubs (again, very rarely) and switches. Nodes can be removed or added with minimal disruption to the network. By implementing a SAN to support your business, you can realize:

- **Improved administration** Consolidation and centralized management and control can result in cost savings. By allowing for any-to-any connectivity, advanced load-balancing systems and storage management infrastructures, you can significantly improve resource utilization.
- **Improved availability** With a SAN, high availability can be provided more effectively at lower cost.
- **Increased business flexibility** Data sharing is increased, while the need to transform data across multiple platforms is reduced.

One of the main advantages of owning and operating a SAN is that it offers a secondary path for file transfers, while keeping the LAN free for other types of data communication. Figure 5.3 shows that the SAN is a separate network from the LAN, and truly provides a secondary path for file transfers.

**Figure 5.3** SAN and LAN





## SAN Virtualization

In order to design the SAN, we must first consider what servers are going to access the actual data that resides on the physical disks. In most instances, there are numerous storage devices clustered together. In these cases, we will need to create a method of accessing and storing data that might reside on several different storage devices, which in many cases may be from different manufacturers. Essentially, we will need to “virtualize” all these devices into a single logical pool of devices. This is known as SAN virtualization.

SAN virtualization products collect all or portions of the physical disks into a single group of resources. This group is then subdivided into logical slices that can be easily accessed by the appropriate servers. To understand how SAN works, it’s necessary to explain the different methods for providing virtualization functionality.

There are basically four divergent virtualization schemes:

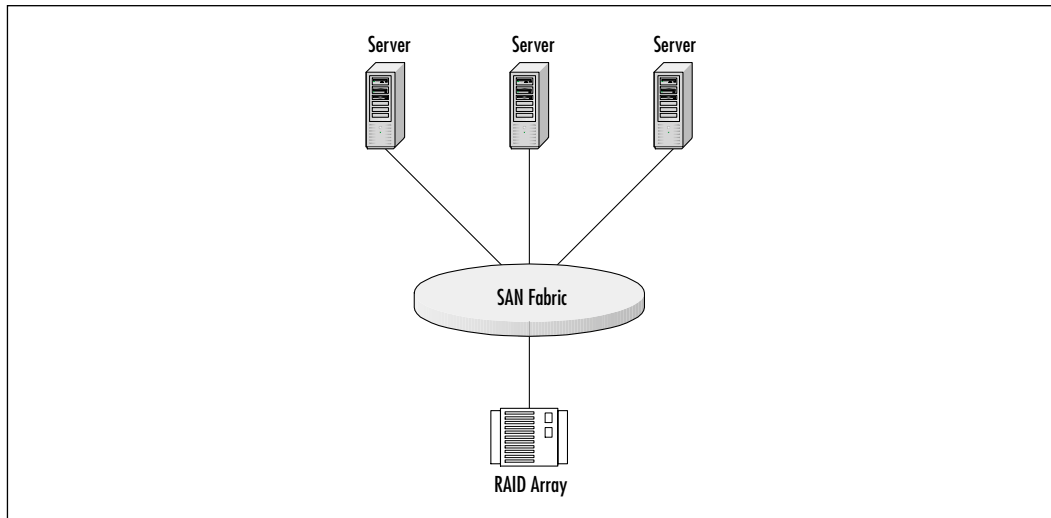
- Multihost arrays
- Logic unit number (LUN) masking
- In-band virtualization
- Storage domain servers

### Multihost Arrays

Multihost arrays are the most simplistic and most common form of SAN virtualization implementation. They put all of the pooling responsibility on the storage subsystem level. This is done using a Redundant Array of Independent Disks (RAID) controller to slice the pool of drives into logical units. Figure 5.4 depicts a multihost array in its most basic form.

In this example, three servers are connected to a single storage device. The servers communicate with the storage device through the SAN *fabric*. The storage device can be broken down into individual hard drives that can be lumped into one single or many logical units of storage space with the use of RAID technology. Administration of this solution is generally accomplished by attaching a computer directly to the array, or by remotely administrating the device through the LAN.

This type of architecture offers a high-availability solution with very good performance, and supports connectivity to numerous types of hosts and operating systems. It does very little to ease security concerns, though, since every host connected to the storage device has full access to the raw data contained within the logical pool.

**Figure 5.4** Multihost Arrays

Upgrading the storage capacity of the system requires the installation of additional disk drives into the enclosure, limiting your upgrade capabilities to the actual form factor or size of the device's chassis. When additional space is needed beyond the capacity of a single device, you will be forced to install a second array, create a second pool, and attach these to your SAN (this is sometimes referred to as *clustering*). This could make for a complex system that limits your centralization and allocation freedom.

## Logical Unit Number Masking

Logical unit number (LUN) masking is a method of providing additional security on multihost arrays. To provide this extra bit of security, a special device driver containing an individualized LUN mask is installed on each of the host systems. The LUN mask is capable of denying the host access to resources that are pre-programmed as forbidden.

Although some packages allow administration to be accomplished centrally, it is still a host-based solution that makes administration complicated and time consuming. In addition, since this solution is also software based, it is sometimes difficult to find support for every platform and operating system, and upgrades may become outdated if new device drivers are not released or routinely maintained.

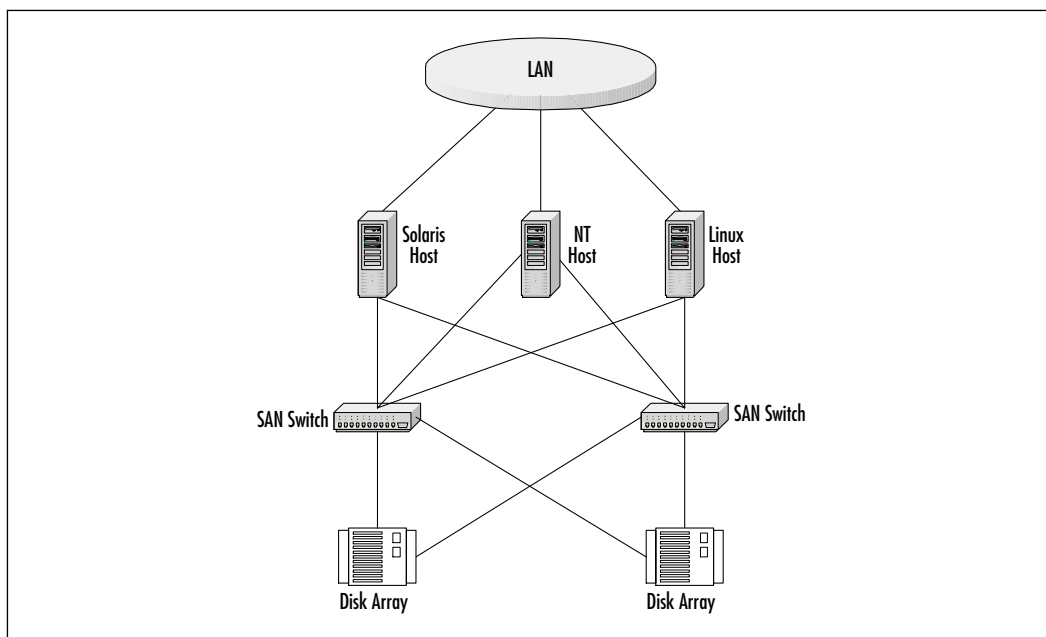
Although this solution was designed to provide additional security, it is far from foolproof. Given that security is controlled by a software element on the host, it is not safe to assume that the hosts do not have access to the supposedly

forbidden data. A malfunctioning or compromised system will still have the ability to cause harm to any portion of data contained in the virtualization pool. If the LUN mask device driver is removed from the system and a new driver that does not contain the LUN mask is installed, there will be nothing restraining the host from accessing data that is supposedly “forbidden.”

## In-Band Virtualization

In-band virtualization refers to a dedicated device or devices, such as a SAN switch or specialized virtualization engine that sits between the hosts and storage arrays to perform all the virtualization functions. Because virtualization takes place in between the hosts and storage devices, it is easy to build a decent amount of high availability into this type of solution. Figure 5.5 is an example of an in-band virtualization solution that provides high availability and operates with multiple network operating systems (NOS).

**Figure 5.5** Highly Available In-Band Virtualization



Since it is a dedicated platform, there is no host software to install and maintain, thus reducing the amount of time and resources used when configuring and maintaining the individual hosts. Platform support offered by these devices can

vary, but is usually very good since they are generally a hardware-based solution and do not require specialized software on the hosts or storage devices.

The products offered range incredibly in price and functionality. Some products are merely software solutions that run on a particular platform and require additional switches and hardware. Other products offer all of these features in a single device. If you find yourself looking into in-band virtualization products and feel overwhelmed by your choices, make sure to keep the concerns listed earlier in mind.

When designing a large SAN, take extra care and look closely at the features offered by these devices. Some of the products will offer improved caching engines that allow your SAN to operate with a good deal of speed and efficiency. Other products might actually hamper your network performance and scalability, depending on their implementation and system requirements.

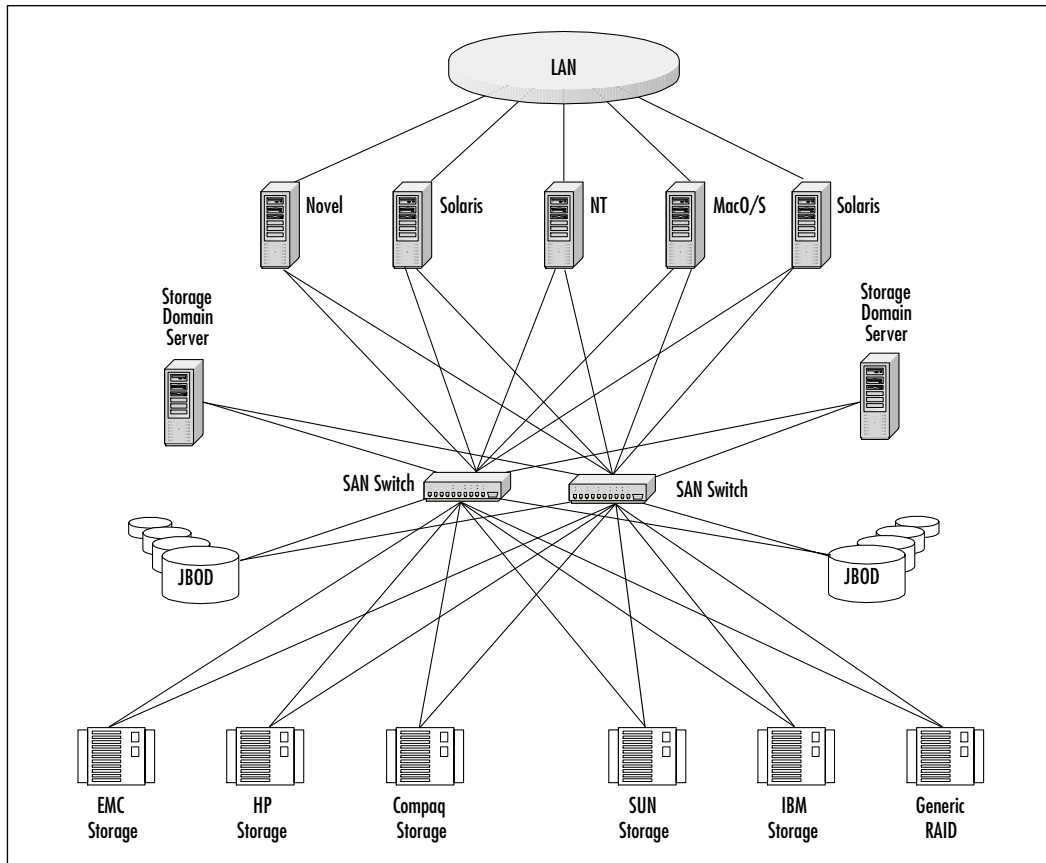
Just as with your LAN, look for fault-tolerant capabilities and solutions, and be careful not to introduce single points of failure into your SAN. Planning and configuring alternate paths throughout your SAN will help to provide continuous availability during device failures or malfunctions.

## Storage Domain Servers

Storage domain servers are dedicated servers that run virtualization software on top of a commercial operating system such as Microsoft Windows NT/2000 or Unix. This type of solution builds off the in-band virtualization solution, and provides for better volume management, interoperability, and security features. These servers permit an architecture that maintains centralized administration, while distributing management functions with other devices within the infrastructure.

Storage domain servers typically offer increased caching and mirroring support, which can help to enhance the speed and fault tolerance of your SAN. These servers also provide the foundation for other advanced features, including services such as server-free backup, which can allow data to be archived without using host devices, thereby freeing up valuable CPU cycles to improve operating efficiency and significantly reduce data backup windows.

Storage domain servers promise to deliver enhanced data services that are unrivaled by other virtualization solutions. Figure 5.6 is an example of a fault-tolerant storage domain server solution that will work with most operating systems, and the majority of storage devices, including groups of disks that have not been configured using RAID, such as just-a-bunch-of-disks (JBOD) storage.

**Figure 5.6** Storage Domain Servers

## NAS versus SAN

Network attached storage (NAS) grew out of the concept that file servers can be used as a service to manage files for customers and their infrastructures. The file server approach became successful due to products such as Novell NetWare and Microsoft Windows NT Server.

With a file server in the infrastructure, large amounts of data storage could be attached to a server and then disseminated to users on a file-by-file basis. As a side benefit, management and backup of that data could be centralized on different servers. Through the course of time, it became more evident that an entire network operating system (NOS) was not necessary to handle file services.

Therefore, a trimmed-down version of the NOS, called a storage appliance, was designed to work with specialized servers. These storage appliances could be

installed within a network to provide the storage for the infrastructure. This is the concept that grew into the NAS market, which is dominated by network appliance and expected to grow to more than \$6 billion by 2003.

SANs, on the other hand, are based on the concept of taking storage devices, and traffic that is storage-heavy, and creating a separate network designed specifically for that type of data traffic. By separating a server from its storage, and placing all the storage devices directly on a network (a Fiber Channel network for this instance) allows many-to-many connections from various servers to storage, and from these storage devices to other storage devices.

By implementing SANs, you will receive the benefits of traditional networking to storage devices, such as increased scalability, availability, and performance. Backups can then be done without affecting the rest of a network, as backup traffic is performed on a separate SAN.

SANs initially have traditionally been based on Fiber Channel Arbitrated Loop (FC-AL) architecture. Many users have implemented SANs for multiple servers to one or two SAN-attached RAID arrays, and in some cases, a tape library.

The Fiber Channel market today revolves around switch vendors (e.g., Brocade, McData, Qlogic, and Vixel) and storage/server vendors (e.g., Compaq, EMC, Hewlett-Packard, IBM, and StorageTek).

The definition of SANs has expanded to include other technologies such as Gigabit Ethernet and SCSI over IP. This has made it possible for new companies such as Gigaset and Nishan Systems to come about. Clients today who talk about their internal SANs are, in some cases, actually describing an environment of NAS boxes communicating across an Ethernet network. Some argue that this does not constitute a SAN.

## Comparing Fiber Channel to SCSI

An all-Fiber Channel SAN offers many performance and administrative benefits over SCSI-connected storage. Multiple users can combine Fiber Channel storage devices with legacy SCSI devices in a SAN environment by using Fiber Channel-to-SCSI bridge products. Both of these implementations have associated benefits and shortcomings.

A SAN is, generally speaking, a shared storage repository. SANs aren't necessarily synonymous with Fiber Channel; in fact, many of these storage devices are RAID based. The channel network can be Fiber Channel, Enterprise Systems Connection (ESCON), or even SCSI. The I/O channel operates on the backend of the server, making file access and data transfers that are independent from the LAN.

This combined high-speed channel can then be scaled up to more sophisticated versions that enable a conglomeration of storage devices (for example, RAID and tape libraries) to communicate over common high-speed, fault-tolerant storage pipelines with multiple hosts. This also allows the devices to connect directly with each other.

As new software and fabric enhancements become available, SANs will be able to support increasingly complex server-to-storage functions such as fault-tolerant access paths, automatic failover, dynamic reallocation of storage devices, the assignment of dedicated storage in a device (also called *logical unit masking*) for hosts and/or operating environments, and high-availability clustering services.

A SAN is composed of three basic components:

- Interfaces (such as SCSI, Fiber Channel, or ESCON)
- Interconnects (routers, switches, gateways, or hubs)
- A protocol (e.g., IP or SCSI) that controls data traffic over access paths that connect multiple nodes

These components, in addition to the attached storage devices and servers, can form what can be considered an independent storage area network. While the SAN is able to support multiple interfaces, Fiber Channel (Fiber Channel Arbitrated Loop and Fiber Channel fabrics) has gained the greatest market in this arena due to its flexibility, high throughput, and fault tolerance.

## The Benefits of Fiber Channel

In an FC-AL-based SAN, you can have up to 126 nodes attached per loop. Switched fabrics are able to currently support up to 16 million device addresses. This modular scaling capability can provide a solid infrastructure and allow for long-term growth. The Fiber Communications Channel is able to support multiple protocols, but has a current bandwidth limitation of 100 Mbps. The Fiber Channel interface can sustain this bandwidth up to 10 kilometers by using long-wave optical interconnects. Fiber Channel switches can also be cascaded to provide remarkable capacity and performance scalability.

In addition, Fiber Channel can be configured rather easily for high-availability environments. Today, most Fiber Channel disk drives are dual ported. This allows for the use of both ports in a dual-loop configuration to provide redundant paths to and from devices, thus guaranteeing access to the device or node if one path should fail. Fiber Channel switches are able to provide fault isolation and hot-swap ability at the port level.

Fiber Channel connections can provide the performance power to handle a multitude of bandwidth-intensive storage management functions such as backup, remote storage (sometime called *vaulting*), and hierarchical storage that frees the LAN resources from these traffic-intensive and bottlenecking services.

## What Are the Limitations of SCSI?

As you may know, there are several incarnations of the small computer system interface (SCSI) such as Wide SCSI, Fast SCSI, and SCSI-3. SCSI is traditionally and currently the interface of choice for computer storage that needs high-speed connectivity for Unix and Windows NT systems. Some characteristics of SCSI, however, limit its ability to enhance LAN storage performance.

For example, SCSI does not support multiple host-to-storage device connections very well. Its main strength is point to point, or directly attached computer to storage device interfaces. There is also a concern with the throughput of SCSI devices. The traditional SCSI data throughput rate of 40 Mbps can quickly bottleneck most of today's database applications and multimedia information transfers. The SCSI limitation of 15 devices maximum per channel is also discouraging to those companies that want to create multiple server to multiple storage device network designs. This limitation on device hookups is further reduced as bus bandwidth increases.

There is also the 25-meter point-to-point connection maximum distance between devices, and if you want to use the faster throughput capabilities of low voltage differential (LVD) SCSI, this drops to 12 meters. This requires storage units to be located close to the server, often within the same enclosure. This is generally unacceptable as it may be a waste of rack space. In addition, configuring a server and its storage within a single chassis usually results in an expensive relationship between the scaling of server capacity and that of upgrading storage capacity. Another drawback inherent with the short cable length is that it prevents storage devices in a centralized environment to perform remote mirroring. SCSI can also take away from LAN resources, as backup traffic from server to server must travel over the LAN. This can place additional strain on the network, and takes away from the bandwidth of users.

## All Fiber versus Mixed Solutions

Fiber Channel was designed specifically to address SCSI limitations. A basic server-to-storage device connection using the Fiber Channel bus (not including a SAN) can greatly improve your overall network and storage access performance. Fiber Channel greatly augments the information transfer throughput for high-



end applications such as digital imaging, video streaming, databases, and computer aided design (CAD). It can also extend connectivity distances of remote backup, data archiving, and site mirroring for disaster recovery purposes.

At current SAN installation sites, Fiber Channel can show immediate and measurable operational benefits over those same connections with SCSI devices. For example, by connecting RAID to a server over a Fiber Channel bus, the higher bandwidth will result in quicker data transfers over greater distances than are possible with standard SCSI. The RAID-Fiber Channel combination helps to improve storage reliability through fault-tolerant operations and redundant data paths.

Fiber Channel switches and hubs are able to provide simplified storage device scalability, hot swapping of storage devices, and isolation of functions. This translates into readily scalable bandwidth and improved system availability.

Fiber Channel provides a combination of bandwidth, performance, high availability, and flexibility in the configuration for servers and disk arrays. By incorporating Fiber Channel connectivity throughout a RAID subsystem, you will be able to surpass SCSI or Fiber Channel-to-SCSI hybrid solutions.

Fiber Channel takes the best features of SCSI and brings it to a high-speed interface. The Fiber Channel physical layer is much faster, but it is still lagging behind many of today's infrastructure transport technologies. Speeds range up to only 100 Mbps, but there are plans to increase these speeds to 200 Mbps and 400 Mbps in the future. Another major difference between SCSI and Fiber Channel is *arbitration* (in other words, who owns the bus). Fiber Channel is faster and equal, not prioritized like SCSI.

Based on the media that is implemented, Fiber Channel can connect devices that are up to ten kilometers (over six miles) apart. SCSI, on the other hand, has cable distance limitations of 25 meters (80 feet). Most SCSI devices support one transmission medium, or a mutually exclusive transmission media. With Fiber Channel, though, you can use copper as well as fiber-optic cable. Remember, though, that the transmission medium used will affect the distances between nodes.

A single shared bus cable further limits SCSI, whereas Fiber Channel can connect nodes by a number of methods such as loops, hubs, and switches. Lower costs will result from having a channel that supports a large number of devices with a single connection.

Fiber Channel hubs can make a loop that looks like a series of point-to-point connections. This is done so that if a cable breaks while connecting a device to the hub, the hub will then remove that portion of the loop from service, but it will still keep the rest of the loop functioning. This makes the adding and removing of devices simple and nondisruptive to data traffic as it flows on the

loop. SCSI implementations do not support hubs, and the addition and removal of hardware must be done when all data traffic has been stopped.

Fiber Channel switch implementations permit multiple devices to be connected through multiple loops, thereby aggregating bandwidth. Therefore, multiple 100-Mbps loop configurations can be managed by one centralized point. Bandwidth can then be allocated according to device demands. Reconfiguration in this architecture is fairly simple. Since SCSI does not support switches, Fiber Channel switches cannot be grouped for performance with SCSI.

Fabrics are composed of multiple switches and enable Fiber Channel networks to scale to very large sizes. This configuration offers extremely high bandwidth. Fabrics can also span very large geographic areas. When other protocols (e.g., TCP/IP) are ported for use in Fiber Channel implementations, they will rely on fabric implementations for their means of transport.

SCSI connections are generally hard to manage in large deployments, and they are especially difficult to diagnose when there are intermittent problems. When configuring SCSI for high availability (such as multiple hosts and multiple arrays), connecting devices is a sophisticated and resource-consuming task. Fiber Channel eliminates split or tap cables and terminators.

There are also some form factor and support issues that are inherent within the SCSI architecture. You see, SCSI cables have 68 wires, whereas Fiber Channel has only four. While it is rare that the wire may fail, connectors can fail, and do. Each Fiber Channel connector has fewer connections on it than does SCSI, and because the Fiber Channel cable is thin, lightweight, and flexible, there is less stress on the connector. In comparison, SCSI cable is thick and not very flexible; therefore, when implementing SCSI in tight configurations, the bending stress may be transferred to the connector.

Fiber Channel architecture supports several fault-tolerant features that are not available or even practical within bridged SCSI solutions. By using Fiber Channel's ability to dual-loop capability with dual-ported Fiber Channel drives, you can easily deploy complete data redundancy without a single point of failure.

Fiber Channel drives and Ultra-2 SCSI drives are comparably priced. However, the installation of SCSI drives is more resource consuming and prone to errors. Fiber Channel SANs can be reconfigured and resources can be reallocated while the chassis is online; SCSI installations require that the system be taken offline when more complex adjustments or device reassignments are needed.

There are some older hybrid implementations around, as not all users are convinced that Fiber Channel is the best end-to-end solution. Some ASPs have opted to integrate SCSI and Fiber Channel on the fabric, while others have

chosen combination Fiber for the front end and SCSI on the back end.

Companies' reasons vary from economics to performance needs.

Some ISPs have spent a lot of money on SCSI devices, so they would like to get the greatest return on that investment. As it stands, most existing SCSI disks will remain directly attached to servers that are outside of the SAN, but there are products on the market that can provide seamless integration between SCSI and Fiber Channel. By being able to use both of these technologies on the SAN fabric, you will be able leverage that investment, especially for devices such as tape libraries.

Other service providers have settled on installing Fiber Channel between their servers and storage device controllers, while maintaining SCSI on the back end; for example, SCSI between the RAID controller and an individual disk drive. You will realize less benefit with Fiber Channel behind the RAID controller, as the RAID architecture is often more important than the device channel.

By installing Fiber Channel on the front end, you will be able to communicate with the storage device over a considerably longer distance because the host and controller are conversing in Fiber Channel with SCSI hidden behind the controller. This configuration enables the infrastructure to gain many of the benefits that come with Fiber Channel while still being able to leverage their existing investments in their current RAID architectures. As we move forward, you will see full Fiber Channel implementations becoming the dominant architecture implementation.

With Fiber Channel implemented in end-to-end connectivity, clients will receive all of the benefits such as scalability, throughput, configuration options, and robustness that this deployment offers. You are also investing in infrastructure technology for the future. However, you can also get many of these benefits with a combination of Fiber Channel and SCSI topologies.

## SAN Management

To truly achieve the full benefits and functionality of SANs, such as performance, availability, cost, scalability, and interoperability, you must be able to effectively manage the infrastructure (switches, routers, other networking devices) of the SANs. To simplify SAN management, SAN vendors have adapted the Simple Network Management Protocol (SNMP), Web Based Enterprise Management (WBEM), and Enterprise Storage Resources Management (ESRM) type standards that can monitor, manage, and alert the components of the SAN.

Many customers will also want or need to manage their partitions of the SAN from a centralized console. The biggest challenge that you and your vendors

face is to ensure that all components are able to work with the various management software packages that are available. Management of the SAN can be divided into various areas that are defined in ESRM. These areas should be implemented across all of your resources that are connected to the SAN to provide a common user interface across all of your resources.

## Capacity Management

Capacity management is the ability to address the sizing of the SAN; for example, how many switches are needed. It also addresses the need to know how much free space, available chassis slots, unassigned volumes, free space in the assigned volumes, number of completed backups, number of tapes, percent utilization, and percentage of the disk that is free.

## Configuration Management

Configuration management handles the need for determining the current logical and physical configuration information, ports utilization information, and device driver information to support the SAN configurations based on the business requirements of high availability and connectivity. It also deals with the need to integrate the configuration of storage resources with configuration of the server's view of them. For example, when a client configures an enterprise storage server, it affects what must be configured at the server.

## Performance Management

Performance management handles the need to improve performance of the SAN, and does problem isolation at all levels (device hardware and software interfaces, application, and even the file level). This approach generally requires a common platform, and independent access standards are implemented across all SAN solutions.

## Availability Management

Availability management takes care of the need to prevent failure of equipment and correct problems when they occur by providing warnings of the key events long before they become critical. For instance, in the event of a path failure, the availability management function may be able to determine whether a link or other component has failed, and assign an alternate path, while notifying an engineer to repair the failing component, thus keeping the systems up throughout the entire process.

# Scalability and How It Affects Your Business

Scalability is of great importance, especially for storage devices. The truth is that even if you design a feature-rich solution and purchase an immense amount of storage capacity in the beginning, it is likely that you or your customers will fill this capacity much faster than you could anticipate.

In some cases, you may only need to add additional hard drives to your storage devices; however, if your total storage solution was not carefully planned for or properly implemented, you may outgrow your solution altogether. As you can imagine, upgrading your systems can become a very cumbersome and time-consuming task. I doubt that you want the expense of removing your original solution that did not scale, in favor of a scalable one.

## Storage in Your Infrastructure

The cold hard fact is that storage fills up. You are probably already familiar with the problem of running out of hard drive space, and have in fact had to make a trip to the local computer store to purchase an additional hard drive. This may have been for your own personal computer in order to install the latest and greatest application, or to copy some important data to your system. Even if you have never touched a computer in your life, you are probably familiar with a lack of storage space in some way.

Running out of storage space happens to nearly everyone in some way or another. Most of us can remember how empty our first house or apartment seemed. Some of us may have even had the thought of “How will I ever fill it?” only to later find ourselves with our house or apartment filled to capacity with little to no room for some of the additional items we would like to purchase.

Some of us are “packrats” and we refuse to throw away anything, filling our living quarters with so much stuff that at times it may even become difficult to move around in our own homes, or even live there! It is amazing how our original notion of our own home’s storage capacity could become shattered so quickly by placing a few additional items in it. It is not so unlikely that we would also tend to underestimate our data storage concerns so easily.

Some years ago, it was unheard of to own and purchase a single hard drive with 100MB capacity, and now we have models capable of storing tens of millions of bytes of information. The only reason we have even attempted to redesign and re-architect the original 100MB hard drive is because we ran out of storage space

with the ever-increasing amount of data stored on our systems. Today's computers are power-hungry devices, and as we have seen the emergence of the Internet, so have we seen the need for more power, speed, and storage abilities.

Luckily, it is a somewhat simple task to install additional storage these days. Most servers have hot-swappable modules that allow you to remove or install each hard drive in a matter of seconds without even powering off the system. The real drawback to this scenario is that there may be a software element that needs changing, and poorly designed hardware and software may require that you power cycle the server in order to actually use the hard drive.

Another issue that was mentioned earlier is that servers, and some other storage devices such as disk arrays, have physical limitations that will dictate how many hard drives can actually be installed in their chassis. If you have designed your entire storage solution around the confined physical space of a few devices, you will most likely need to purchase additional servers or arrays to give you the space to install new hard drives. This can be an expensive solution to what should have been a simple problem.

As you have already read, NAS devices were designed for storage over the network. This means that they should all be very flexible with their total storage capacity, right? The reality is that some NASs are flexible and allow for an almost unlimited amount of storage to be added with few physical limitations, while others are no more than simple storage arrays with a proprietary operating system and network adapter built into them to provide network connectivity.

To decide if physical limitations are an issue, you will need to look at how hard drives are added to a particular storage device. It will not be good enough to read the brochure, because you need to have a visual understanding of the device in question and estimate what its physical limitations may be. This will be important, since one of the largest scalability issues for NAS devices is their physical properties.

In addition to these limitations, there are usually other factors that will limit the number or size of the hard drives that can be installed in a single storage device. When looking for these products, always make sure to check what the maximum storage capacity will be, and look for any caveats to the claim.

For instance, a manufacturer might claim that you can add up to 100 hard drives for a maximum capacity of 1.8 terabytes of data to their product. However, they may not explain that they offer several different sizes of hard drive. In this case, you might order and purchase 450GB worth of capacity to start, and they might deliver 50 hard drives that each have 9GB of capacity.

This might seem fine, since it appears to leave a possible 900GB for future expansion; however, you might not be able to use different-sized hard drives in

the storage device, meaning the true maximum you would be able to upgrade the unit to will be 900GB, not 1.8TB. Therefore, in order to reach the 1.8TB mark, you might need to first remove the 50 9GB drives and replace them with 18GB hard drives. This could prove to be expensive, especially if you do not or cannot recycle the 9GB hard drives in a different storage device.

If you did not clarify the size of the drives to use, and instead asked for a particular amount of storage, it may not be anyone's fault but your own. If anything, remember to be careful with what, or how, you ask for something. "Let the Buyer Beware."

As is usually the case, the most scalable and feature-rich solution is usually the most expensive. SANs are without a doubt the most expensive storage solution you can design, and they do in fact provide for very scalable storage architectures. The same problems relate to individual SAN storage devices as relate to all other storage devices, such as physical limitations and maximum number of hard drives. Remember that a SAN is not a single device, but rather a collection of devices, meaning that the scalability concerns of a single device are diminished since the SAN is designed to allow for numerous individual devices to be connected to it.

A SAN is designed to span great distances, which allow it even more flexibility, since there is not a requirement for the SAN devices to be in close proximity to the hosts that access them. Even if you were to run out of space at a particular location, a device might still provide the same functionality and performance if it were installed elsewhere, as long as it could still be attached to the same SAN.

## Wire Speed and How It Can Help You

Wire speed plays an important role in delivering data to host devices. Whether your environment consists of directly attached storage, NAS, SAN, or a combination there of, you will still have bandwidth concerns that will limit the amount of actual data that can be sent across the wire at any given moment.

Imagine for a moment that you are at home and need to wash your dishes. When you turn on the water faucet, water begins flowing from the water utility company through very large pipes that span through the county and are distributed until the water is finally delivered to your home and faucet. The amount of water pressure available to you at any given time is conditional on two factors.

The first is the size of the piping; only a certain amount of water will be able to fill any given pipe, and thus only a certain amount of water will be able to flow out of your faucet. The second factor is the number of other faucets drawing water from the pipe at the same time. Given a finite amount of water

that must be distributed among multiple homes at the same time, your water pressure will be dependent on the number of simultaneous users accessing the water supply.

If the plumbing was designed and implemented correctly, you should have a good and consistent amount of water pressure each time you turn your faucet on. On the other hand, if it is poorly designed, you might have consistently low water pressure or high water pressure at times of infrequent use, and low water pressure at times of high use. The same scenario can be easily translated to your storage solutions.

You could consider your data to be the water, and the piping as the data delivery fabric that connects your storage devices to hosts. Much like the preceding example, only a certain amount of data will be able to flow through your network, depending on the type of interfaces and devices used in the network. In addition, the total amount of data available to each host is dependent on how many users are accessing the data at the same time. Again referring back to our example, if your solution is properly designed, it will consistently service many hosts efficiently. However, if it is poorly designed, a particular host's bandwidth could fluctuate significantly.

Directly attached storage offers many different methods of delivering data to a host device, such as using SCSI or EIDE to connect a device to the host. Since the purpose of this architecture is solely to connect storage to a single device, there should always be a consistent amount of storage bandwidth available to the host device.

By definition, there will never be an instance of other hosts accessing the storage media other than the device attached to it. Because of this, the amount of data available to the host at any given time is directly related to the capabilities of the hardware and bus technology in use.

With NAS and SANs, data is being delivered across a SAN or LAN, so the available bandwidth is dependent on several factors. The first, and probably most important factor is the configuration and type of network interface installed in a particular SAN or NAS device. The installed network interface will dictate the maximum amount of bandwidth available for the device.

For instance, if the network adapter is a 100 Mbps Fast Ethernet or Fiber Channel adapter, the maximum amount of bandwidth available would be 100 Mbps. Conversely, if the adapter were Gigabit Ethernet, you would have 1000 Mbps of available bandwidth. In addition, if you were using a half-duplex configuration, the available bandwidth would be evenly distributed between both the transmission and reception of data. However, a full-duplex configuration would



yield a pipe that is capable of delivering the maximum bandwidth in each direction of data travel.

The type of network hardware deployed will also play an important role in your SAN or LAN's available bandwidth. If you are using a shared medium such as hubs or concentrators to connect your devices, the total amount of bandwidth available to any device is relative to the amount of data flowing to and from all of the devices within the shared network segment. This means that if your devices are using 100 Mbps connections, there are essentially only 100 Mbps of total bandwidth available to the devices as a whole, not 100 Mbps to this one, and 100 Mbps to that one.

What's worse is that this will tend to create multiple collisions in the network segment, causing data to be dropped and forcing many devices to retransmit their data. In an overly chatty network, this can become a nightmare and make much of the available bandwidth "unusable." As discussed previously, a shared medium is not recommended because it does not allow anywhere near the true wire speed of each individual device. If instead, you were to employ switches to connect your devices, the data path would be much more streamlined and efficient, and allow your usable bandwidth to come close to the maximum speed of your installed network interfaces.

If you have bridges or routers segmenting hosts and the storage devices they access, the speed of data flow could be seriously hindered depending on their configuration, amount of utilization, and inherent speed. For instance, if you have a router with 100 Mbps interfaces segmenting your host and storage devices, and all of your data needs to flow through this router, the conversations will need to share the 100 Mbps pipe provided by the router.

This means that even if you have ten devices capable of delivering 100 Mbps each, they will all be distributed among a single 100 Mbps pipe when traveling through the router or bridge. Moreover, depending on the configuration and inherent speed of the router, it is usually not possible to achieve the true wire speed of the network interfaces. This is especially true with speeds in excess of 100 Mbps.

Even if you were to use a Gigabit Ethernet interface, or an EtherChannel configuration to bond multiple links together, chances are that some overhead will be lost in the processing of the data by the router or bridge. This problem can become compounded if a device has to perform complex tasks to the data traversing its interfaces, such as a router that needs to match all packets flowing through its interfaces against a given set of rules or access lists for the purpose of security.

Although routers and bridges can significantly reduce your performance, the features and capabilities they possess will probably make it impossible or you to rule out using these devices in your SAN or LAN. Whether you decide to use hubs, switches, bridges, or routers in your network, you should always try to design an efficient flow of data and look to solve any problematic bottlenecks you may foresee during the design and evaluation process.

## One versus Many

After reading through some of the capacity concerns, you may have decided to run out with a check and purchase one massive do-it-all storage device with bonded Gigabit Ethernet adapters that provide more than enough bandwidth and storage capacity to last you for 20 years. This might be a good idea, depending on your situation, but it is most likely a very bad idea.

Although we stated earlier that it is important to look into devices that provide scalability features in the devices themselves, you shouldn't assume that you would only ever want or need one device. Remember the saying, "It is never a good idea to put all your eggs in one basket?" If you do happen to buy one device, or put your eggs in one basket, Murphy's Law says that the device will break, leaving you with a worthless heap of technology and no immediate replacement. If not just for fault tolerance reasons, you should consider solutions that incorporate several devices instead of just one. Besides this, there are still several other reasons to consider using many devices instead of one.

Speed, as we discussed earlier, is essential when delivering data to your host devices. If you do not have a fast SAN or NAS solution, it might be pointless to implement it. In addition to the speed concerns we introduced earlier, you should look into some of the mirroring options available with some storage devices. In many cases, it is possible to install several storage devices that can all add up to provide a single solution.

For instance, some manufacturers offer what are known as active/active configurations. What this means is that you can purchase several storage devices that can all be used as a single cluster of devices. This allows hosts to access the same set of data, but on different storage devices, thus reducing the traffic to a particular storage device, and distributing the host connections evenly among each device in the cluster.

Although this could be expensive, it could also provide for an extremely fast and scalable solution. If you are looking to supply data storage services to customers or plan to receive thousands of simultaneous connections to your storage devices, the "many" option is the way to go. In fact, in some instances it may be

the only way to go, since every storage device will have a limitation as to how many simultaneous connections it can handle.

If you do think you will approach this limitation, we do not recommend that you continuously push your storage devices to their maximum level of tolerance, since doing so is bound to cause a problem at some point in time. Instead, simply add another device. Whereas it might cost a little more in the beginning, it will save you problems in the long run.

Adding additional devices does not always mean that they have to be costly active/active pairs, and it doesn't even mean that you need to purchase the exact same make and model of storage device, although this might help you support the device and add additional features. It might simply mean to add storage devices to your solution, and design a logical way of partitioning the data across these multiple devices that allows for a balanced level of access to each storage device.

This might mean that you monitor the data accesses and distribute the data based on a complex model of access frequency and usage patterns, or that you employ a SAN virtualization scheme that includes storage domain servers to allow access between unlike devices. Regardless of how you add devices, or partition and virtualize the data, it is easy to see the benefit that multiple devices offer.

In the end, it could prove that the long-term operating and upgrade costs are smaller when using many storage devices as opposed to using only one or a few. There is also a definite possibility that multiple storage devices will help keep your storage solutions scalable and continue to satisfy your customers with the speed and reliability gained.

## Fault Tolerance Features and Issues

Devices fail for numerous reasons, and it is next to impossible to predict when it will happen. Sometimes a warning will be given prior to a failure, while in most cases the device will function normally and fail suddenly at an unexpected time. If there is one true bet you can make, it is that at least one of your most critical devices will eventually fail. It is this type of realistic thinking combined with future planning that propels systems to be built with high levels of fault tolerance.

Many people think a fault-tolerant system means that you need to buy two of every device instead of just one, and continue to double every device throughout your network until your network itself is effectively doubled. In this type of design, the second device is hardly ever used and usually sits and waits for its corresponding device to fail. While this might seem to be the most fault-tolerant system available, it is obviously the most costly and tends to be a gross waste of

money and resources. Fortunately, there are numerous other ways to deal with fault tolerance, and when it comes to data storage, the possibilities can be endless.

## Shared Resources

One of the largest advantages a SAN has to offer is the true ability to share resources between other server and host systems. In the past, it was possible to share the media between devices, but it was not possible to share the actual data that was stored on the disks. Instead, the storage devices were split into separate partitions that were divided among the hosts. This was never really an issue, since most operating systems were not capable of sharing resources among other hosts. Things have changed, however.

With the advent of SAN, we have redesigned how resources are treated among systems. With SAN, it is possible to share the logical partitions and actual data among several different servers. It may not seem like a huge advancement, but it enables us to do a much better job of providing for fault tolerance.

Imagine, if you will, a group of three servers that are attached to a storage pool through a SAN. Each of these servers might serve three separate databases. Using software on the servers, it is possible to have each of these systems provide fault tolerance for the other two systems. If one of these three servers were to fail, the other two would have access to the live database stored in the storage pool, and actually take over processing functions for their failed partner.

There is no need to restore the database from a backup, because the other servers already have access to the most current version of the database. This provides an excellent way to cut down on time and costs, especially when you consider that there is no need to buy a server that is set aside and dedicated as the failover server.

## Data Backup

One of the most important steps to providing a fault-tolerant system is the periodic backup of mission-critical data. Data backup is usually accomplished by copying the most current data onto backup tapes, which are inserted into specialized devices. The tapes are designed to be resilient in nature, and can be overwritten numerous times without causing damage or errors to the media.

These specialized tape devices come in a variety of shapes, sizes, capabilities, and price ranges. Some are inexpensive and require you to physically insert each individual tape as it is used, while others are more expensive, sophisticated systems that house a robot capable of selecting between hundreds of tapes that have

been preinserted. The benefits of each are fairly obvious; one requires user intervention, while the other offers “set it and forget about it” functionality.

If your data changes or is frequently altered, it is not enough to back up the data lackadaisically. Depending on your situation, it may be necessary to design a strict strategy, and force your systems to automatically back themselves up daily. In some cases, it may be necessary to do this several times in a single day. It takes quite a long time to back up this data, though, so if you are looking to back up your systems more than once a day, you will probably need to look at a more advanced solution such as remote mirroring in conjunction with a less frequent backup strategy. Refer to Chapter 3, “Server Level Considerations,” for more information about backups.

## Remote Mirroring

Remote mirroring is an excellent form of disaster recovery offered by SAN technology. Today, it allows for a complete copy of your data to be contained at a remote location that might be located up to 40 kilometers away. Some of the upcoming product offerings promise even more functionality, and even claim that the distance limitation will completely disappear in the very near future.

This can be a tremendous advantage when one location has been seriously damaged or becomes inoperable due to a serious catastrophe, such as a fire or “act of God.” Although we would like to believe that these situations may never arise, and will never happen to any of us, it is not always the wisest of decisions to take the risk. There are two distinct types of remote mirroring possibilities:

- Synchronous
- Asynchronous

### Synchronous

Synchronous mirroring allows all data to be written to both the primary and backup site simultaneously. This technology is widely available, and allows for near instantaneous disaster recovery when a site has become inoperable. However, the solution requires a good amount of usable bandwidth between the two sites, and can therefore become quite costly.

Due to latency issues and the difference in time it might take to successfully write the data to both sites simultaneously, there is an approximate 10-kilometer distance limitation with which to be concerned. Although these can be significant

reasons to stay away from this technology, it is without a doubt a robust fault-tolerance solution that should be considered if disaster recovery needs to be an extremely quick operation and always up to date.

## Asynchronous

If you want to extend the distance between your mirrored storage beyond 10 kilometers, or use less bandwidth between locations, asynchronous mirroring is the way to go. It works in the same manner as synchronous mirroring, except that data bound for the mirror site can be buffered and transmitted at a slower pace.

Disadvantages to this method lie in the fact that the mirror site is never in full synchronization. If the primary site were to fail, there would be some amount of data loss. This would probably not amount to a huge loss of data, though, and may be an acceptable risk depending on your service level agreements. Besides, this would definitely prove to be a superior solution than not having a mirror at all! The majority of your data would still be safe, and there would still be a minimal amount of actual system downtime.

## Redundant Array of Inexpensive Disks

Redundant Array of Inexpensive Disks (RAID) provides methodology for storing the same data in different places on multiple hard disks. By placing the data on multiple disks, it is possible to take advantage of multiple I/O operations to improve speed and performance. In addition, since the data is stored redundantly and across multiple disks, RAID offers an excellent way of providing fault tolerance. For instance, if a hard drive were to fail, the data would still be stored on other hard drives in the RAID array, thereby providing for minimal to no loss of data.

A RAID solution functions by combining all of the hard drive within the configured array to make them appear as one logical hard disk. To accomplish this, RAID uses a technique called *striping* to break up the logical disk into units that could range from a single sector to many megabytes of space depending on the RAID configuration. The stripes created from this are interleaved and addressed in order. RAID allows you to use small stripes to ensure that all the data is distributed across all of the disks providing better fault tolerance, or larger stripes that can provide better performance in a multiuser environment.

Some forms of RAID will use parity or error checking and correcting (ECC) as a way of verifying and restoring lost data. With this method, when a unit of data is stored on a storage device, a code that explains the bit sequence of the unit is calculated and stored as parity information on the hard drives. When this

unit of data is accessed, the code is again calculated based on the stored unit of data, and this new code is compared against the code that was stored when the data was initially written.

If the two codes match, chances are the data has not changed in any way, and the requested data is transmitted. If the codes do not match, however, the erroneous or missing bits of information are determined and the unit of data is corrected before it is transmitted. If the error recurs even after the power has been cycled, it will be detected as a hardware failure, and depending on the RAID device and configuration, the error will get logged or a notification will be sent to the system administrator.

Almost every manufacturer of mass-storage devices supports some form of RAID. It can be found in directly attached storage configurations, as well as in most SAN and NAS devices. Even if the capability does not come with a particular device, it is usually available as an option, or can be added with additional components. It is in such widespread use that you can even buy a RAID adapter for your personal computer.

RAID has been around for many years, and as a result, numerous versions have been conceived and are available for use. When looking at a specific vendor's device, it is important to check what versions of RAID are supported, as it is rare to find a device that supports all versions. The following is a list of the most common versions of RAID:

- RAID-0
- RAID-1
- RAID-2
- RAID-3
- RAID-4
- RAID-5
- RAID-6
- RAID-7
- RAID-10
- RAID-53

## RAID-0

RAID-0 is the most simplistic form. It provides striping, but does not provide any redundancy of data. The full amount of storage capacity is available for user files, unlike other versions of RAID that will use up a portion of the disk in order to provide redundancy. This solution should only be used when optimum performance is required without a fault-tolerant requirement of the RAID.

This might be a good solution if you are constantly backing up or mirroring your data to another device, or another device is providing your fault tolerance. In this way, you could still have a fault-tolerant solution while achieving optimum performance.

## RAID-1

RAID-1 is often called *disk mirroring*. It provides a true duplicate of all data on at least two hard drives, meaning that data can be read from both hard drives simultaneously, thereby improving read performance. However, since both disks also need to be written to simultaneously, there is not a write performance gained. If you are not frequently backing up your data or have no exterior fault tolerance, it might be a good idea to use RAID-1. The drawback is that half of your available space will be consumed by redundancy, meaning that if you require 100GB of data storage, you will need to install 200GB of actual hard drive space.

## RAID-2

RAID-2 provides striping across all of the disks, and stores ECC information throughout the hard drives. It is rare to find RAID-2 in use, since it is equivalent to RAID 3 and does not offer any advantages over it.

## RAID-3

RAID-3 also uses striping across all of the disks and provides for data redundancy, but dedicates one hard drive to store parity information, which is used to for errors. Data recovery can be performed by calculating the exclusive OR of the information that was stored on the nonparity hard drives. This method of RAID uses all of the installed hard drives when data is accessed, which does not translate into a performance increase. Since only one hard drive is needed to store parity information, most of the actual drive capacity will be available for data. Although RAID-3 provides for some fault tolerance, there are not many reasons to use it since there are more resilient and performance-increased versions available.



## RAID-4

RAID-4 uses large stripes and a separate drive for parity. It offers slightly improved performance over RAID-3, since all the drives can be read from at the same time. However, when data is written to the disk, all drives are used to update the parity drive. As with RAID-3, this gives no performance boost with disk writes. RAID-4 is usually not used, since RAID-5 is generally considered a better solution.

## RAID-5

RAID-5 uses a rotating parity array to overcome the disk write limitation found in RAID-4. This means that there is improved performance for both disk reads and writes. RAID-5 does not store any redundant data, and instead uses the parity information to reconstruct any lost data. RAID-5 is usually deployed in multiuser environments where a medium level of fault tolerance is acceptable. It is the most popular version in use today because it provides improved speed and fault tolerance with minimal loss of usable space.

## RAID-6

RAID-6 is a newer version of RAID that is very similar to RAID-5, except that it stores a second parity scheme that is distributed among all the hard drives. Since there are two sets of parity, if the parity information is lost, it can be restored using the duplicate parity information. Also, with two sets of parity, there is a smaller likelihood of storing incorrect parity information. This means that RAID-6 provides very high fault tolerance that can easily survive most drive failures. Although this is a great solution, for some reason it is hard to find manufacturers that support this version of RAID.

## RAID-10

RAID-10 offers the ultimate high-performance RAID mirroring solution. It uses an array of stripes in which each stripe is a RAID-1 array of drives. This means that if you have 10 drives, you will have five groupings of mirrored drives that can all be written to or accessed at the same time. This obviously provides the same level of fault tolerance, as with RAID-1; however, there is a significant performance increase.

The drawback is that it can be an expensive technology to use. However, if you require high performance and want to rely solely on your RAID solution, this would be the way to go. However, I would argue that an exterior fault-

tolerant solution in addition to a different level of RAID, such as RAID-5, would yield a more flexible solution for an equivalent price.

## RAID-53

RAID-53 offers an array of stripes in which each stripe is a RAID-3 array. This will offer the same level of redundancy and data protection as is offered by RAID-3, but improves the performance by allowing multiple stripes to be read and written to simultaneously. Although it offers superior performance to RAID-3, it is also much more costly.

# SAN Solutions Offered by Various Vendors

Many solutions are available to you when making your decision. Since there are so many, and they are readily subject to change, we have included a sample of some things to look for when implementing your SAN. We used IBM as a template, but by no means do we think that this is the only solution available.

## IBM's SAN Solution

IBM offers a wide and complete range of services and products that include software, infrastructure design and support, and other technologies that are required for you to implement a SAN solution. IBM's SAN solution allows you to:

- Scale the reach of network and allow for manageability with a centralized data center
- Provide access anytime, anywhere, to data irrespective of the platform, source, format, or application type
- Enhance the security and integrity of data on your network infrastructure

## The IBM SAN Strategy

IBM's SAN strategy involves the migration to a SAN infrastructure over time. It tries to deliver its SAN strategy in phases, to leverage new technologies once they are proven, and to help seamlessly integrate SAN technology into a company's IT infrastructure; all this while protecting your investments in application resources, servers, and storage.

IBM SAN technology evolves in three stages:

- **SAN attached storage** This leverages the any-to-any connectivity of SAN technology.
- **SAN optimized storage** This makes use of SAN characteristics and delivers strong SAN solutions.
- **SAN optimized systems** This leverages proven technologies and delivers SAN systemwide solutions.

IBM's SAN solution uses Fiber Channel architecture for connectivity and device-level management. It also provides businesses the basic building blocks that will enable IT resource management and information sharing anytime, anywhere across your storage area networks.

Value can be added to the Fiber Channel infrastructure by adding new storage solutions and comprehensive fabric management, thus helping organizations to manage, track, and more easily share the sophisticated and increasing volume of data created by business applications and the Internet.

## Summary

Your ASP might provide storage solutions for your customers, or you might solely rely on data storage for your own internal purposes. Regardless, your ultimate storage goals and uses will dictate the model of storage you require. If you have minimal centralization and storage requirements, you may want to go with the age-old directly attached storage solution.

This does offer a very simple and successful solution; otherwise, it would not be in such widespread use. If you are instead looking to deliver large amounts of data to your clientele, and need a system capable of performing this task, you will probably decide to use NAS devices that are distributed throughout your network. You might even have separate data and storage concerns that can justify designing an expensive SAN solution to connect several sites together and provide for the most robust set of features.

This, too, is a very viable solution depending on your model. The reality is that all the storage options that we have explained provide for excellent solutions depending on their use and purpose. Likewise, they can also provide for inefficient or cost-deficient solutions when not understood or planned for correctly.

In this chapter, we tried to explain some of the criteria you should consider when designing your storage solution. We covered the characteristics of directly attached storage, NAS, and SAN, in order to give you a better understanding of each and make an informed decision as to which solution best fits your company's goals and budget. We went into some detail as to the features and functionality that each solution has to offer, and explained the advantages and disadvantages of each.

We spoke about scalability issues, in the hope that you will use this information to design a solution that will exist for as long as your company thrives. Finally, we spoke on the issue of fault tolerance, and some of the options that particular storage solutions have to offer. All of these topics were presented to help you build a solution that fits your particular criteria.

In the end, only you know your goals and requirements, and can weigh these against the storage solutions we presented. Be careful in your selection, and always look for a solution that leverages good technology with adequate features that is the "right fit" for your organization rather than the cheapest solution or the "latest craze."

## Solutions Fast Track

### Upfront Concerns and Selection Criteria

- ☑ Currently, there are many differing manufacturers of storage-based equipment, and several methods of delivering storage solutions to your servers and clients.
- ☑ With mass-storage products, some of the major manufacturers may only offer proprietary equipment, while others may standardize their equipment, using a technology such as fiber channel to ensure that their product will work with a similar offering from another manufacturer.
- ☑ Security should always be a concern, but it is especially important given the high visibility of ISPs and ASPs.
- ☑ Outboard security is any type of security feature that is located on the host. It might be an external authentication scheme that is provided by a firewall.
- ☑ You may already own storage devices that use interfaces other than fiber channel, such as small system computer interface (SCSI) or enhanced integrated drive electronics (EIDE) for host connections. It can sometimes prove difficult to port older hardware to some newer storage solutions.

### Directly Attached Storage in Your Infrastructure

- ☑ Server-to-storage access, or directly attached storage, has been in use in much of the history of computing, and still exists in over 90 percent of implementations today.
- ☑ In directly attached implementations, storage devices are directly connected to a server using either interfaces and/or bus architecture such as EIDE or SCSI.

### Network Attached Storage Solutions

- ☑ A NAS is a device that provides server-to-server storage. A NAS is basically a massive array of disk storage connected to a server that has been attached to a local area network (LAN).

- ☑ QoS has the ability to delegate priority to the packets traversing your network, forcing data with a lower priority to be queued in times of heavy use, and allowing for data with a higher priority to still be transmitted.
- ☑ When designing NAS in your network, probably the most effective solution for latency and saturation issues is the location of your NAS servers in relation to the hosts and systems that access their data.

## Storage Area Networks

- ☑ A storage area network (SAN) is a networked storage infrastructure that interconnects storage devices with associated servers. It is currently the most cutting-edge storage technology available, and provides direct and indirect connections to multiple servers and multiple storage devices simultaneously.
- ☑ A SAN can be thought of as a simple network that builds off the familiar LAN design.
- ☑ Distributed computing, client/server applications, and open systems give today's enterprises the power to fully integrate hardware and software from different vendors to create systems tailored to their specific needs.
- ☑ SANs remove data traffic—backup processes, for example—from the production network, giving IT managers a strategic way to improve system performance and application availability.
- ☑ Multihost arrays are the most simplistic and most common form of SAN virtualization implementation.

## Scalability and How It Affects Your Business

- ☑ A SAN is designed to span great distances, which allow it even more flexibility, since there is not a requirement for the SAN devices to be in close proximity to the hosts that access them.
- ☑ Wire speed plays an important role in delivering data to host devices. Whether your environment consists of directly attached storage, NAS, SAN, or a combination there of, you will still have bandwidth concerns that will limit the amount of actual data that can be sent across the wire at any given moment.

## Fault Tolerance Features and Issues

- ☑ One of the largest advantages a SAN has to offer is the true ability to share resources between other server and host systems.
- ☑ Remote mirroring is an excellent form of disaster recovery offered by SAN technology. Today, it allows for a complete copy of your data to be contained at a remote location that might be located up to 40 kilometers away.
- ☑ Redundant Array of Inexpensive Disks (RAID) provides methodology for storing the same data in different places on multiple hard disks.

## SAN Solutions Offered by Various Vendors

- ☑ IBM's SAN strategy involves the migration to a SAN infrastructure over time. It tries to deliver its SAN strategy in phases, to leverage new technologies once they are proven, and to help seamlessly integrate SAN technology into a company's IT infrastructure; all this while protecting your investments in application resources, servers, and storage.
- ☑ IBM's SAN solution uses Fiber Channel architecture for connectivity and device-level management.

## Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the “Ask the Author” form.

**Q:** What is NAS?

**A:** NAS stands for network attached storage, and describes a device that is attached to a LAN and uses a communications protocol to provide file access functionality.

**Q:** What is SAN?

**A:** SAN is a network, much like a LAN, that exists solely for storage-based traffic. It interconnects storage devices with hosts to allow for data access and storage functionality, and incorporates numerous features that allow for complex data-sharing solutions.

**Q:** How can we convince non-IT executives of the need for a storage infrastructure?

**A:** The impact and features a SAN can provide is more far-reaching than your IT budget. SANs can affect your core business, regardless of what that is. If you're in e-commerce, SANs should increase your availability, your system up time, and the functionality that you can provide to your customers. If you're looking at backups, SANs should improve your uptime and your restore time. Assess what your needs are, what benefit you're providing, and you should be able to provide a monetary benefit that's more far-reaching than your IT expenditure.

**Q:** What are some of the concerns when deciding on the right storage solution for my organization?

**A:** You should be concerned with host independence, vendor support, security, legacy support, system availability, and price versus performance when you are planning your storage solutions.



**Q:** What is the difference between synchronous and asynchronous mirroring?

**A:** Both of these techniques allow data stored at one site to be mirrored at another site. Synchronous mirroring writes the stored data to both sites at the same time, which creates a 10-kilometer distance limitation between the sites. Asynchronous mirroring will allow the data to be queued and buffered before transmission to the second site, in order to alleviate network congestion and remove the 10-kilometer distance limitation.

**Q:** What is RAID?

**A:** RAID stands for Redundant Array of Inexpensive Disks, and is a technology that allows data to be placed across multiple disks in an array in order to present them as one single logical disk. Depending on the version used, RAID can use parity and disk mirroring to provide fault tolerance and error checking, and can significantly improve the speed of data access

**Q:** How can I determine if the SAN products I buy are interoperable and conform to open standards?

**A:** You have to look at openness and interoperability on two levels. Just as it is in the LAN world, physical connectivity is going to go away as a problem. Higher up in the protocol stack with management applications, you are going to have to do a reality check. You're not going to see much convergence there for a while, because that's how vendors differentiate. You won't, for instance, see EMC supporting a remote data connection to a Hitachi disk storage system on the other end any time soon.

## ASP Security System Provisioning

### Solutions in this chapter:

- Security Policy
- Security Components
- Security Technologies and Attacks
- Prevention Techniques
- Capturing Evidence
  
- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

## Introduction

Security is a primary concern for many application service provider (ASP) subscribers, whose fear of inadequate security is the biggest barrier to an ASP's growth. In fact, one of the most important catalysts to market acceptance for an ASP is to demonstrate that it is addressing all of your customer's security issues with your application or service.

The notion of security is certainly not new. However, ASPs must now provide many of the security controls and mechanisms that were previously neglected by Internet service providers (ISPs). Many ISPs assumed no responsibility for security, as they were only providing bandwidth to their customers.

With the advent of high-speed, always-on connections such as digital subscriber line (DSL) and cable modem technology, millions of individuals and organizations have joined the Internet community. Of these millions of new hosts, very few have gone to the trouble of securing their systems in any way, shape, or form. Although these hosts may not seem to contain data that would be of much interest to an attacker, they do make for a very easy target. These systems can be used as training grounds to help hone attackers' abilities, or as testing grounds where new techniques can be tested and hardened. Even worse, an attacker might compromise one of these "lowly" hosts just to add it to his or her arsenal of "weapons."

Today, attack technologies are developing in an open source environment that allows nearly any individual to improve upon older or more archaic cyber attacks. There are countless applications and scripts currently available that will allow the average Internet user to launch cyber attacks upon whomever he or she feels like at that particular moment.

With user demand for bigger and better applications at an all-time high, many applications are rushed through production and are not thoroughly tested. This makes for applications that are "buggy" and have "holes" that are susceptible to malicious attack. In addition, very few programmers understand the intricacies of security, and tend to write insecure code that can be easily attacked and compromised.

Since the Internet transcends all geographic boundaries, it is important for us to design tools and implement security solutions on a global basis. In fact, many of today's cyber terrorists are from foreign countries, many of which are trying to gain some shred of notoriety.

Most Internet-oriented publications these days seem to always include an article or story on computer crime or abuse. The recent distributed denial-of-

service (DDoS) attacks are prime examples of potential security problems. In fact, in 2000, the Yankee Group reported that the total cumulative revenue lost due to DDoS attacks that were targeted on Yahoo!, eBay, Amazon.com and other Web sites was in excess of \$1.2 billion.

In the same year, the Computer Security Institute/FBI Computer Crime and Security Study found that 273 organizations reported \$265,589,940 in financial losses as a result of computer-oriented crime in 1999. The Computer Security Institute created a *2000 Computer Crime and Security Survey*, which was produced in association with the FBI. This survey reported that 90 percent of its respondents had detected computer security breaches, and approximately 27 percent had detected DoS attacks.

Here are some other highlights from the *CSI 2000 Computer Crime and Security Survey*:

- Ninety percent of respondents (primarily those considered large corporations and government agencies) had detected computer security breaches to their networks in 1999.
- Seventy percent of respondents had reported a serious computer security breach, other than computer viruses, laptop theft, or employee “Net abuse.” This comprises theft of proprietary information (internal and external), financial fraud, outside system penetration, DoS attacks, and sabotage of data or networks.
- Seventy-five percent acknowledged that they had experienced financial losses due to computer breaches.

The study also mentions that the average annual loss reported over the last three years was huge. The problem is that much loss goes unreported to avoid negatively affecting the standing of the affected organization within its market.

Computer crimes do occur, so obviously the risks are real, and the costs are high. You should strive to minimize these risks by implementing sound security policies and practices to which your users must adhere. When building an ASP, one of your goals should be to protect your systems and develop strong security procedures and policies.

## Designing & Planning...

### **Build Customer Confidence in Your Security System**

To have your customer trust your security system, you should be able to disclose your security policy, especially the procedures for incident response, and provide the customer access to your security logs.

## Security Policy

An ASP needs to develop a general security policy that addresses how it manages and maintains the internal security posture of its infrastructure. Issues such as password management, security auditing, dial-in access, and Internet access are some examples of the areas that should be addressed in a security policy. The policy is the written manifestation of current security requirements and guidelines, as well as procedures that your ASP consistently uses.

Consistent policies will give clarity within the ASP about what steps to take to ensure a minimal amount of security. If the ASP is to see immediate improvement with its security position, establishing security policies is the logical step to follow assessment, and should be initiated as an adjunct to security planning.

As the plan for security management unfolds, the specific elements within the environment may change. As changes occur, the policies should be reviewed and modified to ensure that they communicate the current plan for protecting your ASP environment. Security policies should be reviewed at least every six months to verify the validity of the policy, and they should be updated every time the policy changes regardless of the reason. Therefore, security policies should be a continual work in progress.

## Developing a Security Policy

To develop a comprehensive security policy, you will first need to understand what it is that makes for a good security policy. In general, a security policy defines how an ASP manages, protects, and distributes sensitive information and resources. Any ASP, before connecting to the Internet, should develop a usage policy that clearly identifies the solutions they will be using and exactly how those solutions will be used.

First, the policy should be clear, concise, and understandable, with a large amount of flexibility, and some type of built-in mechanism that allows for periodic revisions and alterations as changes become necessary.

Second, you will need to define the requirements to which the security policy will adhere. To provide this, it will be necessary to draw on your usage policy, and to use it as a guide for defining the security policy. This is necessary to maintain the required functionality while providing the security function. Your requirements should include the external customer demands as defined within your service level agreements (SLAs), external legal requirements concerning security, external supplier security policies, your internal security policies, and other security policies that relate to integration of customer environments with your company.

Third, you need to understand what needs to be protected. This might include, but not be limited to, computer resources, critical systems, sensitive systems, customer and company data, critical data, sensitive data, and public data. To help you evaluate your individual system needs, it would be helpful to make a list of all the nodes in your network, and to designate each of these with a level of security.

For instance, a public machine that poses few consequences if it were to become compromised might be considered low security; a Web server might be considered medium security; and your financial databases might be considered high security. Be careful when designating low-security systems, though. Just because a system may not contain any sensitive data does not mean that they are not a threat; if they have access to devices that do include sensitive data, they might be used as a springboard to access other systems within the network.

Fourth, you need to define the security policy guidelines. To accomplish this, two policies should be written; the first should consist of a high-level policy written from the customers' perspective, and should be a simple document that gets directly to the point. You should base this document on security rationale, and should have very little technical information.

A second low-level policy should also be written for security implementers, and should include detailed technical descriptions of procedures, filtering rules, and so forth. This document should clearly and concisely outline the exact security procedures, and should only be viewable by those who require the information. If such a document were to become publicly accessible, it could be used against your systems maliciously by identifying possible holes in your security policy and thus displaying methods into your network.

For instance, if you are using packet filtering to only allow traffic from a specific network, it might be possible for a would-be cracker to spoof an IP address that is in the accepted range in order to compromise your systems. Because of this, it is best to keep your security policy very secure.

Finally, you must ensure that your security policy is based on actual customer situations, while remaining clear, concise, consistent, and understandable. Furthermore, to ensure a good security policy requires a periodic evaluation of the effectiveness of the current security systems, as well as periodic evaluation of the actual system configurations, or at least the security relevant components.

Sometimes it may even be beneficial to hire a third-party security firm to provide an unbiased evaluation and assessment of your security systems. In many cases, they may discover issues that you did not, and they might be able to suggest possible fixes for some of the issues they encounter.

In addition, it is sometimes easier to sell your customers on your security posture if an evaluation was performed by an outside security organization. It could at least help to instill your customers with confidence in your organization.

## Privacy Policy

An extension of the security policy is the *privacy policy*. Basically, the privacy policy should state what data the ASP considers to be confidential, and how that data can and cannot be used. For instance, you will probably need to define a privacy policy that only allows certain members of your staff to access your own internal data.

At the same time, you will need privacy policies that guarantee that your customers' data is partitioned, and is only accessible by users they have predefined. In this scenario, it will be necessary to govern exactly which users have access to a particular partition of sensitive data, and to deny all other users access privileges. Not all of the data may be sensitive, though, and some of it may not fall under your privacy policy.

Just as it is important to define what information should be kept private, it is also important to define data that will be considered public. This is important, since most ASPs post their privacy policies on their Web sites, or distribute them to their customers in some way. Because of this, it will be necessary to inform your customers of data that will be considered publicly accessible, as well as data that will be considered secure and private.

Unfortunately, a recent study by the Electronic Privacy Information Center (EPIC), a Washington-based privacy research group, indicated that while many Web sites post privacy policies, few actually support their implementation. In

December of 1999, the EPIC released a report entitled “Surfer Beware III: Privacy Policies without Privacy Protection” in which it claimed that only a handful of the 100 most popular shopping Web sites provide only adequate privacy protection for consumers, and many track purchases and online habits.

EPIC also determined that none of the sites adequately addressed the Fair Information Practices, a set of privacy protection principles outlined by the Federal Trade Commission (FTC). Therefore, it is critical to not only develop the privacy policy, but to implement it as well!

## Security Components

As an ASP, to validate both the security policy and the privacy policy, a review of the various security mechanisms and methods used to implement those policies is required. At a minimum, the following security components should be considered:

- Authentication
- Confidentiality
- Incident response
- Security auditing
- Risk assessment

## Authentication

One of the most important methods to provide accurate security is the ability to authenticate users and systems. In fact, all of your security mechanisms will be based on authentication in one way or another. As an example, you will need to authenticate users and nodes that access data on your systems. The authentication might take the form of a username and password, or an access list that governs access from a particular system’s IP address to another system’s IP address.

You may even use a different method entirely, or a combination of methods. Regardless of the method used, it is apparent that without the ability to guarantee or reveal the authenticity of a user or host, it is impossible to guarantee security. In fact, the success of your security mechanisms will hinge greatly on the methods of authentication they incorporate and you employ throughout your network.

## User Authentication

A requirement for any ASP is the ability to positively identify and authenticate users. Depending on the level of security required, the mechanisms to support



this requirement can range from identifying users based on usernames and passwords, to personal identification numbers (PINs) and digital certificates.

### *Usernames and Passwords*

The use of usernames and passwords is one of the most ancient of all authentication schemes. I am sure at some point you have had to enter a username or password to gain access to a resource, or even to log in to your own personal computer. This being the case, you are probably already familiar with some of the security concerns associated with the use of passwords such as not to share them with others and to keep them private.

To accomplish this, you are aware that you are not supposed to write your password on a piece of paper that is taped to your monitor, or that you should not use a password that is easy to guess, such as your first name. However, just because you understand these cardinal rules, it does not always follow that others will too. Because of this, it is always important to set password guidelines for your users, and make certain they adhere to those guidelines.

When evaluating identification and authentication mechanisms, you need to consider both the mechanism and the implementation. A standard user ID and password scheme should have a minimum password length of at least eight characters, and require passwords to be nondictionary words. In addition, the implementation should limit unauthorized access attempts and, at a minimum, after a fixed number of failed attempts, lock out the account for some specified period. If the account is locked out multiple times, it should be locked until an administrator can speak with the owner of the account.

### *Personal Identification Numbers*

A *personal identification number* (PIN) provides another mechanism that you can use to enhance the security of a standard username and password system. In most implementations, users log in to an ASP with their username and password. Once validated, the users are asked to enter their PIN, which is usually a numerical value that is predefined and known only by the user and authentication mechanism. The PIN provides an extra level of access control, but can still be overcome fairly easily.

### *Digital Certificates*

Deploying digital certificate technology would be a more robust access control mechanism. Today, the trend seems to lean toward a digital certificate-based solution that not only validates the user, but also enables the establishment of a

session encryption key to support confidentiality of the transaction once the user is authenticated.

If you use usernames and passwords solely for authentication services, you may be exposing your ASP to an easy attack. If, for instance, an attacker were to gain access to a system by compromising a username and password, he or she would have access to all resources for which the account is privileged. This might allow the attacker access to a single host or numerous hosts in your network. It could also give him or her the opportunity to access and alter data, as well as wreak havoc on your systems and their functionality.

There are numerous methods an attacker can use to bypass password-based security mechanisms, the most popular of which are network sniffing and brute force.

### *Network Sniffing*

*Network sniffing* attempts to acquire clear-text passwords by exploiting the exchange of passwords between systems. If you use unencrypted, clear-text passwords to authenticate users, these passwords are plainly visible to anyone who has access to the data packets containing the password information. If this authentication is taking place across the Internet, it is impossible to guarantee the path these packets will take, and the packets will be visible to nodes and users in any of the networks that the packets traverse.

This means that anyone in between your system and the authenticating party will be able to capture and search these packets for usernames and passwords. Since every one of these packets contains the source and destination IP address, it will also be possible to identify both the system attempting to authenticate, and the system that requires authentication. If this is the case, an attacker may be able to bypass your authentication scheme by providing the correct username and password, thus gaining access to your systems with all of the privileges the account possesses.

Some of the more ambitious hackers will even capture encrypted passwords, and use software to decrypt them. Since the source and destination IP addresses are plainly visible in the packets, they will also be able to identify the systems involved in the exchange of authentication information.

This means that even though you employ an encrypted password mechanism, it is still possible for an attacker to “sniff” these passwords, and gain access to your systems. If you do plan to rely on password encryption, be sure to check into the strength of the encryption scheme employed. With this information, weigh the

chances of a particular password becoming compromised with the value of the systems and data you are attempting to protect.

### *Brute Force*

In addition to sniffing, another widely familiar method of bypassing your password mechanism is to use a brute-force attack. In this case, a software application or script works its way through dictionary entries and word-matching libraries to identify words that can be matched against a given username to allow access into a system.

The word-matching capabilities of some of the software applications are impressive, and with the power of the personal computer today it is sometimes possible to attempt millions of password and username combinations in a single second. This means that if someone has a password that is short, or based on a dictionary word, it is possible for a computer to “guess” the password in a matter of seconds.

To remedy this problem, it is necessary to set very stringent password guidelines, which should include a minimum password length of at least eight characters, and a combination of letters, numerals, and symbols. Still, even if you had a 20-character password that consisted of all these different types of characters, it would still be possible for an attacker to crack the password—it would just take a lot longer. Hopefully, though, you will use additional measures that can alert you to invalid login attempts, and inform you if someone has tried millions of different username and password combinations.

## IP Addresses and Spoofing

When most of us think of authentication, we think of usernames and passwords. However, this is far from the only method of providing authenticity information. There are, in fact, numerous ways to provide authentication services in an IP network; the second most popular method of which is through IP addresses.

IP addresses are used to identify hosts on a network, and allow for a method of addressing packets for delivery to a given host. An IP address can be easily compared to a street address. For instance, when sending a letter to a friend or company, you must first fill out an envelope with their address and include your own address in case there is a problem with delivery or the recipient would like to send a response.

The addresses on the envelope identify a particular location, and are used to deliver mail to the correct home or business. In much the same way, when a computer accesses a host across an IP network, it addresses every data packet it

sends with a “destination” address that identifies the host. It also includes its own “source” address in each packet, to allow for responses to be sent. Because of this feature of IP networking, we are able to identify hosts and networks using their IP addresses.

The problem with is that an IP address can be spoofed. This can be accomplished by modifying the source IP address of each individual data packet sent to a host, or by routing traffic through a third-party organization. Regardless of the method used, spoofing allows an attacker to make it appear as if a given packet or connection came from some other computer or organization.

The organizations spoofed are sometimes very curious and can commonly include NASA, the White House, and colleges or universities. By routing from some other source, hackers can mask any audit trail back to them or bypass security mechanisms.

Probably the most common IP addresses used to spoof data packets are ones that are local to the system being attacked. In some cases, a particular system may be configured to only allow data from nodes that are within the same subnet, or possibly to not authenticate a user by password when the access is from another local device.

In such a case, an attacker might be able to spoof his or her IP address to appear as if the data was coming from a local system, thus bypassing security. This trick has been in widespread use since the early 1990s. Many of today’s firewalls and other security devices incorporate technologies that identify and block spoofed data packets. However, the best method for stopping this type of attack is to implicitly block data packets that originated from the Internet whose source IP addresses match the subnetworks contained in your network.

Access controls are generally associated with identification and authentication, but this may or may not be the case, depending on the type of services being offered by the ASP. Standard role definitions may further limit or control access privileges. As an example, a company may have a corporate or customer logon to an ASP service. This may give access to a number of applications that require further access control mechanisms based on the role of a specific type of user.

## Confidentiality Protection

Confidentiality is usually associated with data encryption mechanisms such as Secure Socket Layer (SSL) or Data Encryption Standard (DES), and targeted at protecting data as it traverses across a network, such as the Internet. An example of this could be a secure Web page that uses SSL to encrypt sensitive information

that a customer provides, or a virtual private network (VPN) tunnel that uses DES to encrypt data that is sent between two sites across the Internet.

Although these are two very different implementations, they both allow data to be encrypted and decrypted by the receiver using an encryption key. This might seem like an excellent solution to confidentiality issues; however, it could introduce latency to your data flow. This stems from the fact that the data needs to be encrypted on one end, and decrypted on the other end.

This means that the speed of the cryptography will be highly dependent on the strength of the mechanism you are using, as well as the hardware or software you employ to handle the cryptography. In general, a more secure confidentiality mechanism will be inherently slower than a less secure method; however, it is always possible to purchase dedicated hardware that can significantly improve cryptographic performance.

It is not good enough to implement any old encryption method and trust that it will prevent anyone from viewing your sensitive data. The fact is that if your data is traveling over a shared medium, such as the Internet, it is highly likely that the data packets can be intercepted and recorded. An attacker may not be able to decrypt your message in real time; however, once recorded, he or she can play back the data flow and dedicate system resources to cracking the encryption key, thus making the data once again intelligible.

This might take hours, or years, depending on certain factors of the encryption mechanism and the amount of resources dedicated to crack the data. Essentially, your decision and implementation will make this task either easy for the attacker or so difficult that it will not be worth the attacker's time.

Because of this, you might decide to employ the strongest level of encryption possible; however, as we mentioned earlier, the stronger the method, the slower the performance, and the higher the associated costs. Ultimately, you will need to be realistic and compare the sensitivity of your data with the need for performance and cost-efficient operation. If you can accomplish this with a hint of paranoia and a dash of prudence, you should be fine.

## Key Length

When evaluating different confidentiality mechanisms, a company should consider both the strength of the mechanism and the implementation. The strength of a mechanism is dependent on several factors. The first is the length of the encryption key. In cryptography, an encryption key is a variable that is applied to plain-text data using an algorithm to encrypt the data.

The key is predefined and shared between both endpoints to give only those systems the capability of encrypting data to be sent, and decrypting the data they receive. If another system attempts to decrypt data without the encryption key, it will be unsuccessful. There is, however, a possibility that someone might be able to crack the key, and this is where the length of the key really matters.

A longer encryption key will be exponentially more difficult to crack when compared to a shorter key. The keys are measured in bits, and each bit can only be in one of two states at any given time: off (0), or on (1). Because of this, the formula for computing the total number of possible permutations of an encryption key with  $x$  number of bits is  $2^x$ .

This means that the total number of possible permutations for a 56-bit key is 72,057,594,037,927,936. Obviously, 72-quadrillion possibilities might make it a little difficult to use a brute-force method to arrive at the correct encryption key. While this may seem like an extremely large number, with today's personal computers, it is actually possible to cycle through all the possible permutations in a matter of months or days.

Supercomputers and specialized devices have been known to crack this level of encryption in a matter of hours, and sometimes within minutes. On the other hand, a 128-bit encryption key would have a possible 340,282,366,920,938,463,463,374,607,431,720,000,000 combinations. This number is so large that it is difficult for us humans to relate to it. Computers, on the other hand, are still capable of cracking code with this number of possibilities; however, it is going to take an extremely long period of time to accomplish this.

Because 128-bit encryption is so strong, there are stringent rules that apply to the export of this technology. Currently, 128-bit encryption and higher is considered unbreakable, and should remain that way for some time.

Encryption keys are not always so easily evaluated, however. For instance, triple-DES (3DES) uses three separate 56-bit keys that are combined when performing the encryption algorithm. In this case, there is not a single 168-bit key; instead, the three separate keys are appended to each other in any possible order. This means that the formula for deriving the total number of possibilities would be  $(2^{56})^3$  for a total of 432,345,564,227,567,616. This number is quite larger than your normal 56-bit DES encryption.

## Types of Algorithms

Besides the size of the encryption key, several other factors determine the overall strength of an encryption technology, such as the type of encryption method

being used. There are two distinct types of key-based encryption algorithms, *symmetric* and *asymmetric*.

### *Symmetric Algorithms*

Symmetric algorithms use the same key for both encryption and decryption. The key can be assigned, or generated randomly. However, in both cases, the key will need to be known by both parties before they will be able to encrypt and decrypt data. With some implementations of symmetric keys, the preshared keys are not exactly the same. However, in these cases, the second key is a derivative of the first key, and can still be cracked if either key is known.

### *Asymmetric Algorithms*

Asymmetric algorithms are also referred to as *public-key algorithms* or *public-key cryptography*. In this encryption method, a public, or known key is used to encrypt data that can only be decrypted using a private, or unknown key. This type of technique is usually associated with very large implementations. The most common use for this type of architecture is to encrypt and decrypt e-mail messages that are sent between two parties.

In this case, the sender finds the recipient's public key, and uses that to encrypt the e-mail message before it is sent. When the recipient receives this message, he or she uses a private key, which might be a password, to decrypt the message. In this way, the public key is known and accessible to anyone who would like to send an encrypted message to the recipient. However, once the message is encrypted, even the person who encrypted the message will be unable to decrypt it without the correct private key.

## Further Cryptographic Considerations

Besides the type of key and its length, several types of factors will determine the overall strength of a given encryption method. For instance, whether a key is user-definable could affect the possibility that a given key could be cracked.

For instance, if you are using a key that was built around a user-definable password, it may be possible to use social engineering to actually figure out the key, without applying any type of brute-force tactics. When considering the ramifications of this, it is probably not a wise idea to use any type of user-definable keys to encrypt or decrypt your data.

Instead of a user-definable key, it might make more sense to use a randomized key. In this way, it is impossible to use social engineering to crack the key.

However, the true randomness of a key might be questionable. Conventional random number generators, like those implemented in most servers and personal computers, are designed with statistical randomness in mind, instead of cryptographic randomness.

In these cases, it may actually be possible to crack a particular key based on the frequency of random numbers; in truth, the numbers are not truly random. On the other hand, a cryptographic random number generator is capable of generating truly random numbers. This is accomplished by using an external source to provide the random effect, such as the noise obtained from a semiconductor, the least significant bits of an audio input, or the intervals between device interrupts or keyboard “clicks.”

In addition to these concerns, you should also look into the cryptographic “period,” or how often the key is changed. If the key is user defined, chances are it will never change until you manually change it. However, if you use randomly generated keys, they will most likely change periodically. They might change based on a predetermined interval or on a session-by-session basis. Regardless, the more frequent the changes, the more secure the data will be.

## Incident Response

As mentioned earlier, you should always design your system with the premise that your systems will be attacked and eventually compromised. This is especially true when you operate an ASP, since your name will be known throughout many circles, and I guarantee someone will want access to the data that you house on your systems. This means you will need to develop a plan to successfully combat an intrusion once it has been accomplished.

Your plan should describe the exact steps to be taken by your staff in the event of an intrusion, and the order in which they should be accomplished. Such a plan should include a method of thoroughly documenting the intrusion and the procedures used to combat the intrusion. This documentation is important and may be used at a later date to further identify, and possibly incarcerate, the perpetrator.

When responding to an incident, the first thing you will need to do is define the attack. There are a couple of questions you should ask yourself, such as “Who is the attacker?” and “What are they attempting to accomplish?” Once this is known, you can begin to combat the problem.

After identifying the intruder, your next step will be to block the attacker from accessing your network and resources further. This might be accomplished



relatively simply, or might be a difficult task, especially if the intruder has been allowed enough time to sufficiently plant him or herself in your systems. If an attacker has been identified, it may be possible to filter the intruder using an access-list in a router, or an additional filter in your firewall.

This should put an immediate stop to the intrusion, but will not provide a good permanent solution. To combat this filtration, the intruder will more than likely use a different IP address, by either employing a spoofing technique or performing the attack from another system to which he or she has access. Regardless of the method, if a different IP address is used, the intruder will be able to bypass your access-lists, and resume the intrusion upon your systems. Because of this, you may need to increase the monitoring of your systems, and make sure that your intrusion detection systems (IDSs) are operating effectively.

Next, you will need to identify exactly how the intruder gained access to your systems in order to enact a solution that will more permanently disable the intruder from accessing your systems. In effect, you will need to “plug the holes” in your system, so that the same method cannot be used a second time to bypass your security and gain access into your systems.

For instance, if an attacker has gained access by using particular username and password, you may need to disable the user account. At a minimum, you should at least change the password on the compromised account.

You will also need to assess the situation very carefully. Again, if the intruder used a username and password combination to gain entry to your system, you must assess whether the intruder might have also gained access or knowledge of other usernames and passwords that can be used to bypass your security mechanisms. Did the intruder have enough time to sniff passwords in the network, or to actually steal data that contains valuable login information?

You should look for any traces an intruder has left behind; especially look for Trojans or backdoors into your network. It will also be very important to address any changes that may have been made to server and device configurations, and look for any access or alteration of data that may have occurred as a result of this intrusion.

Any company can be hit with bugs, glitches, and security incidents. The question is not whether you will be attacked, but rather, *when you are* attacked, will you be able to survive the incident, or repair your systems quickly?

As an ASP, you will more than likely need your own emergency response team. This team will be able to implement and test your security mechanisms on a daily basis, and will be able to provide around-the-clock security for your systems. You will need to plan and deploy your security mechanisms, and keep them

up to date and operating efficiently. You should make it your goal to block most attacks, and identify and neutralize the attacks that penetrate your systems quickly and effectively.

## Security Auditing and Risk Assessment

It will be necessary to reassess your security mechanisms from time to time, and perform risk assessment on all your servers and network devices. You will need to quantify and qualify any security threats, and look for previously undiscovered vulnerabilities that could be used by an attacker to gain entry into your systems. As mentioned earlier, you will need to keep your security systems up to date to effectively combat would-be attackers.

In addition to this, however, we recommend auditing your security mechanisms on a consistent basis. As new devices are added and changes are made to the system, it will be necessary to test your security mechanisms, and be on the lookout for ways to breach it.

When auditing your systems, it will be necessary to audit your individual servers, network equipment, IDSs, and firewalls. This can be quite a daunting task, and will require an individual, or several individuals, with a good deal of security expertise to effectively audit all of these systems. You may already have these individuals on-hand, and it might be their full-time job to perform security analysis and intrusion detection. However, most companies will not be able to afford to employ an entire army of intrusion warfare specialists. In these cases, you may need to resort to other auditing tactics.

There are some software applications that can be used to audit your systems, such as Network Associates' (NAI) CyberCop Scanner. This type of application can simulate attacks on your network and servers, and look for vulnerabilities and ways to compromise your systems. It will then provide the user with a full assessment of your systems and network security mechanisms.

The report is usually prioritized to give an indication of the seriousness of the vulnerability, and, in many cases, the report will even offer suggestions on how to fix or plug certain vulnerabilities. This can be an effective method for periodically assessing your security mechanisms, but if the software application is not current or up to date, it may not attempt tests and intrusions using the most state-of-the-art techniques.

Moreover, the tests that such a software application uses are generalizations, and do not include the same logic a human possesses. In most situations, it will be necessary to also use human judgment to fully assess your particular situation.

It would probably be a good idea to use an external organization to assess your security mechanisms. It is likely that an outside source will have more collective security knowledge, especially if that is the function and nature of their organization. In addition, they will be able to make unbiased assessments and recommendations. It is likely that they will also see vulnerabilities that were not recognized by internal sources.

## Security Technologies and Attacks

ASPs must deploy the best security technologies. Strong encryption is important, whether in the context of an SSL browser connection or a VPN connection. ASPs need to employ authentication systems that are appropriate to the sensitivity of the data, which sometimes may mean username and password combinations, and some instances may even call for hardware tokens, digital certificates, or even biometrics.

It will most likely be necessary to use IDSs and firewalls to protect your systems. In some cases, you may even need to secure the data as it travels between your network and your customer's local area network (LAN). In order to accomplish these tasks, it will be necessary to use highly advanced security technologies that allow you to effectively secure your systems, and ward off attackers.

## Virtual Private Networks

With the proliferation of the Internet today, almost everyone has access to the Internet. High-speed Internet connections are generally simple to purchase, and are easily installed and integrated into an existing network. Yet the question remains: How can we safely transmit our data to a trusted destination across the Internet, and insure that it is not hijacked or read in transit?

The answer is virtual private networks (VPNs). As VPNs are being deployed at break-neck speeds and in almost every company, this book will assist you in determining the proper method of implementing a VPN that fits your needs.

The two basic methods of VPN access are LAN-to-LAN VPNs and remote access VPNs. The LAN-to-LAN VPN is used to create a permanent or "nailed up" connection between two or more sites. This effectively creates a "tunnel" across the Internet, allowing offices and remote locations to share data safely.

The configuration and rules at each VPN endpoint determine what traffic will be permitted to traverse the VPN, and how and/or if it should be encrypted. By combining predefined rule sets with encryption, you can run a satellite office with a single network connection for Internet, office wide area network (WAN),

and Voice over Internet Protocol (VoIP). This provides a great cost savings over the traditional business model, which required separate lines for Internet, WAN connections, small offices home office (SOHO), and voice services.

Remote access VPNs are used to connect individual users (usually dial-up/cable/DSL users) who connect using IP addresses that are unknown or change frequently. These users must run VPN client software on their PCs that can contact a centrally located VPN endpoint, which negotiates authentication, virtual IP addresses, and other connection-specific parameters. This is most commonly deployed for telecommuters who work from home and for remote network support.

Many types of VPN endpoint equipment (VPN concentrators, routers, etc) are capable of terminating both methods of VPN access simultaneously. There are numerous considerations when choosing a VPN concentrator, such as: How many LAN-to-LAN connections are you planning to support? How many remote access connections are you planning to support? How many of these remote users will access the system at any given time? Will these coincide with the site-to-site VPN connections? What type of authentication will you be using? What types and levels of encryption will you be supporting? What types of clients and software will you be supporting? How much future growth will you require?

Once you answer these questions, you can begin to select a VPN device that fits your network. Since VPN concentrators are configured to only accept encrypted, authenticated connections, and do not allow any other connections to their external interfaces, these devices are generally installed in parallel to a firewall. If the concentrator were placed inside your network, you would need to open conduits on the firewall from any source, which would defeat the purpose of VPNs altogether. However, if you are only performing LAN-to-LAN connections and you will always know the source address, then it would make more sense to install the VPN concentrator behind your firewall, preferably in a demilitarized zone (DMZ) (also known as a *bastion network* or a *dirty network*).

From here, you must decide what clients will be supported, and configure the VPN concentrator accordingly. Some concentrators support proprietary client software, while others work with the client software already built into many Microsoft Windows products.

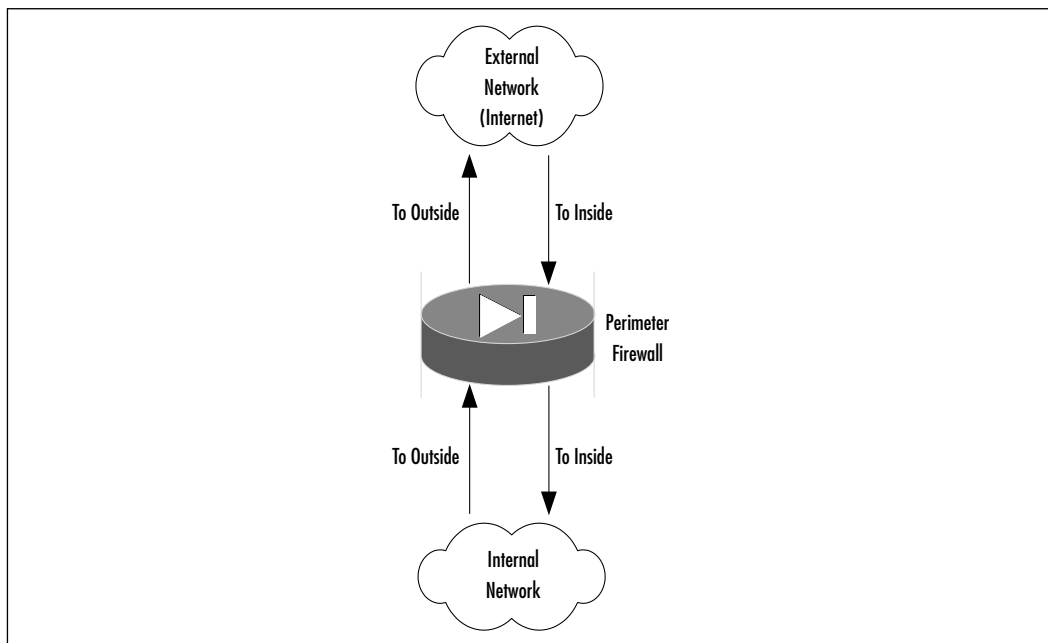
## Perimeter Firewalls

Probably the most common method of providing base-level network security is through a perimeter firewall. A perimeter firewall is a device, or software application, that controls access in to and out of a given network. To accomplish this

successfully, all data must flow through the firewall, making it operate in much the same way as a bridge or router. In fact, most routers incorporate very powerful firewall features.

Most perimeter firewall implementations consist of firewall software that is installed on a server or specialized “appliance.” The server or appliance sits between two or more networks and is capable of permitting or denying data based on a user-defined configuration. See Figure 6.1 for an example of a perimeter firewall.

**Figure 6.1** Perimeter Firewall



There is a variety of firewalls on the market today, and each offers numerous features and functions. Some of the offerings will have robust logging features, and others may have excellent monitoring and reporting functions. There will be a variety of bells and whistles from which to choose. However, the majority of all perimeter firewall products will use at least one, or a combination of, the following methods to allow or deny data passing through its interfaces:

- Stateful inspection
- Packet filtering

## Stateful Inspection

Stateful inspection provides for the most robust of all firewall features. Using stateful inspection technology, each packet traversing the firewall is deconstructed and checked for suspicious activity before it is allowed to pass through the firewall device. This allows the firewall to catch attacks that would otherwise go unnoticed by a packet-filtering device, since it examines the contents of every packet before making security decisions.

Stateful inspection technology is capable of deciphering a packet using all seven layers of the Open System Interconnect (OSI) model. The firewall intercepts each packet, and derives “state” information by building a state and context database. This means that the firewall is actually capable of “understanding” the function of a particular application and conversation.

A firewall using stateful inspection compares the state of each packet against the context of a given application. For instance, if an application requires authentication information, the firewall will see the authentication request being made to the client system when it deconstructs and inspects the packet. If the client system responds to this request with anything other than an authentication reply, the packet will be deemed “out of context” by the firewall, and therefore will not be passed to the requesting server or application. To arrive at this conclusion, the firewall needs to understand the state of previous and current packets, and derive the context of the conversation and applications.

Using stateful inspection technology, it is also possible to gain state information from protocols that are not connection oriented, such as User Datagram Protocol (UDP) and Remote Procedure Call (RPC). Since the firewall builds a database of information regarding all the packets traversing its interfaces, it is able to keep track of packets that are not connection oriented. This provides collective information against which further packets and communication attempts can be compared.

It is obvious to see how stateful inspection provides an excellent security solution. As long as the firewall is capable of understanding the state and context of the applications and communication stream, it is nearly impossible to bypass the security mechanisms using Application-layer attacks. There are, however, a few downsides to this technology.

It is extremely important to keep the application database for a stateful inspection engine up to date. The information contained in this database is used to “understand” your applications and the upper layers of the OSI model (as in Chapter 1, “An Introduction to ASPs for ISPs”). Without the most current

information, the firewall might allow access that it should not, or even disallow access that it should allow.

Since stateful inspection must break down each packet and apply a certain level of artificial intelligence (AI), it can cause a significant performance decrease. Although the speed provided by today's firewalls is enormous, so is the amount of data traversing our networks. Due to its nature, stateful inspection technology is slower than typical packet filtering. However, given the level of security it provides, that is to be expected.

If you need to supply high-speed (over 100 Mbps) throughput, you will need to opt for a load-balanced stateful inspection firewall solution. To provide the throughput you require might prove a bit costly, though. If you require less security and more performance, you should look into high-speed packet-filtering devices.

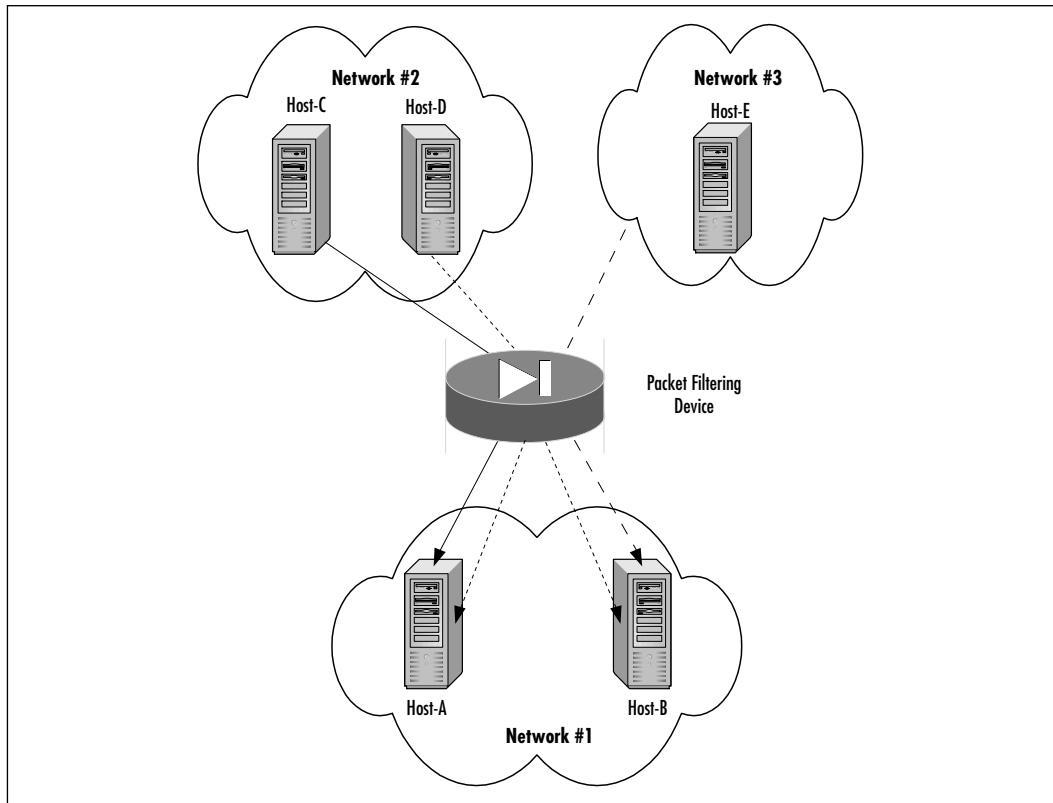
## Packet Filtering

A firewall can screen data as it flows into and out of your network in a number of ways. The most common of these forms is packet filtering. Packet filtering enables a device to permit or deny packets based on the source and destination addresses contained in a given packet, the type of packet, the ports used, and the direction of data flow.

A packet-filtering device accomplishes this using access-lists or preconfigured rule sets that define which networks and nodes data is allowed to flow between. For an example of packet filtering, refer to Figure 6.2.

In this example, we have five hosts, or nodes, and three networks. Host-A and Host-B are both in Network-1, Host-C is in Network-2, and Host-D and Host-E are in Network-3. The access-lists or rule-sets will dictate which hosts and networks can talk with each other, and the packet-filtering device will deny or permit packets based on these rules.

For instance, in our example, we permit Host-C to access Host-A, but deny it access to Host-B. Also, we are permitting any device within Network-3 to access Host-B and allowing Host-D to access Host-A and Host-B. Additionally, since a packet-filtering device has the ability to differentiate between a new stream of data and a previously established connection, we can prevent hosts from communicating with devices unless they are responding to an established connection. Applying this technique to our example, we can configure our packet-filtering device to prevent Host-A and Host-B from initiating connections to other devices and only allow them to respond to previously permitted and *established* streams of data.

**Figure 6.2** Packet Filtering

When Host-C attempts to contact Host-A, the data will first need to flow through the firewall. The configuration of the firewall, its access-lists, and the direction of data flow will be the determining factor in whether the traffic is allowed to pass between these devices.

If our access-list permits these devices to communicate, packets are allowed to flow to Host-A, and Host-A is able to respond since the connection is “established.” The same would be true if Host-D tried to access either Host-A or Host-B; the connection would be allowed. However, when Host-E attempts to connect to Host-A, our firewall will not permit the packets to reach Host-A, since our configuration does not allow access between Network-3 and Host-A.

In this example, we were denying access based solely on the source and destination addresses and the direction of data flow; however, it is also possible to filter packets based on the type of packet. For instance, using the same example, we could modify our access-lists to only allow Host-C to access Host-A when it is using UDP packets.



Conversely, we could have the firewall block all Transmission Control Protocol (TCP) packets bound for a particular node or network. Most packet-filtering devices will even allow us to configure specific port numbers that are allowed or denied. For example, you might want to allow all HyperText Transfer Protocol (HTTP) traffic to pass through the firewall when the destination is Host-A and Host-B.

At the same time, you may want to permit Host-C Post Office Protocol (POP3) access to Host-A. All other traffic flowing into Network-1 should be denied. With a firewall, this can be accomplished easily by filtering packets based on communication ports that particular applications use. HTTP, for instance, uses TCP port 80, while POP3 uses TCP port 110. Packet filtering allows us to deny or permit traffic based on a combination of traffic flow, source and destination addresses, communications protocols, and communication ports.

As you can probably tell, your access-lists can become fairly cumbersome if not planned correctly. In all of the preceding examples, we have only used a total of five hosts; however, in the real world, you will probably be concerned with hundreds or possibly thousands of networks, and an virtually endless number of hosts. When applying packet-filtering rules, there are usually two options: either deny all traffic, except for that which is explicitly allowed, or permit all traffic except that which is implicitly denied.

### *Explicitly Allow Traffic*

Usually the easiest, and definitely the most secure, method to configure packet filtering, and any security mechanism for that matter, is to deny all traffic except what is explicitly allowed. This is the easiest method since the list of hosts allowed into your network is typically much smaller than the number of hosts to which you will need to deny access.

For instance, you might have 100 customers who each need access into your network across the Internet. It will be much easier to permit only these customers access, and deny all others. If you were to instead try to deny the millions of other nodes that you did not want to have access into your system, you would probably be hard-pressed to write an access-list with millions of entries!

Deny everyone and everything, and allow only those functions that are required to run your business. This just makes common sense. This will also help provide the level of security your ASP will need. By denying all traffic that is not required to run your ASP, you will be eliminating thousands, if not hundreds of thousands, of possible ways to breach your security. When you are configuring a perimeter firewall, this is really the only way to go.

## *Explicitly Deny Traffic*

Permitting all traffic except that which is explicitly denied is typically a very bad way to go. There are usually far fewer hosts and networks that need access into your system than those that do not need access. There are, however, a couple of instances where this may not be the case.

For example, many border routers will use packet-filtering rules that allow all traffic unless explicitly denied. This is usually done when there is a firewall behind the border router. In this case, the border router's configuration will usually deny certain networks to eliminate IP address spoofing.

The configuration should also deny all access to the firewall that sits behind it. In this way, the router will allow all traffic to get to the firewall unless it is spoofed traffic, or an attack directed *at* the firewall.

If you are using a firewall to protect LAN segments from other LAN segments internally, it is many times easier to permit all traffic and deny access to specific hosts. This might be especially true in a LAN environment that requires much functionality between different LAN segments.

In this case, there would be far too many permit rules required to allow the level of functionality required. Instead, it is far easier to deny access to particular nodes that need to have additional security. However, if this is the case, it might make more sense to remove the firewall and use an embedded firewall for the servers that need the additional security.

### Configuring & Implementing...

#### **Know Where Your Enemies Are**

A common false assumption is that the enemy is outside your firewall. While you are building an impenetrable wall around your system, fixing your eyes on the external threats from anonymous Internet outposts, those looking to steal or compromise your data will also be looking to enter the backdoor via social engineering or planted ASP employees.

## Embedded Firewalls

Embedded firewalls are software applications that are installed and run on a computer to guard it against attacks. Depending on the embedded firewall solution in use, they can offer the same level of functionality provided by a perimeter firewall, such as stateful inspection and IP filtering techniques. The difference is that the firewall only protects the computer on which it is installed, which allows for a more “personalized” configuration.

Some operating systems come with embedded firewall mechanisms already installed. For instance, most Unix systems include applications that will allow you to configure IP access-lists and rule-sets. For operating systems that do not incorporate such systems, it is usually possible to purchase a third-party application to provide firewall features.

It is even possible to design your own firewall that could be embedded into a given operating system, but we strongly urge against this. It would be difficult to guarantee the compatibility of such an application, and its stability and effectiveness would be questionable. Instead, try to stick with proven firewall solutions that are simple to administrate and offer the level of security you require.

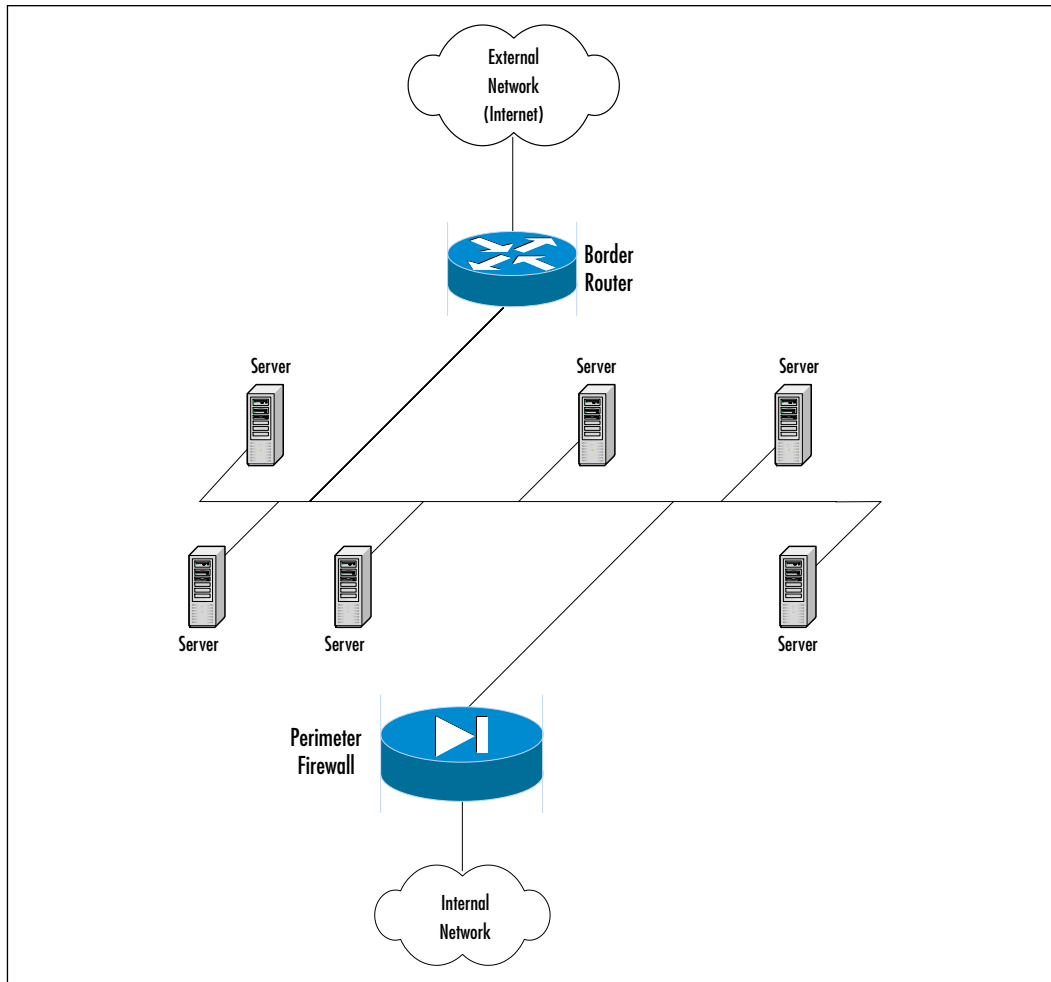
Probably the best feature offered by an embedded firewall solution is that it can protect a system against internal attacks. Since the firewall is installed on the server itself, it can even stop attacks that are coming from the same network segment. All traffic will still need to traverse the embedded firewall.

### *Bastion Network*

Many Web servers are inadequately protected. This is due in part to the design of the Web server and the protocols it uses to communicate with other devices. Web servers are not the only types of insecure servers, though. In fact, most servers that provide “Internet” services are susceptible to attack, such as Simple Mail Transfer Protocol (SMTP) and Domain Name System (DNS) servers. If these devices pose a threat to your internal network, it is possible to place them in a bastion network (see Figure 6.3).

In a bastion network, insecure servers are placed behind a border router, but in front of the firewall. Since these servers are very vulnerable to attack, they would most likely have an embedded firewall installed to protect them. The idea behind this concept is that even if the servers outside the firewall were to become compromised, the networks behind the firewall are still protected.

If this sounds like a good idea to you, you must be ready to repair the damage done to these servers. Even though they may have an embedded firewall

**Figure 6.3** Bastion Network

installed, it may not be able to withstand the barrage of attacks the server may receive. Although this can help to reduce the traffic that is allowed into your network, it is still a dangerous strategy to implement.

When you consider that if one of these bastion servers were to become compromised, the intruder might be able to view all the network traffic flowing to and from your firewall. This information could be used to launch an attack against your internal systems. Moreover, the compromised servers could be used to spoof IP addresses and bypass your firewall and security mechanisms. If you are considering using a bastion network to subvert attacks, be careful, since it might actually increase your chances of an attack.

## Intrusion Detection Systems

No matter how much time and money you spend securing your systems, it is impossible to ensure that you will never have a break-in. In fact, if you were to take this gamble, you will probably end up losing your shirt. Your best bet would be to make it so difficult to compromise your systems that an attacker does not bother, or gives up halfway through the process.

You should also consider taking the stance that your system will be compromised one day, and build a solution that will allow you to identify a break-in quickly, monitor the attacker's actions, protect your most secure data, and eject the attacker from the system. This may not be easy, or cheap, but it will be your best line of defense.

The truth is that for every person who is developing security methodology, there are at least two, ten, or a thousand others who are working to counteract such a security mechanism. The power of the masses will always outweigh the power of the “good,” so you will need to adopt a system of security auditing and intrusion detection in order to best combat these attackers.

Security auditing and intrusion detection systems (IDS) are critical to any ASP, and will give you proactive security monitoring capability. IDSs are analogous to an alarm system that you install in your home. The alarm in your home will look for movement within areas of your home that should not have movement when the system is armed, and will usually monitor all access points into your home.

If someone were to access your home without disabling your alarm system first, he or she would have a huge surprise waiting for him or her; an alarm might begin to sound, or a silent alarm might alert the police that your home has been broken into. Whichever the case, with an alarm system you can rest assured that someone will take notice, and investigate the reason your alarm went off in the first place. IDSs work in much the same way, monitoring entries into your network, and looking for changes of data in places in which changes should not be allowed.

Intruders will usually modify files on a system, either as an unintended consequence of their intrusion, or to further compromise the security of the system in the future. This can be accomplished by installing Trojans or backdoors into your system that allow an attacker further access, or by using applications that will sniff the network for passwords that can be used in a later attack.

An IDS can be programmed to understand what patterns of an attack or intrusion look like, and can monitor individual systems for unwarranted changes to the system, such as a computer that has had its data altered mysteriously. Some

IDSs will even build a database of the files on the computer, using a mathematical hash based on a filename, size, modification date, and contents.

When a pattern of attack is recognized, or an alteration is made to a device, the IDS will have a mechanism for alerting a system administrator. Some of these devices will send out real-time information that can be viewed and monitored, and others will attempt to contact individuals through an e-mail or a pager notification. Still others will automatically alter the network topology or device configurations to block the attacker in real time.

For example, the RealSecure IDS offered by Internet Security Systems (ISS) will work in real time with a CheckPoint firewall solution to filter an attacker's IP address immediately, blocking the attack before an administrator has even had time to react to the suspected intrusion.

This kind of advanced capability can significantly improve your response time; however, if it is not configured correctly, it could just as easily detect "wanted" traffic as an attack on your systems, and filter out particular user's functionality. This could obviously be a bad thing, so you should always take care when configuring your IDS. It should be tailored with your system and data flow in mind, and you should always keep its attack database as current as possible.

The capabilities and function of your IDS will depend on your expectation of performance and the IDS device you decide to purchase. You will most likely need to look for devices that work in real time and are capable of handling the type of throughput your network usually delivers. You will also need a system that can give you real-time IDS alerts and information.

You would not, for example, want your home alarm to wait a couple of days to alert you when your home has been broken into. That would defeat the purpose of owning the alarm in the first place. The same is true for your IDS. In order for it to be fully effective, it will need to alert the correct individuals quickly, or take immediate action to block an attack.

It is also extremely important to keep such a system up to date. Since the IDS relies on well-known attack "signatures" and is constantly watching for patterns that match these signatures, it will be necessary to keep the IDS current so that it understands new patterns of attack and signatures. This is necessary, since most attackers will attempt to use up-to-date technologies and techniques to break into your system.

If this is the case and an attacker uses a new method to compromise your system, unless your IDS understands the pattern of attack and signature, it may not alert anyone that your system has been compromised. Instead, it may see this attack as unimportant, once again defeating the purpose of the system.

If you keep your IDS as current as possible, you will minimize these possibilities significantly. This is not to say that you may be the target of some type of brand-new attack that is not recognized by your IDS, but chances are that sometime during that brand-new attack, data will be altered or something malicious will be done to your systems. This modification is usually recognized by your IDS. At that point, you will need to audit your security and logs to fully understand and track any changes the attacker has made.

## Types of Attack

There are numerous methods that an attacker can use against your systems. It is possible for an attacker to steal or alter data, or break into your organization to cause some other type of damage. Not all break-ins result in theft, though. Lately, there has been an enormous amount of other attacks that were once considered to be unimportant since they did not cause any damage. Many of the most common are port scans. These can lead to other attacks down the line, or they just might be unintentional attacks from products such as PcAnywhere.

## Applications Attack

An application attack is a direct attack on a particular application or operating system. The purpose of this type of attack is to render an application useless, or to gain access to a computer. There are many reasons an attacker might want to gain access to a computer in your network. It might be done as a way to “test” one’s skill, or to purposefully cause damage to a system. Unfortunately for most, it is usually the latter.

An attacker will perpetrate such an attack by taking advantage of vulnerabilities within a particular application. Most applications have vulnerabilities that can be taken advantage of in some way or another, and will usually allow an attacker to compromise a system completely. Some of the most popular applications that are easily compromised include:

- DNS/BIND
- HTTP
- SMTP
- FTP

By using programs that can scan systems or networks for open ports, attackers can determine what applications are accessible over the IP network. Armed with

this information, an attacker can select systems that are easily compromised and begin to inflict serious harm.

The attacker might decide to compromise a host and use it as an “attack platform.” In this way, the intruder can attack other systems using the compromised host. The host could even be used to perform a DoS attack.

If an attacker gains control of a host on your system, he or she is free to do anything he or she likes to the computer. This means the attacker could wipe out all of the data contained on the host, or just alter it a little. The attacker could delete all the user accounts, or just sniff the passwords as users log in.

Many times, attackers will want not want to do anything this destructive, so as to go unnoticed. In these instances, the attacker will usually embed hidden code on the compromised host that will perform a particular function.

## Denial-of-Service

Denial-of-service (DoS) is a type of attack that deprives a user or an entire organization of their services and resources. A DoS attack is essentially a flood of erroneous data that forces a device, or devices, to process so much data that the system will be unable to respond to “real” client requests. In effect, an attacker using DoS tactics is capable of shutting off network resources. In actuality, the network devices are still working, but are overloaded to the point where they are no longer functioning correctly.

The loss of service might range from a single host that has been affected, such as an e-mail server, to a complete loss of all network resources. The severity of the attack will be based on the device or devices that are being attacked, how the network is designed, and the amount of DoS traffic. For instance, if your clients access servers and resources that lie behind your firewall, and it is subject to a DoS attack, it will be difficult, if not impossible, to access any resources behind the firewall for the duration of the attack.

What’s worse is that even if you attempt to filter out the data that is causing the attack, it may not stop the DoS. This will be the case if the device still needs to process the packets in order to deny them access into the network. The act of processing the data will actually cause the loss of functionality.

A DoS rarely results in theft or loss of data; instead, it is done to temporarily cease services for a given target. Although there is usually not an “intruder,” a DoS can cause a company to lose valuable time and money while combating the problem, and little work can be accomplished during the attack. Obviously, your customers will not like this lack of functionality, and may demand refunds, even though it was “technically” not your fault.



Several methods can be used to perform a DoS attack; however, any packet that is allowed to reach your network or devices could be used to execute a DoS style of attack. Since most attackers lack the skill required to design and implement new methods of attack, they usually lean toward existing tools and methodology to perpetrate a DoS.

There are several “popular” methods of performing a DoS attack on a system, such as:

- Buffer overflow attack
- Synchronization (SYN) attack
- IP fragmentation attack
- Smurf attack
- Fraggle attack
- Infrastructure attacks

## Buffer Overflow Attacks

Buffer overflow attacks are one of the most common types of DoS attack. A buffer overflow attack targets individual software applications, or operating systems, and sends more data than the applications buffer can allow. When this occurs, programs tend to respond adversely, and in most cases, the application will stop functioning correctly. In these instances, the system may have to be rebooted before the application will resume its normal operation.

One of the most popular buffer overflow attacks has been to send e-mail messages with an attachment that has 256 or more characters in the filename. When these messages were delivered to Netscape or Microsoft Mail applications, they would crash the application running on the server.

Probably the most popular buffer overflow attack is to send oversized Internet Control Message Protocol (ICMP) packets to a device. This is known commonly as “The Ping of Death.” When the device or server receives these packets, it overloads the device, and in many cases, it will cause the operating system to crash.

Other buffer overflow attacks have been designed to combat particular applications. For example, there is a program available called WinNuke. This program allows the person executing the application to enter the target IP address of a computer that is running a Microsoft Windows operating system. The program then exploits a large bug in the Windows operating system and causes unexpected errors. The most common of these errors is a complete system crash. The

system can be easily power-cycled to resume normal operation, but there is a chance that some data loss may occur as a result of the crash.

Unfortunately, nearly every application will have a software bug or component that can be exploited with a buffer overflow technique. Your best bet to combat this problem is to keep your software, and especially your operating systems, up to date. Chances are that any exploits that were discovered in the last release will be fixed in the newer release.

You should also install the necessary bug fixes and security upgrades as they become available. Hackers are usually very versatile, though, and may design a way around even a recently released software application.

## SYN Attacks

SYN attacks take advantage of the structure of the TCP protocol. Since TCP is a connection-oriented protocol, it needs a method of initiating, acknowledging, and ending a session. When a session is initiated, the SYN field is used within the data packet to identify the sequence of the message exchange. This request is received by the target device, and is stored in its buffer to facilitate further connection setup.

The receiving device then acknowledges receipt of the SYN, and awaits further packets. However, if the initiating device fails to respond, the original packet remains in the buffer until it expires, which is usually about 45 seconds.

The problem with all this stems from the fact that the buffer used is usually very small, and can fill up quite easily. When this occurs, other packets will be dropped until there is more room in the buffer. In effect, your device might end up processing only bogus SYN requests, and legitimate connections will be dropped for the entire duration of the attack.

Since the attacker only transmits packets and does not need to receive acknowledgments from the target system, it is common to see SYN floods that use spoofed source addresses. This makes it extremely difficult to trace the actual attacker, and equally as difficult to filter the attack using access-lists. If you did decide to filter the source address of these packets, the attacker could just as easily spoof a different IP address and bypass your security mechanisms.

Moreover, the attack will usually come from several, or possibly hundreds or even thousands of different IP addresses at once. If this is the case, it would probably prove far too difficult to deny all these hosts access.

Instead, you will probably need to adjust your buffer sizes and timeout values. This will allow your device to hold more packets, expire them relatively quickly, or a combination of both. Although this might put an end to small SYN floods, if

an attacker launches an extraordinary attack on your systems, it may be impossible to tweak your buffer size and timeout values enough to resist the attack. If an attacker were capable of creating this much traffic, however, he or she would probably use a different method of attack.

## IP Fragmentation Attack

IP fragmentation attacks, or teardrop attacks, take advantage of the IP protocol and packet size constraints. The IP protocol requires that a packet to be divided into segments if it is too large for the target or next-hop device to handle. When these packets are divided, or fragmented, an offset value is used to identify exactly how to reconstruct the packet once all the fragments have been received.

An attacker can take advantage of this fragmentation by sending fragmented packets that have offset values that are either too small or too large. This causes the target device to reassemble the packet with either too little or too much data. Unless the target device has a method for dealing with this situation, it could cause the system to crash. If this is the case, the system will most likely need to be power-cycled to restore functionality.

Many firewalls have mechanisms that will recognize these types of attacks, and deny the packets access into your network. However, whether you use a firewall or not, it is probably also a good idea to keep your operating systems and device drivers current, and install the latest security fixes and patches on them.

## Smurf Attack

Smurf attacks are among the most popular DoS methods. An attacker will send an ICMP echo, or packet inter-network proper (ping) request to a network broadcast address. When the router “owning” that broadcast address receives the ping, it is usually configured to forward the ping requests to all of the nodes within the network block. This means that all of the network hosts will respond to the ping request, sending hundreds of replies to the source IP address of the ping request.

If an attacker changes the source IP address of the initial ICMP request to the IP address of a target device, the hundreds of replies will be directed at the target device instead of the attacker. Furthermore, since the attacker is capable of amplifying a single ping by a factor of many hundreds, it is possible to send these ICMP requests to numerous broadcast addresses and amplify the number of pings significantly.

If we look at this scenario carefully, it is easy to see how this attack can produce a significant amount of data that will affect a target system. Assume for a moment

that the attacker is capable of sending a stream of ICMP requests totaling 768 kilobits per second (kbps). Let's also assume that these ICMP requests are forwarded to 200 host devices that all send a response to the target system at the same time.

This means that all of the replies combined will generate 150 megabits per second (Mbps) of traffic that will be directed at the victim. If we assume that the parties involved each have a 100 Mbps connection to the Internet, it is easy to see that this attack would cause a DoS for both the target system and the intermediary system. The attacker would remain safely hidden, and would have instigated this mess with a measly 768 kbps of traffic.

To stop your system from being the intermediary system in a Smurf DoS attack, it will be necessary to turn off *directed broadcast* functionality in your routers. If this feature is turned off, the router will not forward an ICMP request to the network nodes when the destination address is the network broadcast address. Essentially, the ping would have little to no effect on your system, preventing the attacker from using your network as a source of amplification.

If you are the victim of a Smurf attack, you will need to filter this traffic before it enters your LAN. If possible, you may choose to add an access-list entry in your border router that denies all ICMP packets coming from the intermediary network. In some routers, this may not prove helpful, though; the problem stems from how the router denies ICMP packets.

Most routers will need to respond to the denied ICMP packet by sending a "destination unreachable" message to the source of the ICMP packet. If this is the case, the router will need to actually process the packet, and send a response. If the number of packets per second (pps) is very high, as it probably will be, this could cause severe performance problems within the router, and in some cases cause the device to crash altogether.

Newer routers and software, such as Cisco routers using Internetwork Operating System (IOS) version 11.1 and higher, are capable of foregoing the "destination unreachable" message, and instead dropping the packet on the interrupt level. This means that the router does not need to process the packet, and simply drops it, leaving its resources available to process legitimate packets.

Another method of denying Smurf attacks is to use a Committed Access Rate (CAR) to limit the amount of packets are allowed to pass through a given network interface. This is easily accomplished on a Cisco series router with IOS version 11.1 or higher. This is discussed in further detail in the "Attack Prevention" section later in the chapter.

Adding access-lists and CARs to your border router may free up network resources; however, your network bandwidth is probably going to be consumed by the amount of traffic flowing in through your connections. Even if you filter this traffic at your border routers, you may not have enough bandwidth to allow legitimate traffic in or out of your access points.

This could especially be the case if you are dealing with a well-organized attack. In this instance, you may need to filter the traffic even further upstream. This situation might even require you to call some of your friends in the ISP business to help you filter the traffic at the network access point (NAP) level.

## Fraggle Attack

We already covered Smurf attacks, so we need to cover the other attack that involves little creatures of television fame: the Fraggle. A Fraggle attack is simply a Smurf attack that has been rewritten to use UDP echo requests instead of ICMP. An attacker will spoof the source IP address of the UDP echo request in much the same way as a Smurf attack, and the effect is the same.

In fact, the methods used to combat this type of attack are identical to those used to counteract a Smurf attack, except, once again, the packets are based on the UDP protocol.

## Physical Attacks

Physical attacks refer to a method of denying service by “physically attacking” part of your infrastructure. For instance, if an attacker has physical access to your servers and network devices, he or she might be able to turn one of the devices off, or unplug a necessary fiber-optic cable.

This might be accidental or malicious behavior. If the person really wanted to do damage, he or she could physically harm some of the devices, by say, pouring coffee in them and short-circuiting the hardware. These types of attacks have the possibility of being the most severe.


Obviously, this type of DoS can only be done by someone who has access to your equipment. At a minimum, you should be able to secure your equipment by locking it up, and only giving certain individuals the key. You might, however, be a very large company, and this type of security concern will need to be handled by live security guards and advanced closed-caption monitoring technology.

Whatever the case, make sure that you understand the implications involved when numerous individuals possess physical access to your equipment, and remember to keep your eyes open for internal sabotage.

## Distributed Denial of Service

Distributed denial of service (DDoS) is one of the newest and most troubling types of attack an ASP must face. This type of attack is perpetrated to cause the same undesired effects offered by DoS attacks, but on an even larger scale. Mainly, it is implemented to cease service and resource offerings of a particular host or network. However, it is accomplished using a “distributed” method. Refer to Figure 6.4.

**Figure 6.4** DDoS Architecture



A DDoS attack uses numerous computers to launch a coordinated DoS attack against one or more targets. Using client/server architectures, sometimes referred to as master/slave, an attacker is able to multiply the effectiveness of a given attack using multiple systems simultaneously.

In a DDoS attack, an attacker must first use a client system to perpetrate an attack on a victim. The attacker will compromise a host, and install a special application called a *handler* on this system. This allows the host to control many other systems. Other hosts are then compromised, and the *agent* software is installed on them.

The agent can then be used to generate a large stream of packets that can be directed at an intended victim. It is a fairly simple concept, one where the attacker controls the client, which in turn controls *handlers* or *masters* that control the *agents* or *slaves*.

Attackers will generally compromise several hosts to be used as handlers, and hundreds or possibly thousands of other hosts to be used as agents. This is accomplished by scanning a large number of networks for hosts that have a particular vulnerability that will allow the attacker to compromise the systems. The attacker will then compromise each host individually and install the required DDoS tool on it.

These hosts can then be used to scan more networks, and compromise other systems. Finally, the attacker will have built a vast arsenal, and may choose to flood a victim with so many packets of data that the victim will almost certainly suffer a very serious DDoS attack.

There are a several applications that an attacker can use to help automate the entire process. There are, in fact, scripts available that will compromise a system and install an active DDoS daemon in a matter of seconds. You could imagine how this allows an attacker to easily generate an arsenal consisting of thousands of agents very quickly.

Once this arsenal has been created, the attacker can use an application to coordinate the attack using his or her entire army simultaneously. Some of the most popular of these DDoS toolkits include:

- Trinoo
- Tribal Flood Network
- Tribal Flood Networks 2000
- Stacheldraht

## Trinoo

Trinoo is a set of tools that allows an attacker to use randomized UDP ports to flood a network with packets of data. This makes it very difficult to filter, unless you do not need UDP packets in your network—this is rarely the case, however. Trinoo has caused many DDoS incidents in the past, and in severe instances, networks that were targeted by attackers using Trinoo took days to fix.

By installing the Trinoo tools on master and slave hosts, an attacker is capable of remotely controlling the master systems using TCP port 27665. In turn, the master hosts communicate commands to the slave hosts using UDP port 27444, and the slave hosts communicate with the masters using UDP port 31335.

It is then possible to filter these TCP and UDP communication ports out of your network, and prevent your systems from becoming masters or slaves. However, this is not entirely true. Since these ports are only the default values, it is possible for an attacker to change the communication ports that Trinoo uses and thwart your security.

Since many attackers may not bother to change these values, it might still be useful to filter these ports out of your network. This will at least prevent the “common thug” from controlling your systems.

## Tribal Flood Network

Tribal Flood Network (TFN) is another DDoS tool that is implemented using “master and slave” architecture. TFN has been used to launch distributed attacks by flooding a target system with ICMP, SYN, or UDP packets, and it is even capable of delivering distributed Smurf attacks. In addition to these capabilities, it is possible to gain immediate root, or system administrator, access to a system that has TFN installed.

Obviously, you will not want your systems to be used to attack a target, so it might be helpful to filter TFN communication from your network. The problem with this is that TFN uses ICMP\_ECHOREPLY packets for two-way communication with its arsenal of compromised masters and slaves. ICMP\_ECHOREPLY packets are packets that are sent when a device is responding to a ping, or ICMP\_ECHO request.

This makes it difficult to discover and block TFN communication with a firewall or packet-filtering device. In many instances, the only way to successfully filter TFN communication in your network is to filter out all ICMP packets. If this is done, you will not be able to ping in to or out of your network, making it difficult to troubleshoot problems. This might present other problems as well, if other processes or devices use ICMP packets for communication or monitoring.



Since TFN is capable of flooding a target system using numerous styles of DoS attacks, it can be very difficult to battle. For instance, if you were flooded with UDP packets using random communications ports, you might filter UDP packets out of your network altogether; however, the attacker could easily switch to another style of attack such as using SYN packets, and resume the attack.

Your best bet is to attempt to block all styles of attack using a unified security system that includes the use of firewalls, border-router packet filtering, CARs, and IDSs. Still, the variety offered by TFN, coupled with the speed that new types of attacks are developed, makes for an explosion that is waiting to happen.

## Tribal Flood Networks 2000

Tribal Flood Networks 2000 (TFN2K) expands on TFN to make a DDoS tool that is almost undetectable, and unstoppable. TFN2K was designed to offer even more flexibility and control over the hosts within its arsenal, and its DDoS style of attack. In fact, TFN2K is even capable of flooding a system with corrupt packets of data that will often times cause a system to crash.

TFN2K is also more difficult to detect when compared to other DDoS tools such as TFN. For one, TFN2K can be configured to communicate with its arsenal of compromised hosts using TCP, UDP, or ICMP packets. In addition, the communications ports can be altered, making it impossible to filter the communication traffic.

Your best line of defense against TFN2K is to not allow the attacker to compromise your hosts in the first place. To accomplish this, you will need to harden your systems and network, and set up egress and ingress packet-filtering techniques that block spoofed IP addresses.

If one of your systems is compromised, and has been turned into a master or slave device, it will be nearly impossible to detect the intrusion until it is too late. By that time, your systems may be used to launch a DDoS attack on someone else.

## Stacheldraht

*Stacheldraht* is the German word for barbed wire. It is another DDoS tool based on the master/slave model and has been used to orchestrate thousands of host systems in distributed attacks. Stacheldraht, much like TFN, is capable of attacking a target using many different methods of attack, and is capable of spoofing IP addresses.

An attacker using Stacheldraht can flood a target using UDP, TCP (SYN), or ICMP packets, as well as Smurf techniques. This makes it just as difficult to stop a Stacheldraht attack as it is to stop a TFN attack.

Stacheldraht uses TCP port 16660 and ICMP\_ECHOREPLY packets for communication between the master and the slave devices. Although it is possible to filter these protocols and ports out of your network, the attacker can easily change the ports used to communicate.

An attacker is able to remotely control a master device using a client that connects to these devices using TCP port 16660. This TCP port can be filtered out of your network easily and should; however, the attacker may just decide to use a different value.

Stacheldraht also offers encryption between the attacker and the master servers, making it more difficult to discover by firewalls and IDSs. The Stacheldraht slaves are even capable of updating their DDoS daemon software automatically, and are therefore typically using the most current version.

The methods for stopping a Stacheldraht attack are about the same as with a TFN attack. First, make certain to harden your systems and network to prevent your hosts from being compromised. Second, use advanced firewall and packet-filtering techniques to successfully protect your systems from becoming a victim or a target.

## Prevention Techniques

Hopefully, by now you already understand some of the methods an attacker can use to break into your systems and cause harm. You should also understand how to safeguard against some of these attacks. At a minimum, you should be planning to install a firewall, and use packet-filtering techniques to guard against some types of attack.

However, if an attacker is clever and truly wants access into your system, he or she will find a way to be successful. Although it will be impossible to stop every type of attack, it is still important to employ some tactics that will improve your chances for survival.

If you have secured your system successfully, an attacker might even grow tired of attempting to break into your systems, and turn his or her attention to somebody else. Although this may not allow you to make friends with the attacker, it will thwart his or her invasions the majority of the time.

For these reasons, you should look into securing your systems as much as possible, and use only technologies that provide true protection for your network and systems.

Prevention is always the best policy. I am certain that you would much rather prefer that attackers not penetrate your defenses, rather than chase them around

and clean up after the mess they leave behind—unless you enjoy playing “cops and robbers” in real life.

The following steps will help you secure your systems and block many types of attack:

1. Filter all RFC 1918 address spaces.
2. Apply ingress and egress filtering.
3. Apply rate limiting.
4. Use TCP Intercept to prevent SYN floods.

## Filtering RFC1918 Address Spaces

As IP networking and the Internet began to come into widespread use, it became obvious that some companies used IP addresses for systems that were never intended to connect to the Internet. This meant that many of the dwindling IP addresses were wasted on private companies that used the addresses only to route internal traffic.

To counteract this problem, the Network Working Group created RFC 1918. This document outlined Internet Best Practices, mainly the use of certain IP address ranges within private networks.

In addition, the Internet Assigned Numbers Authority (IANA) reserved the following addresses for private use:

- 10.0.0.0–10.255.255.255 (10.0.0.0/8)
- 172.16.0.0–172.31.255.255 (172.16.0.0 /12)
- 192.168.0.0–192.168.255.255 (192.168.0.0 /16)

In effect, organizations that wanted to use IP networking for their own internal purposes could now use an addressing scheme based on one of these three reserved address spaces. Since the nodes were never intended to access the Internet, it did not matter if the addressing scheme used by one organization overlapped with another, as long as there was not an overlap in any of the connected networks. This meant that these reserved addresses could never be used on the Internet, since they were free to be used by all.

Attackers quickly took notice, and began to use these reserved address spaces to spoof IP packets. An attacker may choose to spoof the IP packets for any number of reasons, many of which we have already discussed. Regardless of the reason, he or she will typically use an IP address from the reserved address space.

This means that you can stop the majority of spoofing attacks by filtering these reserved address spaces at your Internet border router. This can usually be accomplished very easily by installing filters on your border router to deny all packets entering your external interface with a source address that is in the range of the reserved addresses described in RFC 1918.

Although the method may vary, all routers will allow you to configure packet filters in some way or another. Figure 6.5 shows the configuration of a Cisco router that denies spoofed packets and permits all other packets.

**Figure 6.5** Cisco 7500 Series Router Configuration with Anti-Spoofing Access-Lists

---

```
!  
hostname Border-Router  
!  
interface FastEthernet1/0/0  
 ip address 199.199.199.200 255.255.255.0  
 ip access-group 150 in  
!  
interface FastEthernet2/0/0  
 ip address 200.200.200.200 255.255.255.0  
!  
access-list 150 deny ip 10.0.0.0 0.255.255.255 any  
access-list 150 deny ip 192.168.0.0 0.0.255.255 any  
access-list 150 deny ip 172.16.0.0 0.15.255.255 any  
access-list 150 permit ip any any  
!
```

---

In the configuration shown in Figure 6.5, any packets entering the external, or Internet interface will be denied access if their source address is a private address. All other “legitimate” packets will be allowed to enter the internal network, and all packets regardless of their source address will be permitted to leave the internal network. This will successfully stop RFC 1918 addresses from entering your network. In addition, you might want to apply the same access-list in reverse, as shown in Figure 6.6.

**Figure 6.6** Cisco 7500 Series Border Router Configuration

---

```

!
hostname Border-Router
!
interface FastEthernet1/0/0
 ip address 200.200.200.200 255.255.255.0
 ip access-group 150 in
 ip access-group 150 out
!
interface FastEthernet2/0/0
 ip address 201.201.201.201 255.255.255.0
!
access-list 150 deny ip 10.0.0.0 0.255.255.255 any
access-list 150 deny ip 192.168.0.0 0.0.255.255 any
access-list 150 deny ip 172.16.0.0 0.15.255.255 any
access-list 150 permit ip any any
!

```


---

In this example, RFC 1918 addresses are blocked from leaving and entering your network. This will ensure that your internal systems will be unable to send spoofed source packets to other systems on the Internet. This may in fact help you to prevent your systems from being used as an intermediary system in a DoS attack.

## Ingress and Egress Filtering

Consider Figure 6.7 for a moment. In this example, the only traffic that should pass through the ISP edge router is packets with a source address that is assigned to the customer's network. Conversely, the customer network should only accept traffic with source addresses other than its own network block. This means that any other packets with source addresses that do not match these rules should not be allowed into the corresponding network. These packets would, by nature be the result of a spoofing attack.

Think about it for a moment. If a packet with a source address other than the customer's network, and any other networks behind it, were to enter the ISP edge router, where would this packet have come from? In our example, the customer network is comprised of a single network block, and it has not been subnetted or partitioned in any way.

**Figure 6.7** ISP Core-to-Customer Network

This means that the only “legal” addresses are those that the customer owns or has been assigned to use. Therefore, the only packets we should allow to leave the customer’s border router and enter the ISP edge router are those with a source IP address that matches the customer’s assigned range. All other packets should be filtered out.

The same is true in reverse. If a packet with a source address that matches the customer’s assigned IP scheme enters the customer’s router via the ISP edge router, how is this packet legal? It isn’t, because the only organization allowed to use this IP address *is* the customer. In this instance, the customer’s own network block should not be allowed to exit the ISP edge router on its customer-facing interface, nor should it be allowed to enter the customer’s border router. Instead, the address range should be filtered and denied access.

This type of filtering is referred to as *ingress* and *egress* filtering. Ingress filtering is used when these packets are filtered as they enter a network interface, and egress filtering is used when we filter these packets as they leave the interface. This is an extremely effective method of stopping spoof attacks. In fact, if ingress and egress filtering techniques were applied to every edge and border router across the Internet, it would be nearly impossible for anyone to spoof packets. Imagine that!

Ingress and egress filtering can be done on just about any router in existence. Figure 6.8 depicts the ISP edge router configuration based on our example in Figure 6.7, and Figure 6.9 depicts the customer border router from the same example.

**Figure 6.8** ISP Edge Router (Cisco 7500 Series) Ingress and Egress Filtering

```
!
hostname Edge-Router
!
interface GigaBitEthernet1/0/0
```

Continued

**Figure 6.8** Continued

---

```

ip address 198.198.198.198 255.255.255.0
!
interface FastEthernet2/0/0
ip address 199.199.199.199 255.255.255.0
ip access-group 110 in
!
access-list 110 permit ip 200.200.200.0 0.255.255.255 any
access-list 110 deny ip any any
!

```

---

**Figure 6.9** Customer Border Router (Cisco 3600 Series) Ingress and Egress Filtering

---

```

!
hostname Border-Router
!
interface FastEthernet1/0
ip address 199.199.199.200 255.255.255.0
!
interface FastEthernet2/0
ip address 200.200.200.200 255.255.255.0
ip access-group 110 in
ip access-group 120 out
!
access-list 110 deny ip 200.200.200.0 0.255.255.255 any
access-list 110 permit ip any any
access-list 120 permit ip 200.200.200.0 0.255.255.255 any
access-list 120 deny ip any any
!

```

---

## Rate Limiting

Most routers can be configured to limit the amount of data that will be processed for a particular time interval. This is known as *rate limiting*. It is usually configured

by identifying the type of traffic that should be affected, and setting a maximum rate, usually in kbps, that a router will process.

Rate limiting can be an extremely useful tool for preventing flooding attacks that are usually the result of DoS or DDoS tactics. In these types of attacks, an attacker will attempt to take down an Internet connection by flooding a network with erroneous traffic. By limiting the amount of data that your router can process, it is sometimes possible to overcome even large flooding attempts.

Cisco Systems routers are capable of limiting the input or output rate of a particular interface using Committed Access Rate (CAR). This is accomplished using the *rate-limit* command on the interface that requires limiting. When configuring this, you will first need to create an access-list that defines the exact type of traffic, including the source and destination addresses, that you would like to rate-limit. Once this is done, you will need to decide which interface you need to filter traffic on, and use the *rate-limit* command to activate it on the interface.

When using the *rate-limit* command, it will be necessary to configure the following criteria:

- **Packet direction** This is the actual direction that data should be filtered. For example, you might want to filter data entering your system (input), or data exiting your system (output).
- **Average rate** This is the “normal” amount of traffic that should not be affected by rate limiting. When your total predetermined traffic is measured and falls below the *average rate*, this traffic is considered to “conform” and will be transmitted if so configured. However, if the predetermined traffic exceeds the *average rate*, it is considered to “exceed” the rate limitations, and can be dropped depending on the configuration.

You should take great care in setting the *average rate* value, and use a long-term average to arrive at a suitable figure. If you set the value too low, the router will rate-limit normal traffic, and if you set the value too high, the router may allow too much data to enter your system. Because of this, you should probably monitor the type of traffic you wish to rate limit, and use a long-term average to set the value.

- **Normal burst size** How large a particular burst of traffic can be before the sender is considered to exceed the allocated rate.
- **Excess burst size** How large traffic bursts can be before all traffic exceeds the rate limit.



## Flood Attacks

As you have learned, many methods of DoS attack will flood your network using ICMP packets. Although it is possible to filter these packets out of your network completely, ICMP is frequently used by legitimate Internet applications. By filtering ICMP packets completely out of your network, you may remove the functionality of some of the applications on which you rely.

Instead, you might want to consider limiting the amount of ICMP packets allowed into your network. In this way, you will still prevent many types of DoS attack, while regaining the functionality of some of your Internet applications. Figure 6.10 shows an example of how to limit the rate of ICMP packets entering a Cisco router.

**Figure 6.10** Cisco 3600 Series Router with ICMP Rate Filter

---

```

!
hostname Border-Router
!
interface Serial0/0
 ip address 199.199.199.200 255.255.255.0
 encapsulation ppp
 rate-limit input access-group 175 256000 8000 8000 conform-action
   transmit exceed-action drop
!
interface FastEthernet1/0
 ip address 200.200.200.200 255.255.255.0
!
access-list 175 permit icmp any any echo
access-list 175 permit icmp any any echo-reply
!

```

---

In Figure 6.10, the *rate-limit* command has been applied to traffic flowing into Serial0/0, which might be a T1 operating at 1.44 Mbps, or a T3 operating at 45 Mbps. The type of traffic to be filtered—in this case, ICMP—has been defined by *access-list 175*. The *average rate* has been set to 256 kbps, and both the *normal* and *excess burst* sizes have been set to 8 kbps.

This should stop ICMP flooding attempts, while permitting enough ICMP traffic into the system for normal operation. Your system will most certainly be

different, so it is important to get a good idea of your “normal” usage patterns prior to configuring rate limiting.

Although the preceding example only filtered ICMP packets, the same tactic can be easily converted to limit TCP or UDP packets. In fact, since an access-list is used to define the type of traffic that will be rate-limited, it is possible to specify the communications ports of a given protocol. Rate limiting could even be used to limit the flow of all traffic entering or exiting an interface, regardless of the protocol used.

## SYN Attacks

It is possible to prevent most SYN attacks on your system using CARs to limit the amount of TCP traffic bursting allowed on your system. To accomplish this, you will need to configure rate limiting to allow for the full bandwidth of your connection, but reduce your *normal* and *excess bursting* sizes. Figure 6.11 shows a Cisco router configured to prevent HTTP (port 80) SYN attacks.

**Figure 6.11** Cisco 7500 Series Router Using CAR to Prevent SYN Attacks

---

```
!  
hostname Border-Router  
ip cef  
!  
interface FastEthernet0/0/0  
    cef distributed  
    ip address 199.199.199.200 255.255.255.0  
    rate-limit output access-group 175 2500000 250000 250000 conform-  
        action transmit exceed-action drop  
!  
interface FastEthernet1/0/0  
    ip address 200.200.200.200 255.255.255.0  
!  
access-list 176 permit tcp any host eq www established
```

---

**NOTE**


---

Before rate-limiting SYNs, it is very important to find the amount of “normal” or legitimate SYNs in your network. If the *rate-limit* values are set too low, you may block legitimate traffic. Once CAR has been configured correctly, you should be able to use the *show interfaces rate-limit* command to display the *conformed* and *exceeded* rates for a given interface.

---

If the *exceeded* rates are incrementing, and you are not under an attack, you will need to alter the *rate-limit* values to correct the problem. It will also be a good idea to keep an eye on these statistics, since the amount of legitimate SYNs may grow, and they will help alert you when you *are* under attack.

## TCP Intercept

When an attacker launches a SYN attack on your organization, the flood is usually directed at your servers. Using SYN packets that have unreachable source addresses, the attacker is able to overwhelm your servers and force them to deny legitimate packets. In order to combat this problem, you will need to take steps to reduce the amount of illegitimate SYN packets in your network, thus allowing the legitimate packets to be processed by your systems. This can be accomplished on a Cisco router using the *ip tcp intercept* command.

TCP provides a method by which a router can “intercept” incoming TCP SYN packets and verify that the source address is reachable. If the packet appears to be legitimate, it is allowed to enter your network. If it is not, however, it won’t be allowed to reach the destination server, and will be dropped by the router. In Figure 6.12, a Cisco router is configured using TCP Intercept.

**Figure 6.12** Cisco 7500 Series Router Using TCP Intercept to Filter SYN Packets

---

```
!
hostname Border-Router
!
interface FastEthernet0/0/0
 ip address 199.199.199.200 255.255.255.0
!
interface FastEthernet1/0/0
```

---

Continued

**Figure 6.12** Continued

---

```
ip address 200.200.200.200 255.255.255.0
!  
ip tcp intercept list 125  
!  
access-list 125 permit tcp any 200.200.200.0 0.0.0.255
```

---

In Figure 6.12, TCP Intercept was enabled using only the default values. Although this might work well in most environments, you may need to adjust some of the characteristics of TCP Intercept to suit your organization. The following is a list of settings that can be used to alter TCP Intercept:

- TCP Intercept mode
- TCP Intercept timers
- Drop Mode
- Aggressive Mode Thresholds

## TCP Intercept Mode

TCP Intercept can be configured in two modes, the first and default mode is known as *active intercept*. Using this method, the router intercepts all incoming SYNs and responds to the request by sending an acknowledgment (ACK) and request (SYN) back to the originating device.

If the router receives an acknowledgment back from the originating device, the communication is considered legitimate. The router then sets up the connection with the destination server, and allows for communication between the devices. If, on the other hand, the router never receives an acknowledgment from the originating server, no packets will be sent to the destination server, and the router would have successfully filtered an illegitimate connection request.

TCP Intercept can also be configured in *watch* mode. In this mode, the router is passive, and allows the connection requests to be sent directly to the server. The destination server will then respond to the request and begin to set up the connection. The router will watch this communication, and reset the connection if the originating device has not responded back in a certain amount of time that is defined by the intercept timer. This will allow the server to release the connection, freeing up resources for other legitimate connections.

## TCP Intercept Timers

The default time that a router will wait for an acknowledgment back from an originating device before resetting the connection is 30 seconds. However, you may need to adjust this setting if you have legitimate connections that are being dropped.

This can be accomplished using the `ip tcp intercept timeout` command. For example, to set the timeout value to 60 seconds, you would issue the following command in global configuration mode:

```
Router(config)# ip tcp intercept watch-timeout 60
```

## Drop Mode

When your system is flooded with SYN packets, TCP Intercept will move into *aggressive* mode. This is triggered when the total number of incomplete connections exceeds the *max-incomplete* threshold, or the number of connections received within a 60-second time interval exceeds the *one-minute* threshold.

When one of these thresholds is surpassed, the router will begin dropping packets in order to keep the amount of traffic allowed into the network below the configured maximum threshold. By default, this will be accomplished by dropping the oldest connection to allow for the new connection. This can, however, be changed to drop random partial connections instead by using the following command in global configuration mode:

```
Router(config)# ip tcp intercept drop-mode random
```

## Aggressive Mode Thresholds

Depending on the needs of your organization, you may need to adjust the TCP Intercept aggressive mode thresholds. If you rely on the default values, there is a chance that legitimate requests may be dropped as a result of slow network connections, or high network utilization. If this is the case, four thresholds can be adjusted to alter the aggressive mode behavior of TCP Intercept:

- **One-Minute Low** This controls the number of connection attempts, over a one-minute period that will trigger aggressive mode behavior. The default is 900 connections.
- **One-Minute High** This limits the number of connection attempts allowed in a one-minute period. When a new connection attempt is made, and it exceeds this threshold, a connection will first need to be

dropped before the new connection is allowed. The default is a maximum of 1100 connections.

- **Max-Incomplete Low** This controls the total number of concurrent connection attempts that will trigger aggressive mode behavior. The default is 900 connections.
- **Max-Incomplete High** This limits the number of concurrent connection attempts allowed. The maximum number of incomplete connections is 1100 by default.

## NOTE

Before configuring TCP Intercept, you should monitor your network and take a long-term measurement of the SYN packets in your network. This data will help you arrive at suitable values to use in your TCP Intercept configuration. Once TCP Intercept has been configured and is operating on your router, you can use the *show tcp intercept connections* command to view the number of incomplete and established connections.

The *show tcp intercept statistics* command will also display the TCP Intercept statistics. The information gathered using these commands should help you alter your TCP Intercept configuration to maximize your resources and minimize the threat of SYN flood attacks.

## Capturing Evidence

If your organization has been the victim of an attack, it will be very important to capture and preserve as much evidence as possible. Any evidence you may be able to gather might prove useful in locating an attacker, and preventing further attack. In addition, you will need to present evidence to a law enforcement agency in order to prosecute an attacker.

If you have firewalls and intrusion detection systems on your network, these will probably have logging functionality built into them. Logging is not always enabled by default, however, so it is very important to check the device and configuration to ensure that logs are being successfully recorded. In addition to this, there are usually several options from which to choose, ranging from minimal to maximum logging methods with many degrees in between.

Maximum logging is usually beneficial, since you will want to preserve every detail possible once an attack has been launched on your system. If it does not cause a performance decrease, and you have enough room available to keep sufficient logs, we would recommend using maximum logging.

Be careful when setting maximum logging, since it will usually generate a significant amount of logs, and most devices are not capable of storing much information. If you have a lot of traffic flowing in and out of your network, and are logging each connection, the device may only hold enough room for an hour's worth of logs. If this is the case, you may need to significantly reduce the level of logging you have configured, so that you can store your logs for a longer period of time.

## Syslog

Syslog is a software daemon that runs on a server to allow for logging of messages and events. Most IP network devices are capable of transmitting logging information to a preconfigured Syslog server. The advantage of this is that you can install a centralized server where all or some of your devices can store logging information.

Since the logs are stored on a server and not the individual devices, the location of the log files is always known, and you can easily add additional storage capacity, in the form of hard drives, to facilitate storage of additional logs. This means you can increase the level of logging without the concern of running out of space on each individual device.

Instead, when your Syslog server begins to fill, you can choose to add additional hard drives, back up the log files to tape to free up space, or remove the data that is no longer needed. Obviously, this gives you a great deal of flexibility for device logging.

Syslog originated on the Unix platform, and is available for almost every operating system. If you own a Unix system, getting Syslog to operate is probably just a matter of configuring and running the daemon. Even if you have a non-Unix server, such as a Windows 2000 device, it is easy to locate a Syslog application that will operate with your system.

Because of the simplicity of the Syslog daemon, and the flexibility it provides, we recommend using a Syslog server in your network, and holding your log files for at least a month. In this way, it will be possible to look for small or recurring attacks that may have otherwise gone unnoticed.

## Packet Capturing

If you discover that you are under attack, it will be important to capture additional information that your network devices or Syslog server may not record. If possible, you should attempt to capture all of an attacker's data packet traffic for sample analysis. You can do this using a commercially available packet capture utility such as Wild Packets' Etherpeek, or Network Associates' Sniffer Suites. In addition to these tools, most Unix systems come equipped with a packet-capturing program.

Linux and SUN operating systems include an application called `tcpdump` that can be used to capture packets in real time. It is very simple to use, and can be initiated by issuing the following command when logged in to a Unix server:

```
Tcpdump -I interface -s MTU -w name-of-capture-file
```

In addition to this, Solaris operating systems include the `snoop` application, which can be initiated using the following command:

```
snoop -d interface -o name-of-capture-file -s MTU
```

In both of these instances, you should make sure there is enough disk space available to capture the packets before initiating the command. In addition, it will not be good enough to issue the command on any server within the network. Instead, you will need to choose a device that is either being attacked directly by the attacker, or has the ability to see all of the attacker's packets as they traverse the network.

For example, if you are using a hub to connect several servers to the network, each server will see the same traffic as its neighboring server, thereby allowing you to sniff from any of the devices. However, if you connect the servers with a switch, each server is partitioned into its own collision domain, and will therefore not see traffic that is targeted at a different server.

To alleviate some of these problems, you may need to capture packets as they enter and exit your network, such as at the router, or configure a switch to send all traffic to a specific port that can be used to "sniff" the packets in your network.



## Summary

For an ASP to survive, it must deliver solid and functional solutions to its customers, and be capable of delivering these with a certain level of security and privacy. In fact, the average ASP customer is becoming increasingly inquisitive of the security measures that the ASP uses to guarantee privacy and the security of their data.

With the availability of hacking tools, the increasing number of attacks, and the large amount of publicity they have garnered in our society, it is no wonder why everyone is so concerned—you should be, too.

Would you purchase a car that had no way of locking the doors, and no way of preventing an intruder from stealing it? I doubt it, just as I doubt a customer will want to purchase a solution from you, knowing that you will not guarantee or provide any method of security with your services. If you cannot guarantee that your systems will be up and running tomorrow at 100 percent of their intended functionality, and that your customers information and data will be held private, why would a customer want to purchase a solution from you? Instead, you must be able to guarantee a certain level of security and privacy in order to sell your products to a customer. Even more importantly, you will need to effectively implement these policies.

In this chapter, you learned why it is important to develop both a security and privacy policy. You have also learned that these must be updated frequently, and that your security mechanisms need to be tested and audited at regular intervals.

We explained various components that you can use to implement your security policy, such as login procedures, digital certificates, cryptography, and security logging.

Delivering customers viable and secure solutions is the most important goal for ASPs. Without it, ASPs will soon find themselves out of business. For ASPs, sharing information securely with their customers and creating solutions that are secure are paramount. Remember that delivering and receiving information is at the center of its existence for ASPs. Anything that threatens this information or the processing of that information will directly imperil the ASP.

Some of the major concerns that face an ASP are confidentiality, accuracy, timelines for the information or the availability of solutions, and threats that have to be countered by security measures. Security management can help an ASP by defining and implementing countermeasures for security risks.

The ASP that is able to monitor and handle these security risks will have a quantifiable competitive advantage.

# Solutions Fast Track

## Security Policy

- ☑ An ASP needs to develop a general security policy that addresses how it manages and maintains the internal security posture of its infrastructure.
- ☑ A security policy defines how an ASP manages, protects, and distributes sensitive information and resources. Any ASP, before connecting to the Internet, should develop a usage policy that clearly identifies the solutions they will be using and exactly how those solutions will be used.
- ☑ An extension of the security policy is the *privacy policy*. The privacy policy should state what data the ASP considers to be confidential, and how that data can and cannot be used.

## Security Components

- ☑ As an ASP, to validate both the security policy and the privacy policy, a review of the various security mechanisms and methods used to implement those policies is required.
- ☑ One of the most important methods to provide accurate security is the ability to authenticate users and systems.
- ☑ A PIN provides another mechanism that you can use to enhance the security of a standard username and password system.
- ☑ Confidentiality is usually associated with data encryption mechanisms such as Secure Socket Layer (SSL) or Data Encryption Standard (DES), and targeted at protecting data as it traverses across a network, such as the Internet.

## Security Technologies and Attacks

- ☑ ASPs must deploy the best security technologies. Strong encryption is important, whether in the context of an SSL browser connection or a VPN connection.
- ☑ The two basic methods of VPN access are LAN-to-LAN VPNs and remote access VPNs.

- ☑ A perimeter firewall is a device, or software application, that controls access in to and out of a given network.
- ☑ Stateful inspection provides for the most robust of all firewall features.
- ☑ Embedded firewalls are software applications that are installed and run on a computer to guard it against attacks.
- ☑ Distributed denial of service (DDoS) is one of the newest and most troubling types of attack an ASP must face. This type of attack is perpetrated to cause the same undesired effects offered by DoS attacks, but on an even larger scale.

## Prevention Techniques

- ☑ As IP networking and the Internet began to come into widespread use, it became obvious that some companies used IP addresses for systems that were never intended to connect to the Internet. This meant that many of the dwindling IP addresses were wasted on private companies that used the addresses only to route internal traffic.
- ☑ Ingress filtering is used when these packets are filtered as they enter a network interface, and egress filtering is used when we filter these packets as they leave the interface.
- ☑ Most routers can be configured to limit the amount of data that will be processed for a particular time interval. This is known as *rate limiting*.
- ☑ It is possible to prevent most SYN attacks on your system using CARs to limit the amount of TCP traffic bursting allowed on your system. To accomplish this, you will need to configure rate limiting to allow for the full bandwidth of your connection, but reduce your *normal* and *excess bursting* sizes.

## Capturing Evidence

- ☑ If your organization has been the victim of an attack, it will be very important to capture and preserve as much evidence as possible. Any evidence you may be able to gather might prove useful in locating an attacker, and preventing further attack.

- ☑ Syslog is a software daemon that runs on a server to allow for logging of messages and events.
- ☑ Linux and SUN operating systems include an application called tcpdump that can be used to capture packets in real time.

## Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the “Ask the Author” form.

**Q:** Is your ASP able to provide redundancy and load-balancing services for firewalls and other security-critical elements?

**A:** Depending on the design and implementation of your network, you may be able to answer “yes” to this question.

**Q:** Can your ASP handle external test attacks on at least a quarterly basis, and internal network security audits at least annually?

**A:** This is so that you can unequivocally state to your clients that your network is at least tested, and can deny at least a certain level of intrusion.

**Q:** Can the ASP obtain audits for customer network security?

**A:** This will help to ensure that other your other customers will not compromise the ASP backbone.

**Q:** Can the ASP provide a documented policy for the operating system that runs the Web and other servers?

**A:** This will help explain your policies and procedures to the customer. This will help when you have clients who need to know exactly what will happen to their data once it gets to your application hosting site.

**Q:** If the ASP runs customer applications on physical servers that are located outside of its network, does it have a documented set of controls that it will use

to ensure that there is separation of data and security information between customer applications?

**A:** This is highly important to gain an edge over other ASP-based companies. If you can effectively point out where your applications are and how they are handled when they get there, you should have the ability to ease the customer as to the security of your services.

**Q:** Does you provide application or transaction-based intrusion detection services?

**A:** This question will explain how you implemented your security policy. If it is by application, that may mean that there is a security check that takes place during the usage of an application. If the policy that you implement is transaction based, this means that every calculation or information change will require a new security check.

**Q:** Does your ASP perform background checks on personnel who will have administrative access to servers and applications?

**A:** This falls under the realm of social engineering, and may be the weakest link in the chain for many companies. If you cannot trust your people, there is truly no way to secure your data.

**Q:** Does your ASP have a documented process for evaluating operating systems and applications, and what is the process for installing security patches and service packs?

**A:** This is very important to many high-security type companies. Many times, these companies are looking for some form of stability and processes, rather than an ad hoc, network-on-the-fly, environment.

**Q:** Does your ASP have the ability to show its documented procedures for intrusion detection, incident response, and incident escalation/investigation?

**A:** This is very important for the tracing and prosecution of network trespassers.

**Q:** Is your ASP a member of the Forum for Incident Response and Security Teams, or uses a security service provider that is?

**A:** This is like a certification such as ISO 9000. What this proves to your client is that you are committed to having a secure network and application infrastructure.

## Management and Monitoring

### Solutions in this chapter:

- The Effect of Outsourcing
- What Service Levels Should the Service Provider Consider?
- The Realities of Customer Compensation
- How Service Providers Have Responded
- The Operation Support System Model
- Broadband Access Changes the Market
- Quality of Service
- Management Systems for Your ASP
- What Tools Do You Need to Automate TMN?
- The ASP Transformation
- Pricing Models and Billing
- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

## Introduction

According to a recent survey by Current Analysis, customers rank support capability, cost and pricing structure, service level agreement (SLA), and other management and monitoring capabilities as the most important decision criteria in selecting an application service provider (ASP). USi, one of today's leading ASPs, further declared that the true full-service ASP, after the initial deployment of its product, should also diligently keep up with maintaining the ongoing performance of applications. This means continuous network and applications management, the tightest security, and 24x7xForever customer support. In other words, an ASP should take total responsibility for the full life cycle of the service offering.

There are two major tasks central to the ongoing management of an ASP. The first service component for application management is that an ASP must have expertise pertinent to the applications it is offering. The ASP will need to respond to customer application problems, meaning that the ASP must go back to source independent software vendors (ISVs) if an application failure requires code modification.

The second service component for application management is more challenging, and involves end-to-end customer care and service guarantee. An ASP is the customer's single point of contact for application performance. The ASP has to be responsible for all failures or problems, including those emanating from any of the underlying service layers that support hosted applications. The best help desk or customer care practice is to issue a single trouble ticket for any problem encountered with a hosted application. An ASP needs to be either in control of the data center and network layers of its service, or have a mechanism established with its service providers to troubleshoot infrastructure-related problems that may affect application performance.

## The Effect of Outsourcing

With the explosion of distributed applications and database systems, customers are paying more attention to the performance of their service provider. When the Internet first gained a foothold in the corporate network, it allowed companies to scale to a wide geographic range. ISVs began offering packages designed to meet the needs of companies that were struggling with the strains of building a highly available, and scalable, infrastructure. In essence, these packaged technologies were able to help customers leverage cost effective, redundant infrastructures that were too cost prohibitive in the past.

As with all changes, there are challenges that one must face. By implementing outsourced application packages, many companies lost their ability to control the performance and reliability of their networks. As you probably can attest to, this leads to unhappy clients for you and your customer. As time progressed, this became a very substantial issue, but how do you outsource and still maintain control?

## Service Level Agreements

Carrier services these days are embedded with management capabilities that enable clients to receive an acceptable set of metrics that you as a service provider must maintain. So, what is the glorious document that will help change the business? The service level agreement (SLA). SLAs allow the customer to set minimum (and maximum) limits to be met, or there will be consequences and serious repercussions. There are three main areas in almost every SLA:

- **Planning** Determining the wide area network (WAN) service levels.
- **Verification** Monitoring the service levels to guarantee fulfillment.
- **Troubleshooting** Isolating issues when service levels are not delivered.

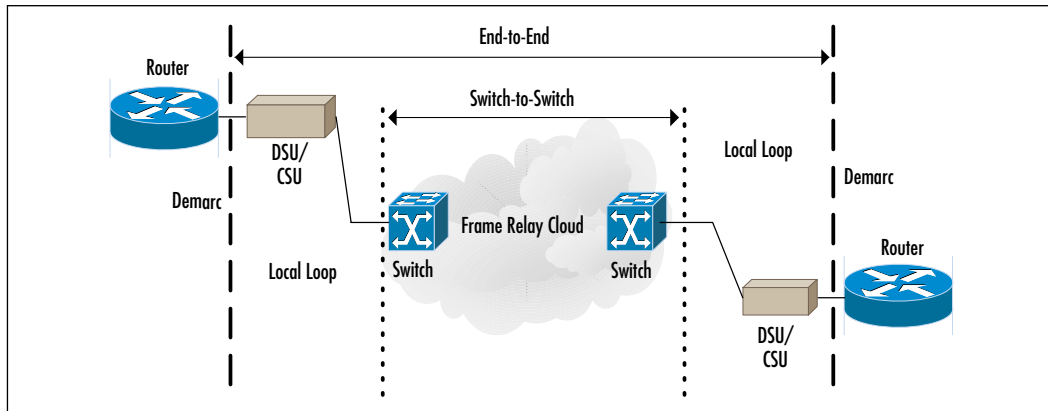
## Some Common SLA Guarantees

What are the common guarantees given to ASP customers these days? What is setting these service providers apart from their competition? I think that it comes as no surprise that most service providers offer:

- High availability and system uptime
- Bandwidth (and more bandwidth)
- Latency assurances

There are some key pieces of information that will have a direct impact on these SLA issues. One of these issues is where the measurements are taken. Do you take these measurements from end to end (from the customer premise equipment (CPE)), or from within the Frame Relay cloud (from switch to switch). The reason that this has a large impact on the SLA is due to problems that can arise in the “last mile” (or local loop). In a switch-to-switch deployment, the last mile is not taken into account; therefore, many customers find it more meaningful to measure from end to end. Figure 7.1 shows a simple end-to-end topology.



**Figure 7.1** Simple End-to-End Topology

There is another key area of concern in finding a measurement system that is independent from the network that is being sampled. A switch (or router) within the network cannot provide all of the vital statistics that will give meaningful WAN service level data. Implementing a device that is not biased toward router or switch architecture is the only way to receive valid network statistics in the end-to-end model.

You have to remember that the presentation of the data is almost as important as the data itself. Reporting methods that are clear and concise are necessary to give your customers the performance guarantees that they are anticipating. This statistical data is the only way that you can truly validate the value that is added by your service.

## What Are the Basic Components of SLAs for Frame Relay Circuits?

Frame Relay involves a number of system parameters that go beyond the standard parameters that can be monitored by the Simple Network Management Protocol (SNMP). Some of these elements cover the entire network, segmented networks, or even single circuits. The level at which an SLA can be defined depends entirely on the business need of the circuit. For example, SLAs that cover individual devices or components usually allow for less downtime than those that cover the entire infrastructure do.

SLA components are generally implemented inconsistently from company to company, even though there are fairly standard ways to calculate reliability. When

you are trying to determine SLAs, you must understand the details implied for each of these measurements:

- **Network availability** This is generally measured for one month and is comprised of the following equation:

$$\frac{(\text{hours in a day}) * (\text{number of days in month}) * (\text{number of locations}) - (\text{network down time})}{(\text{hours in a day}) * (\text{number of days in month}) * (\text{number of locations})}$$

- **PVC availability** This is generally measured for one month and is comprised of the following equation:

$$\frac{(\text{hours in a day}) * (\text{number of days in month}) * (\text{number of PVCs}) - (\text{PVC downtime})}{(\text{hours in a day}) * (\text{number of days in month}) * (\text{number of PVCs})}$$

- **Average network delay (round-trip)** This is generally measured for one month and is comprised of the following equation:

$$\frac{(\text{cumulative sum of samples taken end-to-end})}{(\text{number of samples taken})}$$

- **Average PVC delay (round-trip)** This is generally measured for one month and is comprised of the following equation:

$$\frac{(\text{cumulative sum of samples taken of PVC delay})}{(\text{number of samples taken})}$$

- **Effective throughput (PVC)** This is generally measured for one month and is comprised of the following equation:

$$\frac{(\text{egress frame count})}{(\text{ingress frame count}) - (\text{number of frames above committed burst size}) - (\text{excess burst size})}$$

- **Response time (mean)** This is generally measured as a monthly average. It is calculated when the trouble ticket is recorded and is measured until personnel respond:

$$\frac{(\text{total time in hours to respond})}{(\text{total number of trouble tickets})}$$

- **Time to resolution or repair (mean)** This is generally measured as a monthly average. It is calculated when the trouble ticket is recorded and is measured until the ticket is closed to the customer's satisfaction:

$$\frac{\text{(total time in hours to respond)}}{\text{(total number of trouble tickets)}}$$

## What Service Levels Should the Service Provider Consider?

Service providers need to be extremely careful in their negotiations with their customers. As many of the larger carriers know, a minimum number of sites should be negotiated. You should not enter into an SLA without this minimum site guarantee. You should also try to exclude items that will be out of your control. Be meticulous, as there are potential fiscal repercussions if you do not meet these service levels.

### Designing & Planning...

#### Items that Are Generally Excluded in an SLA

Items that you may want to exclude include:

- Acts of God
- The customer DSU/CSU
- The customer router
- Other customer access devices
- Customer-induced downtime
- Externally provided local loop
- Scheduled maintenance

Your customers will want to negotiate, and remember that they will move to another provider who will provide services and levels that they want. Maintain a strong point, but weigh the costs of what the customer wants against the potential loss that could occur by not getting the client. Again, most customers will want:

- Network availability
- PVC availability
- Average network delay
- Average PVC delay
- Effective throughput
- Response time
- Time to resolution or repair

## Network Availability

Most clients will want you to commit to a monthly guarantee of at least 99.5 (more often, 99.999) percent uptime. This guarantee generally includes all of the devices that are within your infrastructure, that connect to the local loop, or connect to the CPE. An uptime of 99.5 percent equals 3.6 total hours of downtime per month per site.

### Designing & Planning...

#### **The Difference between Network-based and Site-based Availability**

There is a distinction between network-based and site-based availability. For instance, if you have a client with a 10-site network, 99.5-percent network availability would allow for a total of 36 hours of downtime. If the SLA is based on site availability, then a site can only experience 3.6 hours of downtime. This is an important distinction when you are computing downtime.

## PVC Availability

Because the availability of network- or site-based SLAs usually does not meet business requirements for many of your clients, many companies will look to permanent virtual connection (PVC) availability, which restricts the amount of

downtime to single PVCs. This amount of availability is critical for networks that run applications that are sensitive to network delay or droppage. PVC availability includes (and excludes) all of the components that are within network availability.

## Average Network Delay and Average PVC Delay

Many potential guarantees are available; most of them depend on your network capabilities. Many of the largest companies guarantee a delay (round-trip) no greater than 300 milliseconds. You may be able to provide guarantees based on access line speeds, which can offer much lower delays for T1 and 64 kbps.

### Configuring & Implementing...

#### Measurement of Metrics Testing

Sometimes, you will hold the customer accountable for testing the measurement of delay. A word of caution, however: Often, customers will use packet Internet groper (ping) to test the delay during times of low traffic. There are two problems with this testing method: ping measurements include router delay, and pings have low network priority.

## Effective Throughput

You can interpret effective throughput in any way you wish. Some service providers base this category on the percentage of delivered frames based on a Committed Interface Rate (CIR) or frames that are labeled discard eligible (DE). Other providers base this calculation on the committed burst size rather than the excess burst size. You may be able to exclude configurations where the destination port is not configured to handle the bandwidth of the CIR. Some things that you can try to exclude include:

- Data that is lost during scheduled maintenance
- PVCs or other connections that were added or reconfigured during that month
- Any month that a client does not transmit an agreed-upon amount of data

## Response Time

Response time can be whatever number of hours that you and the client agree upon. There is a pretty standard method that says that you will respond within four hours of reported outage. This also depends on the location of the service provider from the maintenance center. Usually this maintenance only covers CPE, as your facility will be handled on an internal basis.

## Time to Resolution or Repair

Again, this is whatever number of hours that you and the client can agree upon, and depends on the type of failure and/or application that is running. For instance, if this is in support of a database for a client, restore time could include the retrieval of offsite backups. You should be very specific when defining time to resolution or repair.

## The Realities of Customer Compensation

Should a network outage occur, you should be able to quickly diagnose and repair the problem before it affects your clients. Many of your customers realize that they will never recoup all of the losses that will accrue if your system goes down. SLAs are not going to make your customers rich; they are trying to use your resources to make their business viable. Therefore, what they are interested in is reliability.

Many of your customers will want to know if you can find and fix issues (and potential issues) before they are affected. They will also most likely want to know if you will proactively fix issues, or wait for them to call and inform you. They will also wonder if you have the resources to meet the demand of the time to resolution or repair that is included within their SLA. In the customer's mind, compensation for downtime is not the correct answer, nor will it ever be. They just want you to take care of them, so that they in turn can take care of their clients.

## Designing & Planning...

### Are You Maintaining Your SLAs?

You will be asked to provide reports on a regular basis as to whether you are maintaining your SLAs. Generally, these reports will cover the metrics upon which you and your client agreed. In the past, many of these reports were skewed, with many of the metrics set out in a confusing type of way.

## What Will Your Customers Look for in Their Implemented SLA?

What will your clients look for in these reports on SLAs? Here are some things that your clients will ask you to do:

- Continually check that the WAN is capable of handling the services that they are providing.
- Verify that service levels are being maintained. This request may require your ability to show monitoring in real time.
- If services are not being met, then there must be an immediate path to resolution. This may be entirely your responsibility.

## What Are the Guidelines for Implementing the Monitoring Necessary to Handle These Tasks?

Baselining the network is a very important task, and will assist you in determining where potential problems could arise for you and your clients. Here are some of the common pitfalls that you may encounter:

- You will need to determine traffic patterns for your client's network connections. By understanding the application usage and peak-time utilization, you can better tune these connections to enhance efficiency.
- Help your customer understand your network; the core, and your policies. This will alleviate many of the common misconceptions that occur between you and the customer.

- Make sure that you are providing the best service, as well as maintaining your SLAs with your clients. It is imperative to maintain customers and grow your network. Remember that there is a lot of competition out there.
- Explain your monitoring and reporting infrastructure. Many clients will appreciate a well thought-out monitoring solution. This helps to show your customers that you are committed to their interests.
- Provide the baseline metrics for your network. This will help to give the customer an idea of what your infrastructure is capable of supporting, as well as the levels of efficiency that you can offer them. Make adjustments when necessary.
- Analyze your network and its reliability at least once a week. If you experience a poor-performance week, you may be able to save your SLA metrics on a week-to-week basis, as the reports usually go out once a month.

## Where Is Your Weakest Link?

Remember the saying, “A chain is only as strong as its weakest link.” Well, that really doesn’t equate to the ASP model. You see, ASPs aren’t even that strong. Many, if not all, of the components necessary to make an ASP viable are somehow inherently flawed when implemented in the overall picture. There is no way to create 100-percent uptime for each component within the ASP model. This doesn’t mean that the ASP model is bad; it means that you have to be more careful in your planning and deployment methods.

There is an equation that can help assist you in the planning and design phase of your ASP called total service availability percentage (TSA%). The TSA % is calculated using the equation  $(TSA\% = SA\%1 \times \dots \times SA\%N)$ . This equation is derived from the following pieces:

- Network provider [(WAN and Internet facilities)] 99.5 percent
- Infrastructure provider [(data center/system uptime)] x 99.5 percent
- Application management services [(application/fail-over services)] x 99.5 percent

Total service availability = 98.9 percent

As you can see, the individual components are guaranteed to have no more than 50 minutes of downtime a week (due to the 99.5 percent uptime guarantee).



However, if each component fails at a different time, there will be more than 2.5 hours of downtime, which is unacceptable. You may be able to define your availability so that it is measured in consecutive hours. What this means is that if service is restored within the 2.5 hours, the network has maintained the 99.5-percent uptime.

## Network SLAs

SLAs at the physical layer have been around for some time. For those companies that support their own networks and service providers, this level of monitoring allows them to report on their own performance by whatever metrics they feel are necessary. This level of granularity can assist you in negotiating the usage of other private networks or service providers, and still maintain your SLAs.

Many tools are available to assist you in the monitoring of your network. Several large vendors, such as Nortel Networks and Cisco Systems, include suites of tools that can allow you to monitor the performance of your network. You can use these tools with other third-party packages to automate and enhance the performance of your network.

## System Level SLAs

Many tools are available to monitor the systems in the data center environment. These tools are generally used to collect usage statistics and the percentage of uptime for devices. These packages will also inform a centralized management station of the number of outages, the length of these outages, the mean time between failures (MTBF), and the mean time to repair (MTTR).

Many service providers rely on default tools that function well with specific applications and servers. For example, providers that use Microsoft Windows Terminal Clients will usually rely on Citrix Resource Services Manager and Microsoft Systems Management Server (SMS).

Other server vendors such as Hewlett-Packard (HP), Sun, and Compaq offer their own tools for performance management and availability. These tools can be used in conjunction with or separately from default application tools. In fact, many ASPs build their performance-reporting capabilities around higher-level third-party vendor products.

## Application SLAs

You can also monitor the applications that the end users use. To do this, you will need to implement “smart agents” that are deployed at various collection points

within your server infrastructure. These agents can give you extremely accurate information as to the performance of these applications, but the measurements themselves can fluctuate depending on the usage by the end user. This makes it harder to interpret SLA agreements.

Several vendors make applications that you can use in these environments, including BMC Software, Compuware, and Candle. FirstSense Software was one of the first SLA management vendors to join the ASO Industry Consortium, and it has modified an existing Enterprise Monitoring Package to measure end-to-end response times. This package is also able to collect data from multiple providers, so that the end user can get a composite picture of SLAs from these companies.

## Making Your Company More Customer Oriented

So, if everyone is playing the same game, and all the toys are the same, what separates the service providers? By making their model more customer oriented, service providers can offer SLAs for things such as:

- Emergency response
- Response time guarantees
- Call center availability
- Remote troubleshooting

It isn't as though the customer doesn't care about availability, bandwidth, and latency; what they are looking for are the extras, the intangibles if you will. Service providers are moving to a more customer-centric model. They want you to know that the customer always comes first and is always correct. Many service providers today can give the same level and type of service. In order to differentiate themselves from their competitors, they need to maintain customer loyalty and build from that.

As the corporate infrastructure has evolved, so have the dynamics of the corporate network. What you are more apt to find in these changing times is an internal staff that handles and maintains very little of the overall network, remaining entirely within their walls or boundaries. External staff is comprised of the outsourced applications and infrastructure support. When you combine these two teams, you can encompass the range of support, including intranet-based Enterprise Resource Planning (ERP), electronic mail (e-mail), messaging,

scheduling, desktop support, operating systems, remote access, security, and other miscellaneous company needs.

## How Service Providers Have Responded

With all of the mission-critical applications that are available, many service providers are now offering services that are more advanced than the typical “leased line” connectivity that had been their bread and butter for so long. Leased lines were the lifelines to companies that needed direct access to their sites, and to their applications.

When Frame Relay (and to some extent, Switched Multi-megabit Data Service (SMDS)) became available, it was able to offer a service that was connectionless, thereby reducing the complexity of the service provider’s provisioning and circuit management. You were now able to increase your bandwidth for Internet connectivity and increase your uptime by adding meshed links and low-speed fail-over links to reduce network downtime.

There were other benefits with Frame Relay, including lower cost, and the capability to add layers of redundancy—which left no reason to go with leased lines.

Connecting a company to the widespread Frame Relay cloud did have at least one major advantage. With the Frame circuit in place, a company effectively gave up control of its circuits to the carriers. It was also hard to retrieve statistical data such as performance statistics, availability, and throughput. There are exceptions, such as private Frame Relay networks.

Remember I said that there would be consequences and repercussions? What I meant is that if a service provider failed to meet the required service levels, there is usually some type of financial compensation, or reimbursement for the customer. This is part of the reason that Frame was so popular with the carriers. Without the ability to see performance statistics, there was very little that the customer could produce to show quantifiable statistics.

As you can see, with no real way to implement checks and balances, along with the explosive growth of outsourced applications, there was a huge amount of customer dissatisfaction. Consequently, there was a major need for customer service, *the* differentiator in the market.

## Acceptable Performance

Service providers began to work with (as opposed to just for) their customers. They did this by helping customers design their infrastructure and their WAN

connectivity. This exposed the service providers to the requirements that are faced within the enterprise, which helped to put into perspective the reliance that corporations were putting into the Internet, and therefore, the provider.

With all of this realization came a boom in applications that were more aware of the network. These applications were comprised of programs such as ERP, supply chain management (SCM), content, VPN, and thin clients that were able to reduce network traffic and enhance user experiences.

## The Added Bonus

There were benefits to this client interaction. Service providers were able to fully utilize their infrastructure by overbuilding or undersubscribing their network, which helped them meet the SLA, and therefore helped to keep the customer happy. Service providers began to notice that users were more than willing to pay higher rates, if there was guaranteed service (as in the case of an overbuilt network). That made it easier to justify the enhancements to the network to maintain SLAs.

Originally, service providers were working on the POISSON model, in which they would oversubscribe their network, take their profits, and turn it back into a larger infrastructure. This was somewhat of a catch-22 (really more of a vicious cycle) that continues to this day. So, what changed the model away from the oversubscription model? Customers were more than willing to pay higher rates in order to receive guaranteed service levels. As discussed earlier, there will be a penalty if you are unable to meet your SLAs. It really doesn't require an MBA to understand that the more penalties you pay, the less effective your cash model will be.

## The Operation Support System Model

The Operations Support System (OSS) model usually refers to a system (or systems) that can perform the management necessary to maintain and monitor your SLA requirements. This model takes the following items into account:

- Performance management
- Inventory control
- System engineering
- Design
- Support

In the Beginning... (I just like to say that), OSSs were mainframe-based, stand-alone devices that were implemented to assist the telephone companies in doing their jobs. These systems were designed to automate manual processes for efficiency and reliability. Today, service providers need to manage more sophisticated devices and environments.

As with all things, the more things change, the more they stay the same. With all of these dynamic service providers growing and changing, they noticed that there was a need to update these OSS tools that were lagging behind the technology. Companies needed OSS tools that would add to their efficiency, and therefore help their bottom line (or return on investment (ROI)).

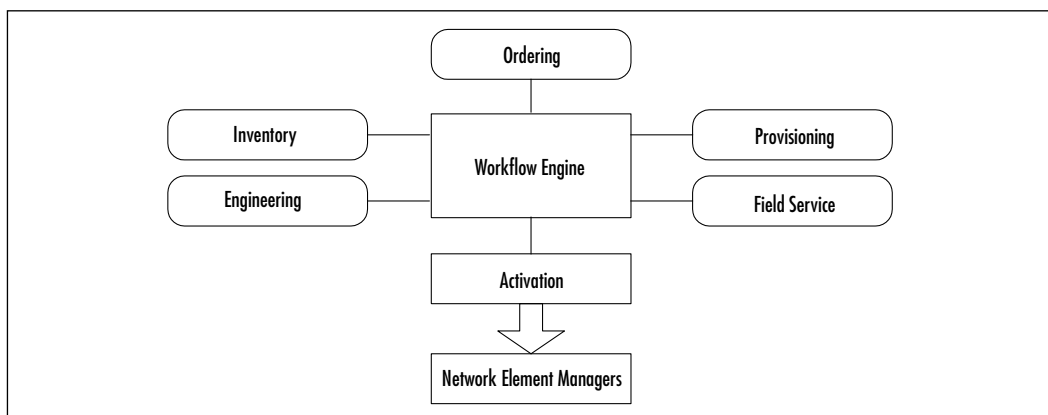
## What Are the Basics of OSS?

In order to truly understand OSSs, you must first become familiar with some of the fundamental systems that are involved. These systems handle the functions of ordering, service fulfillment (such as voice, data, and other IP-based services), inventory, circuit provisioning, and activation.

### The Workflow Engine

The workflow engine is the core of an incorporated OSS solution. The engine allows the service provider to better manage the flow of traffic and disseminate it to disparate systems. This engine helps to enable the service provider to complete tasks in a timely and efficient manner. Some OSS vendors have packages that incorporate the workflow engine in their offerings; other companies specialize in this area (Figure 7.2).

**Figure 7.2** Process Workflow System



## Ordering

Ordering is one of the more important parts of the OSS model. This is where you manage the information that is necessary to provide your services. This process will allow you to monitor and manage your clients and your relationships with your suppliers and partners.

Many of today's ordering suites include some form of graphical user interface (GUI), which makes training less necessary and less costly. It also helps you to complete more orders quickly and accurately, and can even allow you to give your clients the ability to provision their resources from you through Web-enabled technologies.

When an order is started (and completed), it will generate a number of tasks and procedures that interact with other pieces of the OSS model. For instance, when a customer orders, the ordering system will usually check the inventory and then process the work through the workflow engine.

## Inventory and Allotment

An inventory system is used to manage information about your infrastructure. When an order is placed, there must be sufficient resources to handle it. There is an immediate inventory take in network design and provisioning, so the inventory and allotment piece must be able to connect and interact with other management pieces.

## Engineering and Provisioning

Engineering and provisioning systems allow providers to manage, monitor, and reallocate resources within their infrastructure. These systems are often integrated with the network design portion of the OSS model. This is often referred to as the "Design and Assign" system.

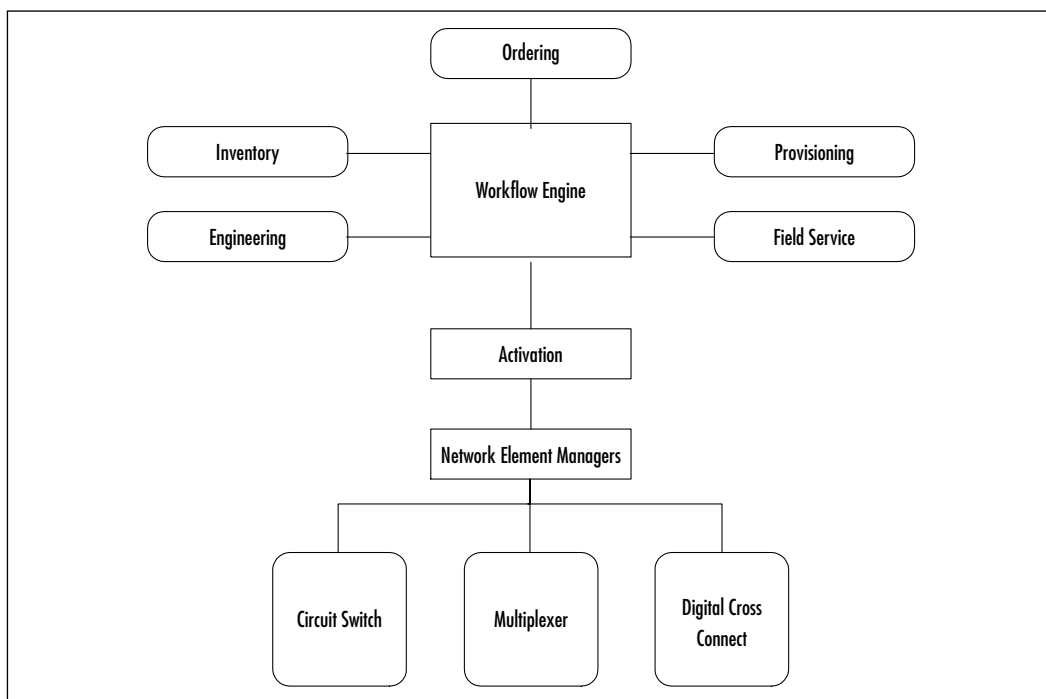
## Activation and Service Management for the Field

When a service has been ordered, engineered, and provisioned, the services will then be installed and activated. Activation is comprised of several steps. Does new equipment need to be installed? If the answer is "yes," then you will need to allocate field resources to handle the installation and configuration. When the installation is complete, the technician must then contact the central office so that there can be turn-up.

If the answer is “no,” then you may be able to automatically—er, automatically—turn up the circuit. In these instances, due to the integration of the packages, there can be an ordered circuit that may never need human interaction in order to be turned up.

Many of today’s networks and OSSs are designed with some sort of built-in management, whether it is Common Management Information Protocol (CMIP), Transaction Language 1 (TL1), or SNMP. With these management tools, an OSS activation system can work as a “Manager for Managers” by supervising the various devices within the network (Figure 7.3).

**Figure 7.3** The “Manager for Managers” System

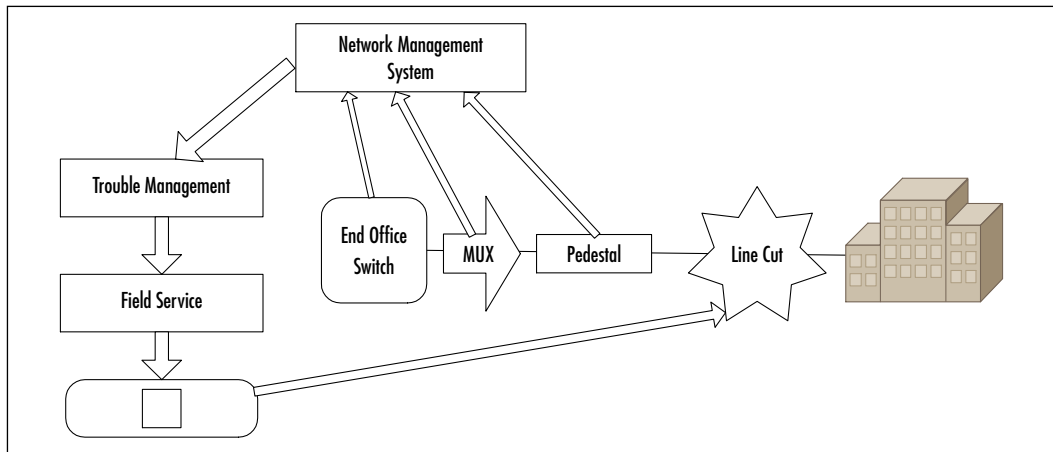


## Network Management and Support

In the grand scheme of things, OSS functionality does not end with service activation. In fact, once the installation and activation are over, there is still a lot of work left to do. This work falls into two main areas: network management and support. Network management systems are in charge of monitoring the network. Management generally uses protocols such as SNMP and CMIP to communicate between network components. These protocols collect data on the performance

and utilization of the devices located within the network. If there is a problem, these systems will notify the proper support staff, usually through a centralized Network Operation Center (NOC). The NOC verifies the problem, and assigns the proper support resources to troubleshoot and repair the issue. The NOC is also able to use the other elements of the OSS and reroute traffic around these trouble spots (Figure 7.4).

**Figure 7.4** Trouble Management System



## What Is OSS Interconnection, and What Does It Mean?

In 1996, a Telecommunications Act was created to deal with OSS interconnection. When referenced this way, *interconnection* refers to the policies that are needed by the Regional Bell Operating Companies (RBOCs) to allow their competitors at least limited access to their customer databases and OSS information gathering, such as pre-ordering, ordering, and provisioning. Pre-ordering is the method that a Competitive Local Exchange Carrier (CLEC) who has received customer consent uses to request information about the customer from an RBOC.

The Federal Communications Commission (FCC) does not permit RBOCs to enter the long-distance market until they can provide an interconnection that is able to provide competition. RBOCs and Incumbent Local Exchange Carriers (ILECs) have interfaces that are able to provide interconnections to CLECs. Interconnections are extremely complex and time sensitive, and the communications industry has been trying to resolve these issues for years.



## What Are the Challenges Facing Interconnection?

One of the largest issues that is slowing the sharing of information from RBOCs is that their OSSs are normally proprietary systems that were not designed to share information between disparate systems. RBOCs have spent a lot of money on these systems, and they don't want to discard them and ruin their ROI. Therefore, they need to find a way to maintain these proprietary systems and still meet government directives to share information.

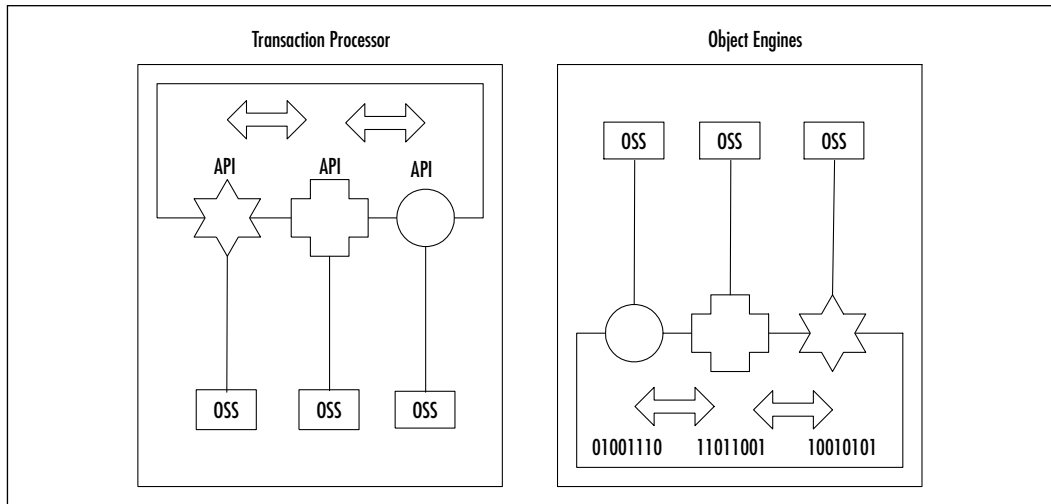
## Upgrading the OSS

As we said earlier, RBOC systems weren't designed to share customer data. They weren't able to store and distribute the data they collected from customers who received services from CLECs. To be able to handle the next generation of interconnection technology, RBOCs systems will need to be able to respond to incoming interconnections to fulfill CLEC requests for customer data.

There are many different approaches to making OSSs integrate between RBOCs and CLECs. There is the ability to create middleware or "glue code" or transaction processes (TP) that work in between these systems to transfer information between disparate systems. This glue code is usually comprised of common application programming interfaces (APIs) that can incorporate and manage data translation and dispersion (Figure 7.5).

When workflow systems are used in conjunction with glue code, you are able to provide many dynamic APIs that manage tasks and data traffic flow; all while the TP is handling the data's conversion. Object-based engines, such as those that use technologies such as Object Management Group's (OMG) Common Object Request Broker Architecture (CORBA) or Microsoft's Distributed Component Object Model (D-COM), are able to summarize application interfaces into defined, yet dynamic, software objects to intercommunicate.

There is currently no standard on how to integrate OSSs. In fact, RBOCs usually rely on technologies that are already available to exchange information with customers and InterExchange Carriers (IXCs). Most of these interfaces were not intended to be used as interconnection platforms. These platforms are generally considered to provide some of the most efficient and most economical benefits for RBOCs, because a large amount of the necessary glue code is already in place. Using this method usually requires the use of electronic data interchange (EDI). EDI was created to share documents between businesses, but is now commonly employed for ordering and pre-ordering.

**Figure 7.5** Integrating OSS Technologies

## Efficiencies in Your OSS

Many of today's OSS solutions are considered commercial off-the-shelf (COTS) packages. These applications are able to offer some out-of-the-box utilities and are intended to be modified to meet customer needs. This customization could allow your company to integrate management capabilities and enable your customers to take advantage of your services, thus adding efficiency.

## Remaining Flexible

Due to the dynamic nature of the networking world and customers' wants and needs, it is a core requirement that you remain flexible your solution. Try to remain as vendor neutral as you dare; that way, you are not locked in to a platform that is no longer able to support you and your market. Find technologies that will allow you to respond immediately to changes in the marketplace, whether those changes come from marketing, new technology, or some regulatory requirements.

## API Functionality and Gateways

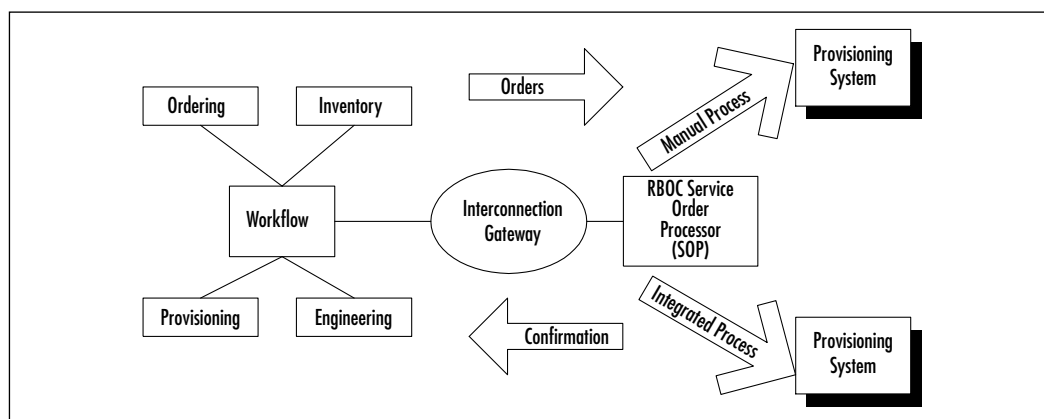
There are numerous API and gateway packages available to the OSS market. These products were intended to assist CLECs in the development of interfaces that are necessary to interconnect RBOC OSSs. There is an industry organization that primarily devotes itself to the implementation of standards-based practices in

the telecommunications market. This forum, called the TeleManagement forum, tries to assist carriers in the deployment of methods such as the Telecommunications Management Network (TMN) model, and has helped to create the guidelines for a Common Interconnection Gateway Platform (CIGP). CIGP is a way to create vendor-neutral, and therefore nonproprietary, technologies that are common across the market, which can help CLECs create interconnection interfaces.

Gateways and APIs are used to manage these interfaces between CLEC and RBOC OSS interfaces, and are implemented to maintain data integrity and security between carriers and customers when data is exchanged. Remember that CLECs and RBOCs customers will want the utmost security, as they have to give permission for their information to be shared. This information is normally transferred between carriers using Universal Service Order Codes (USOCs). There are literally thousands of codes, all of which are very cryptic.

Today, gateways are able to read these USOC codes and match them to CLECs that offer a catalog database, so that they can generate product offerings to customers. As you can see, this is far more efficient, and allows customers to access services that they are most likely to use (Figure 7.6).

**Figure 7.6** The Interconnection Process



## Supporting Your Data Services


Your OSS solution must be able to support increasingly complex clients. With new technology being presented and older technologies gaining in popularity or use, your infrastructure must be able to accommodate things such as Frame

Relay, cable, digital subscriber lines (DSL), and Asynchronous Transfer Mode (ATM), all used in conjunction with IP.

## Provisioning Data Service

With the advent of broadband technology, service providers now need to address bandwidth between two locations very carefully, as there will be certain QoS and SLA needs associated with these connections (Figure 7.7). After the equipment is provisioned, the provider must determine the layout for the mapping of services to connections.

**Figure 7.7** A Basic Service Order Work Request



## Activation of Data Services

In an effort to support end-to-end efficiency and an automated process, you, as a service provider, must be able to pass information to the network management layer (NML) in order to activate your client connections. The NML can activate the proper devices with very little user interaction. In fact, today's OSS is capable of providing activation in real time with the communication that can be achieved between the NML and the service management layer (SML).

## Broadband Access Changes the Market

That may seem like an odd statement, but it is entirely true (maybe too true). You see, broadband access has changed the way we do business, and how we live at home. At this moment in time, DSL and cable are surpassing every other method of access across the United States. This isn't to say that Frame or other connections are going to disappear; it is really saying that, like everything else, things change.

Many of today's service providers are struggling with the deployment of these technologies. It's not because they don't have the bandwidth; it's because it is difficult to maintain and upgrade your infrastructure if you are unable to see your current copper allocation (for the local loop) and resource availability. One of the ways that a central office (CO) can handle these issues is to have an up-to-date, dynamic inventory of provisioning, as discussed earlier in the chapter.

## Getting Access to the Masses

The CO must have a good management platform, and be able to provide the connection to the customer. As stated earlier, DSL and cable are two of the more popular access methods within the United States. In order for a service provider to incorporate DSL within its infrastructure, there is the need to integrate two components:

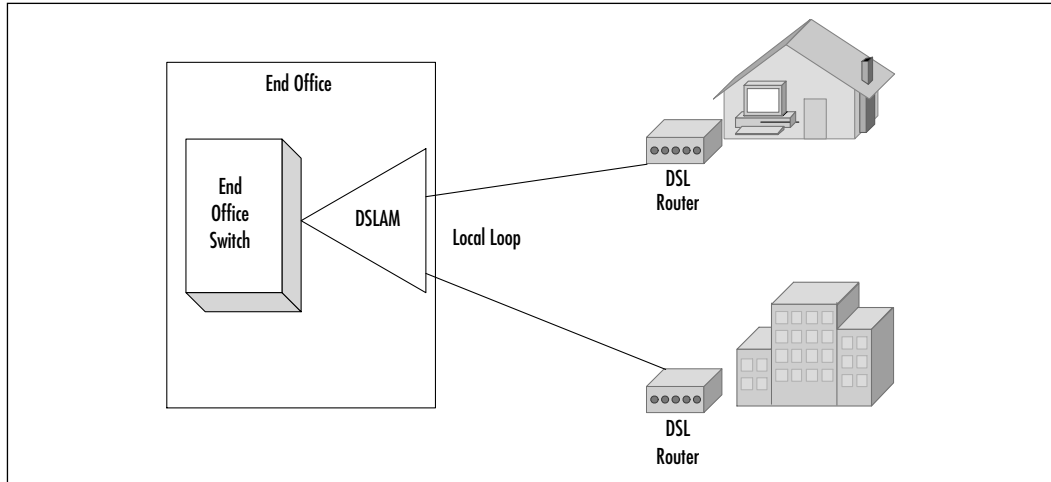
- A splitter
- DSL Access Multiplexer (DSLAM)

A splitter distributes voice traffic to the Plain Old Telephone System (POTS) cloud, and data traffic to the DSLAM. It is becoming more apparent that the splitter is not long for the planet, as the demand for "all-in-one" solutions is rising.

A DSLAM is able to communicate with the DSL router (I really don't like the DSL modem terminology, because DSL is digital; there is no modulation/demodulation on the circuit) that is located on the customer's premises. The

DSLAM aggregates multiple DSL connections into a Layer 2 device that are able to offer high performance and various multiplexing schemes (Figure 7.8).

**Figure 7.8** The DSL Access Multiplexer (DSLAM)



## Configuring & Implementing...

### Some Concerns for DSL

DSL is not a finished technology; in fact, there is still a lot of polishing that needs to be done. As you may know, DSL works on specific frequencies. Line interference caused by other devices on the line that use the same frequencies will render DSL inoperable. In addition, many users may not have wiring that will support DSL, or they may be beyond the distance that DSL can achieve.

## Quality of Service

Quality of Service (QoS) is a measurement of the service value. Measurement of QoS is very subjective; it depends on the technology on which it is implemented to see if there are acceptable levels of performance. For instance, if you have a dedicated 56 k connection (say for a small office in Alaska) that only does service

delivery, you will experience a high quality of service because the traffic for specific services are the only sessions that can use these links.

On the other end of the spectrum, if you have a call between two sites, and the users are not talking, then you are wasting valuable bandwidth. Remember that in the networking world, where time is measured in milliseconds and microseconds, even one second of unused connection time is a huge waste of bandwidth and resources. This will not help in your overall QoS levels.

There is a way to improve QoS across the board. IP is one of the technologies that can use multiple paths to get the most out of all available bandwidth. IP uses only as much bandwidth as it needs, which allows it to use multiple paths. The drawback with IP is that it uses best-effort delivery as its method of operation, which can lead to lost traffic and poor QoS to the customer.

You will need to maintain a high level of QoS to maintain and attract new customers. Therefore, you should implement and manage your solution so that it is capable of meeting your customers' expectations. QoS will vary from customer to customer, so tailor your SLAs to reflect client needs; for example, a bank that may need to implement high-speed transport (ATM) and VPNs.

## Management Systems for Your ASP

Many of today's service providers use (at least at some level) the Telecommunications Management Network (TMN) model. The TMN model provides the outline for attaining interconnectivity and communications across diverse platforms and environments. TMN was developed by the International Telecommunications Union (ITU) as a tool to help support, manage, and deploy services. TMN was originally based on the common management information service element (CMISE).

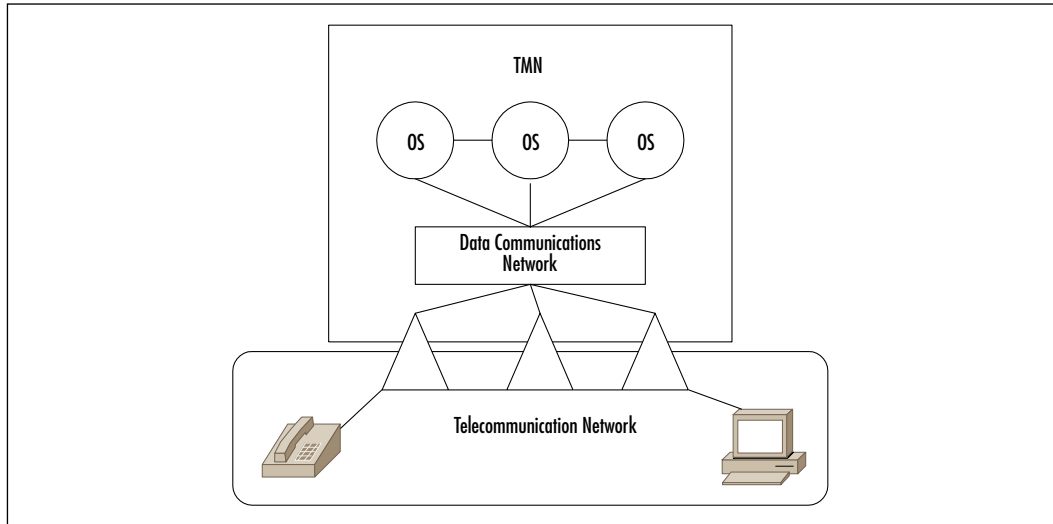
### The TMN Outline

The TMN model outlines what is necessary to make your network infrastructure flexible, scalable, manageable, and highly available. TMN defines standard ways of handling management tasks and communications across networks. TMN allows you to distribute the appropriate levels for growth, efficiency, and communication performance.

The principles brought forth by the TMN model can be incorporated into the network (Figure 7.9). Remember that at its most basic levels, the infrastructure is composed of routers, switches, circuits, and so forth. In TMN terms, these

components are referred to as network elements (NEs). TMN is there to enable communication between the OSS and NE.

**Figure 7.9** How the TMN Model Fits into a Telecommunications Network



## TMN Standards

When service providers implement TMN standards, their services become interoperable; if all providers used the TMN model, then all infrastructures would be able to communicate. TMN uses principles that are object oriented and standard interfaces to define communication between management nodes. This standard interface is called the Q3 interface.

TMN is defined by the ITU M.3000 recommendation series. This is derived from the Seven Layer Open System Interconnect (OSI) model, and includes:

- **Abstract Syntax Notation One (ASN.1)** Provides the syntax rules for data types.
- **Common Management Information Protocol (CMIP)** Defines the management services that are traded between nodes.
- **Guideline for Definition of Managed Objects (GDMO)** Creates the model for organizing and describing managed resources.
- **Open Systems Interconnect reference model (OSI)** The Seven Layer OSI reference model (as discussed in Chapter 1, “An Introduction to ASPs for ISPs”).



TMN has been embraced and propagated by other standards bodies, including the Network Management Forum (NMF), Bellcore, and the European Telecommunications Standards Institute (ETSI). NMF and Bellcore are trying to accelerate the deployment of the TMN model by creating generic outlines for establishing requirements. This isn't limited to just these organizations; the Synchronous Optical Network (SONET), the Interoperability Forum (SIF), and the Asynchronous Transfer Mode Forum (ATMF) are all moving toward TMN-compliant management interfaces.

Within the TMN model, management functions are performed by operations that are included within the Common Management Information Services (CMIS). Network-managed information and rules are contained in a package called the Management Information Base (MIB).

The MIB processes management. Processes that manage the information are called management entities. These entities can take one of two roles, that of manager or agent. The manager and agent processes will send and receive requests and notifications by using the CMIP.

In addition to the TMN-layering structure, the ITU also splits the general-management functionality offered by systems into five key areas:

- Fault
- Configuration
- Accounting
- Performance
- Security

This is also referred to as FCAPS. This categorization is a functional one and does not describe the business role of a management system within the network. The idea of FCAPS stems directly from the ITU recommendations and describes the five different types of information handled by management systems. Portions of each of the FCAPS functionality will be performed at different layers of the TMN architecture. For instance, fault management at the element management layer (EML) is detailed logging of each discrete alarm or event. The EMS then filters the alarms and forwards them to an NMS that performs alarm correlation across multiple nodes and technologies to perform root-cause analysis. A subset of the FCAPS functionality is listed in Table 7.1.

**Table 7.1** A Subset of the FCAPS Functionality

Fault Management	Configuration Management	Accounting Management	Performance Management	Security Management
Alarm Handling	System Turn-up	Track Service Usage	Data Collection	Control NE Access
Trouble Detection	Network Provisioning	Bill for Services	Report Creation	Enable NE Functions
Trouble Correction	Autodiscovery	—	Data Analysis	Access Logs
Test and Acceptance	Back up and Restore	—	—	—
Network Recovery	Database Handling	—	—	—

## The Building Blocks of the TMN Model

The TMN model is represented by several building blocks, all of which can combine to provide an overall personification of the management issues and roles of TMN. Figure 7.10 illustrates the TMN building blocks.

**Figure 7.10** The TMN Building Blocks

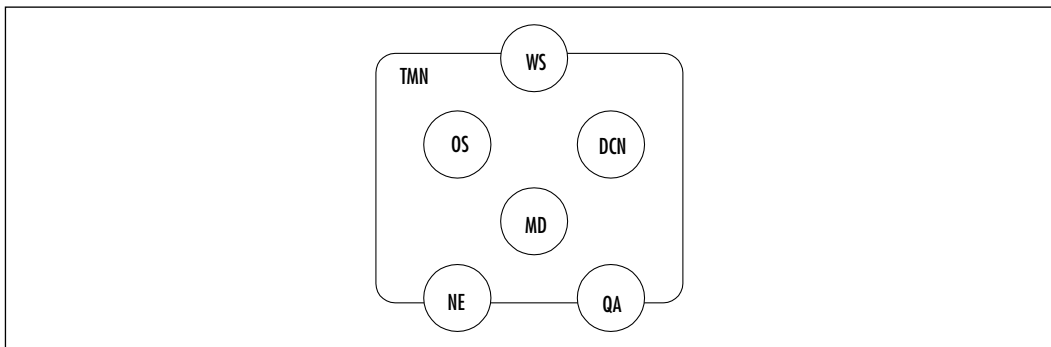


Table 7.2 lists and describes each TMN component and the role it performs within the TMN model. In some cases, these roles may be performed in conjunction with other system components. The mediation device (MD), for example, also provides some of the functionality that is defined as part of the operations systems (OSs), Q adapters (QAs), and workstations (WSs). Conversely, the OS may also provide some of the MDs, QAs, and WSs.

**Table 7.2** Roles of Components in the TMN Model

<b>System Component</b>	<b>Description</b>
<b>Operations Systems (OS)</b>	The operations system component performs operations and system functions. This can include operations that monitor and control telecommunications and management functions. The OS component can also provide some of the mediation, Q adaptation, and workstation responsibilities.
<b>Mediation Device (MD)</b>	The mediation device performs negotiations between local TMN interfaces and the OS information model. The mediation role may be needed to ensure that the information, scope, and functionality are presented in the exact way that the OS expects. Mediation functions can be implemented across hierarchies of cascaded MDs.
<b>Q Adapters (QA)</b>	The Q adapter enables the TMN to manage network elements that have non-TMN interfaces. The QA will translate between TMN and non-TMN interfaces. For instance, a TL1 Q-adapter translates between a TL1 ASCII message-based protocol and the CMIP, the TMN interface protocol; the same way that the Q-adapter translates between SNMP and CMIP.
<b>Network Element (NE)</b>	A network element contains manageable information that can be monitored and controlled by an operations system. An NE must have a standard TMN interface to be managed within the scope of the TMN model. If an NE does not have a standard interface, the NE can still be managed via a Q adapter. The NE provides the OS with a representation of its manageable information and functionality. As a building block, the actual NE can also contain its own OS function, as well as QA function, MD function, etc.
<b>Workstation (WS)</b>	The workstation performs the role of translating information between TMN format and a displayable format for the user.
<b>Data Communication Network (DCN)</b>	The DCN is the communication network that is located within a TMN. The DCN represents OSI Layers 1 through 3.

## How the OSI Functions in the TMN Model

The TMN model is designed to define a message communication function (MCF). All building blocks with physical interfaces require an MCF. An MCF is

able to provide the protocol layers necessary to connect a block to a DCN (for example Layers 4 through 7). An MCF can provide connectivity to all seven OSI layers, and it can provide protocol convergence functions for interfaces that use some other layer configurations.

## Manager and Agent Roles

As stated earlier, the TMN function blocks can act in the role of manager and/or agent. The manager/agent are the same as those that are used for CMIP and OSI management. In other words, a manager process issues directives and receives notifications, and an agent process carries out directives, sends responses, and produces events and alarms. A building block may be viewed as a manager to one peer, even though it is viewed as an agent to another peer.

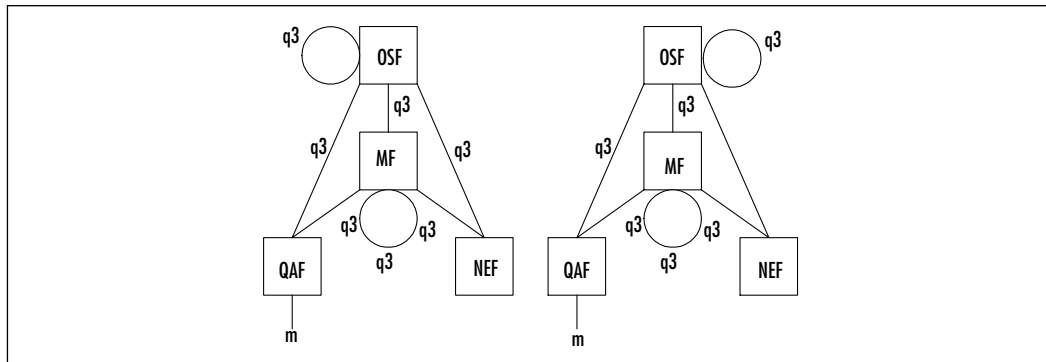
## The Standard Interfaces

In the TMN model, there are specific interfaces between two TMN components that need to communicate with each other (Table 7.3).

**Table 7.3** Standard TMN Interfaces

Interface	Description
Q	The Q interface is the interface that exists between two TMN functional blocks that are within the same TMN domain. The Qx will then carry the information that is shared between the MD and the NEs that it supports. The Qx interface exists between the NE and MD; QA and MD; and MD and MD. The Q3 interface is the OS interface. Any component that interfaces directly to the OS uses the Q3 interface. Therefore, the Q3 interface is between the NE and OS; QA and OS; MD and OS; and OS and OS.
F	The F interface is located between a WS and OS, and between a WS and MD.
X	The X interface is located between two TMN OSs that are located in two separate domains, or between a TMN OS and another OS in a non-TMN network.

There are two other reference points, G and M, that are outside the scope of TMN. They are between non-TMN entities and other non-TMN portions of the WSF and QAF, respectively. In Figure 7.11, each line represents an interface between two TMN components.

**Figure 7.11** Standard Interfaces between TMN Components

## The Logical TMN Model

The Logical Model of the TMN supplies layers that define the management level for specific functionality. These functions can be implemented at many levels; from the highest level, which manages corporate or enterprise goals, to a lower level, which is defined by a network or network resource. Starting with the bottom level, these hierarchy layers include Network Elements, the Element Management Layer (EML), the Network Management Layer (NML), The Service Management Layer (SML), and the Business Management Layer (BML). When the management is defined at the lower layers, additional management applications can be built on this foundation (Table 7.4).

**Table 7.4** The Logical Layers of the TMN Model

Layer	This layer is concerned with:
<b>Business Management Layer (BML)</b>	The Business Management Layer was created for high-level planning, budgeting, goal setting, executive decisions, business-level agreements (BLAs), etc.
<b>Service Management Layer (SML)</b>	The Service Management Layer uses information presented by the Network Management Layer to manage contracted services for existing and potential customers. This becomes the basic point of contact with customers for provisioning, accounts, quality of service, and fault management. This layer is also a main point for interaction with service providers and other administrative domains. It maintains statistical data to support quality of service, etc. OSs in the SML interface with OSs in the SML of other administrative domains via the X interface. OSs in the SML interface with OSs in the BML via the Q3 interface.

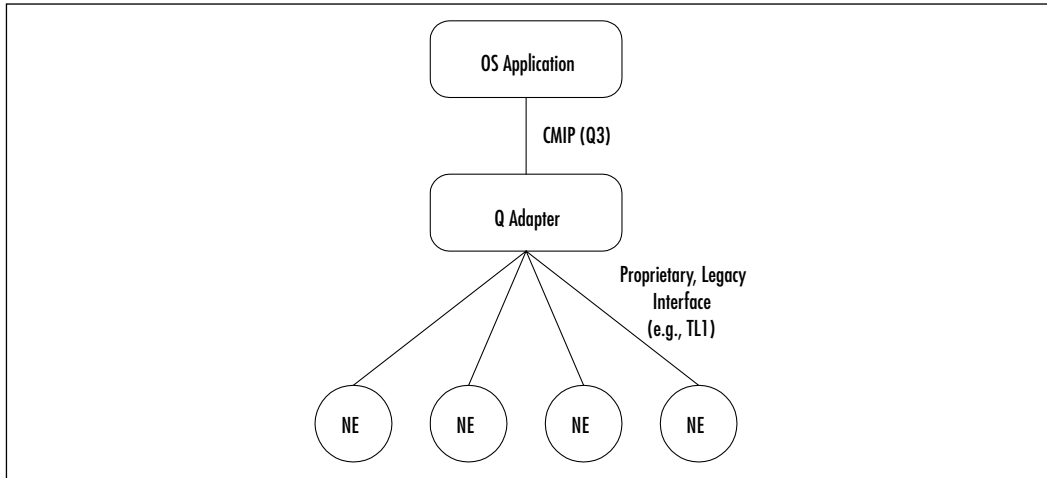
Continued

Table 7.4 Continued

Layer	This layer is concerned with:
<b>Network Management Layer (NML)</b>	The Network Management Layer is visible to the entire network, and is based on the NE information that is presented by the Element Management Layer OSs. The NML manages individual NEs and all NEs as a group. The NML receives the first managed view of the network. The NML then coordinates all network activities and supports the demands of the SML. OSs in the NML interface with OSs in the SML via the Q3 interface.
<b>Element Management Layer (EML)</b>	The Element Management Layer manages each network element. The EML is composed of element managers, or OSs, each of which is responsible for the TMN information for certain NEs. In general, an element manager is responsible for a subset of the NEs. An element manager can manage network element data, logs, activity, etc. In a Logical view, MDs are in the EML, even when they are physically located in some other logical layer, such as the NML or SML. An MD communicates with an EML OS via the Q3 interface. An EML OS presents its management information from a subset of the NEs to an OS in the NML through the Q3 interface.
<b>Network Element Layer (NEL)</b>	The Network Element Layer presents the TMN information for an individual NE. Both the Q-adapter, which adapts between TMN and non-TMN information, and the NE are located in the NEL. Therefore, the NEL interfaces between the proprietary manageable information and the TMN infrastructure.

There is also the Element Management System (EMS) that manages one or more of a specific type of telecommunications NE. Typically, the EMS manages the functions and capabilities within each NE, but does not manage the traffic between different NEs in the network. To support management of the traffic between itself and other NEs, the EMS communicates upward to higher-level NMS. The EMS provides the foundation to implement TMN-layered operations support system (OSS) architectures that enable service providers to meet customer needs for rapid deployment of new services, as well as meeting stringent QoS requirements. Figure 7.12 demonstrates how a Q adapter translates between CMIP/Q3 and Proprietary Interfaces.

**Figure 7.12** A Q Adapter Translates between CMIP/Q3 and Proprietary Interfaces



## What Tools Do You Need to Automate TMN?

A multitude of tools are available to automate the task of building TMN agent or manager applications. You can deploy and tailor the TMN agent and manager toolkits to match your company's GDMO/ASN.1 MIB representations. These products should have the following features in order to take advantage of the TMN model and to most productively support a TMN infrastructure (Table 7.5):

- **Automated prototyping** These tools can compile GDMO/ASN.1 information models and produce model-specific interfaces and other reports.
- **Conformance to all TMN standards** This is very important with the implementation of service, data, and managed object layers of the NMF API, and support for specific and generic application types.
- **Dynamic information modeling** This feature allows you the ability to add or change the network configuration, or modify functionality without reinstalling or recompiling applications and implementations.
- **Management Information Base (MIB)** These are the building blocks that help the developer to construct a GDMO/ASN.1 information model for any managed network.

- **Platform-independent interfaces and tools** These tools are used for testing and simulating the behavior of a manager or agent, not on its implementation. These tools are able to act in the role of agent, manager, or both.
- **Q adaption capability or compatibility** This is the ability for the interface to integrate legacy NEs (such as TL1 message-based equipment types) as well as enterprise network systems (SNMP-based information).
- **System management functions (SMFs)** These tools can generate, filter, forward, and log incoming events and alarms.

**Table 7.5** The Five-Layer TMN Network Management Architecture

Layers	Description
<b>Business Management Layer (BML)</b>	Manage the overall business; e.g., achieving return on investment, market share, employee satisfaction, community and governmental goals.
<b>Service Management Layer (SML)</b>	Manage the service offered to customers; e.g., meeting customer service levels, service quality, cost, and time-to-market objectives.
<b>Network Management Layer (NML)</b>	Manage the network and systems that deliver those services; e.g., capacity, diversity, and congestion.
<b>Element Management Layer (EML)</b>	Manage the elements comprising the networks and systems.
<b>Network Elements Layer (NEL)</b>	Switches, transmission, distribution systems, etc.

\* FCAPS have different tasks at each layer.

## The ASP Transformation

To transform from an ISP to an ASP, you will need a service management solution that is designed specifically to manage the unique functions and processes of ASPs with carrier-class reliability and scalability.

To fully leverage scale economies and provide aggressively priced services, large ASPs typically use a common hosting infrastructure shared among their customer base. This model requires service management solutions dedicated to



handle the unique requirements of an application-hosting environment consisting of a large library of applications from multiple vendors hosted on multiple operating systems for multiple enterprise customers with large volumes of users.

In addition, the ASP value proposition and target market requires that a service provider's service management infrastructure be capable of developing and managing flexible service offerings comprised of any combination of applications, resources, and services with a variety of flexible pricing capabilities and service level options.

Furthermore, the market requires that the ASP provide these capabilities under conditions of high growth and escalating transaction volume. The emerging ASP cannot afford to experience the growing pains typically encountered by developing service technologies. You must provide carrier-class reliability and scalability, as enterprise customers will demand bulletproof dependability from service providers managing their mission-critical systems.

All of these ASP service management requirements impose considerable and unique challenges to an ASP that standard enterprise or ISP tools are incapable of resolving. Additionally, it may be impractical for ASPs to develop this service management system internally. The resource requirements for the continual development, enhancement, and maintenance of a middleware software system capable of supporting a large number of sophisticated operational functions under severe growth conditions are substantial, and thus require skills outside the core competency of an ASP.

One of the biggest challenges ASPs face is proving that they can manage and maintain enterprise applications with flexible, high-quality services. ASPs are striving to win the trust of users with new management systems that promise easier service provisioning, a variety of billing options, and more detailed SLAs.

Businesses can expect ASPs that offer enterprise application hosting services to introduce features such as better SLAs and service options. Service providers believe that users are willing to pay more for better services.

## Industry Examples of Successfully Deployed ASP Management Tools

Futurelink, Corio, and USinternetworking (USi) are some of the industry leaders that have already been deploying management tools that will lead to enhance service offerings.

Futurelink, based out of Irvine, California, offers application hosting, thus allowing users more flexibility with a product called ASP Workbench, a new

management system from Xevo. Futurelink is deploying ASP Workbench to offer its customers a simplified application hosting service provisioning and new billing options for those services.

This management platform lets an ASP synchronize application configurations and user application parameters on a centralized database. This means that when a customer signs up, Futurelink will gather the customer's application and user parameters that apply to a firm's entire employee base. Therefore, when a new hire comes on board, it only takes seconds to get that user into the system.

While Futurelink is installing Xevo's ASP Workbench, Corio is fine-tuning its automated provisioning system. One of its biggest selling points for ASPs is that they can get business users up and running quickly, which is why Corio is automating its provisioning process. Corio is using software from Chainlink that has been customized to easily bring customers online. Chainlink works off a list of processes defined by Corio. Corio is able to create feature preferences and how application patches and upgrades should be distributed for Corio customers.

For instance, by automating the service setup and maintenance process, Corio would only have to deploy a software patch once for all customers who subscribe to Corio's PeopleSoft application hosting services. In addition, bringing users online from a single company would mean only entering in each user's ID and password instead of each user's application preferences.

Corio also uses a Netegrity product called SiteMinder to integrate different applications such as those from Siebel Systems into a Citrix-based environment. Corio has a central repository that stores all of the customer accounts so they can use a single sign-on and reach all the applications to which they subscribe. Without SiteMinder, customers would have to log on each time they wanted to access a different application.

USi, an ASP in Annapolis, Maryland, is also emphasizing its management and monitoring systems. USi has an extensive Tivoli network management system, which it customized to monitor network outages and bottlenecks. The ASP is also developing SLAs that will be based on a new transaction monitoring system, and will offer business users more detailed service guarantees.

## ASP Infrastructure Operations

What if service management systems could become smart enough to generate a service portal for every user, and then tie that portal to personalized allocation of data center and network capacity? What if an ASP end user could select a new application just by clicking on a service portal icon? What if every time that user

logged on, the entire data center and network devices could know who he or she is, know the application, and allocate just the right server power, OS memory, and database links for the task at hand? What if performance monitoring and subscriber profile systems could conspire to tweak network bandwidth knobs, protecting premium users from momentary traffic spikes?

Questions such as these have prompted various vendor partnerships, acquisitions, and new product launches this year, making such scenarios less and less far-fetched. Indeed, many of the players argue that ASP service growth itself will prove far-fetched if these scenarios do not become commonplace.

Provisioning for an application infrastructure provider has to touch so many components that you simply have to automate the workflow. Otherwise, you will have too many operations people and not enough profit.

For every user sign-on to an ASP network, the workflow can involve provisioning resources including firewalls, application servers, server hardware processor cycles, OS memory, backend databases, load-balancing devices, encryption and compression devices, routers, switches, and edge access devices.

At the same time, OSSs and BSSs must be provisioned to support a whole range of customer care, usage record keeping, billing, network device monitoring, performance analysis, troubleshooting, and other service delivery and business applications.

## Network Operating System

Xevo Corporation ([www.xevo.com](http://www.xevo.com)) announced its XevoWorks Partner Program with some marquee names across the ASP infrastructure chain:

- Cisco Systems Inc. ([www.cisco.com](http://www.cisco.com))
- Citrix Systems Inc. ([www.citrix.com](http://www.citrix.com))
- Compaq Computer Corp. ([www.compaq.com](http://www.compaq.com))
- Great Plains Software Inc. ([www.greatplains.com](http://www.greatplains.com))
- Onyx Software Corp. ([www.onyx.com](http://www.onyx.com))
- Marimba Inc. ([www.marimba.com](http://www.marimba.com))
- Microsoft ([www.microsoft.com](http://www.microsoft.com))
- Portal Software Inc. ([www.portal.com](http://www.portal.com))
- Progress Software Corp. ([www.progress.com](http://www.progress.com))

Xevo describes Workbench as a platform for service management and one-touch provisioning automation. The main elements for this package, such as service definition and bundling, provisioning, metering, and call record mediation, are designed to serve as the head coordinator of all the provisioning processes executed by its partners' products. This level of coordination is necessary because ASPs are serving up specific services to specific users.

Unlike a Web site host, an ASP or a business-to-business extranet host needs to know and define end users and their access and application privileges. Once you start naming users, you must manage the provisioning of servers, load balancers, databases, operating systems, and network resources. That range of systems illustrates a fundamental shift in the definition of an "operating system."

Operating systems used to mean that we were talking to the hard drive and managing memory. Now, the network and OS will become a collection of components. If you are accessing applications, the entire operating infrastructure of the ASP must be managed. With each addition of a new application, server, switch, or end user, new information must be shared across subscriber management, provisioning, billing, and other OSSs and BSSs.

The ultimate goal is to build a unified system that automatically and dynamically builds and provisions a packaged service in response to customer clicks on a service portal icon even across multiple data centers and service sources. Standards-based interfaces are beginning to make that possible by allowing communications between provisioning systems and the applications.

## Pricing Models and Billing

ASPs can use various pricing models. There are different kinds of users and applications. Certain pricing models might be more appropriate for certain users/applications.

It is easy to measure time spent on a system; most operating systems will tell you how long a user has been logged on. However, OSs will not tell you how long the user has been on for a particular application. Part of the pricing model may be based on how often a customer accesses an application by PC or via a wireless device. Pricing is a matter of how a service provider packages an application, as well as the costs of delivering that service.

Usage-based billing is also receiving a lot of attention; however, companies are struggling over what to measure, and how to measure it. If a company decides to measure usage, it must measure packets and relate the number of packets to some level of utilization. There are many methods of measurement. The interesting

thing is, the way that the user data is gathered does not necessarily equate to how you present it back to the customer. If I say you use 75 units, how do I measure the units; is it the amount of bandwidth you use? The number of computer cycles you use? A formula that summarizes of all those? How much disk space you use?

It is not a big challenge if a company uses recognizable metrics such as computer usage or processor usage; these metrics can be gathered from the operating system itself. However, it becomes a challenge when a company tries to equate packets to usage. If I was using some type of graphic application that required me to move a lot of big graphics files back and forth between the ASP, shouldn't that be more expensive than if I was just sending text files back and forth? The usage thing is kind of a slippery slope, too, because you don't want to get so complicated that it turns people off. A majority of surveys indicates that most people will not pay for transactions that way.

In addition, if a thin-client mechanism being used in an ASP model does not allow the ASP to measure what packets go with what application the customer is using, pricing for a particular application is impossible.

Pricing by transaction is also gaining momentum. Still, defining a transaction and being able to capture the transactions for the billing system is no small task. Some applications could be open to pricing by the amount of data stored within; for example, the number of customers stored within an application for a dentist office.

Threshold pricing is another possible variation; for example, users pay a flat fee for usage up to a certain threshold. Beyond that, they would pay a small fee per unit (CPU cycles) used.

Probably the most common pricing model today is to charge a flat fee per month, often on a per-license/per-user basis. For the larger applications such as ERP software, some pricing occurs per seat/per license within the software (Table 7.6).

Many believe that it is only a matter of time before pricing moves away from that. However, there is the tendency to make pricing too complex; the incredible success that service providers have enjoyed over the last couple of years is not because of the complexity of the billing, but rather because of its simplicity. The answer is not to make every management system so complex and so terrifying for the end user that he or she will shy away from signing up for the service.

**Table 7.6** ASP Pricing Models

Pricing Model	Pricing Elements	Target Customer
Upfront Charges plus Monthly	This model uses upfront fees for integration, consulting, and/or customization. There is the option of transferring software and hardware assets to the customer after a defined period of time. Monthly fees are based on usage, number of users, or number of concurrent users.	This model is geared toward larger companies with more complex implementation needs.
Bundled Monthly Charges	This model uses upfront costs that are included in monthly payments.	This is pushed to smaller companies with minor integration and installation needs, or companies interested in conserving capital.
Flexible/Revenue-Sharing	This model uses upfront costs, but there are no fixed monthly costs. Ultimately, the pricing is driven by the total success of the site (including number of users, and transaction volume). There are either minor or no upfront costs. Total pricing is based on the transaction volume and ASP service performance.	This is used for Dotcoms that are running e-commerce or B2B commerce sites.

Important to pricing is user management. ASPs must be able to track users to bill for applications. An organization can have users who operate an application every day, or regular users who sign on once a month. Some may only use the application once a year. Should all of these pay the same amount to use the application?

Many experts say that the billing system needs to understand the user distinction. Such a metric defines the types of users and assesses whether a particular user has reached the predefined threshold for one of those types.

You need to have a hook into the application so that you know the users who essentially have been given logons to the application. If you want to be able

to measure a user's actual usage of the application, you would need, again, a hook into that specific application that will show you when the user logged on, when he or she logged off, or how a transaction is executed.

## Configuring & Implementing...

### The 95<sup>th</sup> Percentile

Service providers generally measure customer circuit usage by using a calculation method called the 95<sup>th</sup> percentile measurement. Through this method, two separate 95<sup>th</sup> percentile values are computed. The first is based on inbound data samples collected; the second is based on outbound data samples collected. For billing purposes, the higher of the inbound or outbound 95<sup>th</sup> percentile value will be used.

For companies that use multiple physical circuits, through a media such as Ethernet circuits, there are two values that the 95<sup>th</sup> percentile will use for billing purposes, using whichever is the larger value. The first method uses the sum of the inbound 95<sup>th</sup> percentile values; the second method uses the sum of the outbound 95<sup>th</sup> percentile values.

For traffic that flows into the routers from both the inbound and outbound directions, the 95<sup>th</sup> percentile value is calculated by using data usage samples that are collected over a given period of time (usually one month). This method uses statistics to compute a value that reflects a common point at which 5 percent of the samples are greater in value, and 95 percent of the samples are lesser in value than the 95<sup>th</sup> percentile value. To say it another way, the 95<sup>th</sup> percentile calculation establishes a "threshold of acceptance" value for customer billing purposes. This method essentially discards the peaks in customer usage and establishes a more representative circuit usage value, for billing purposes.

## Billing

As an ASP provider, you will face various billing issues that are likely to be among your greatest challenges. Regardless of what is offered, the billing systems must go through many changes before they can effectively meet your billing needs in this new ASP business model.

If you decide to bill based on the number of customers, you will need a way to measure how many customers are in the database and be able to feed that data into your billing system. If you want to bill based on the number of transactions, which is even harder, your system will need to capture the number of transactions, such as each time an ASP customer presses Enter. A company would need to determine what it wants to measure, figure out how to capture it, and feed that into the billing system.

Your billing application would have to be designed (or modified with an API) so that, for instance, every time a customer executes a predefined transaction, the application would then log or identify the user, thereby enabling the billing system to catch and sum the transaction in a usage log.

The software vendors must take their applications that now run in a standard client-server environment and change them to some extent to be able to run in an ASP environment. One of the modifications that they need to make is putting some hooks or some specific data feeds into their applications to allow you to bill based on some of these unique characteristics.

There are still miles to go before these application providers can offer their applications in a hosted environment. It is not just a matter of putting the applications on servers and giving people access to them.

Directory services are the way to manage an installation of numerous servers. Many applications, though, are not directory enabled.

It took Bell Laboratories the better part of 100 years to get telephone systems into a format that was reliable to handle millions of customers uninterrupted. Software as we know it is going to have to go through a massive transformation before the same can be said about software applications, especially in the ASP model.

## Managing Billing with Partners

Beyond tracking usage, ASPs are faced with the difficulty of tracking the financial relationships with their revenue partners.

An ASP must have relationships with many intermediaries to provide full service for the customer. If the intermediary collects the money, it would pass the money to the ASP that provides the service and bears the cost of delivering it. The ASP will pay a commission to the intermediary for that. On the other hand, the ASP can charge intermediaries for a service it provides, such as billing services or virtual customer service.



Here is an example of how one client handles revenue between partners: Distributors or value-added resellers (VARs) sell the ASP application. These distributors do not have any infrastructure to provide or bill for the application. At the end of the month, the ASP produces bills for the end customer, but instead of placing its logo on the bills, it inserts the VAR's logo, producing VAR-branded bills or invoices.

At the end of the month, the ASP also sends a bill to the VAR. The VAR collects money from the end customer. The ASP tells the VAR how much it owes, giving the VAR the total bill with a discount that figures in the VAR's commission.

The difficulty in this process is that the billing system needs to run another cycle, so to speak, with the same data: once the ASP finishes the bills for the customer, it must produce bills for the VAR, a phase that usually does not exist in typical billing systems.

This partner relationship is complicated because the ASP pays the VAR a commission based on the actual payment of the customer. The level of commission is different for each VAR according to what they negotiate, it is different for each application, and it is different depending on the seniority of the customers (how long they have been using the service). An ASP may tell the VAR that the first year the customer uses the service the VAR gets 5 percent, but the second year, it may get only 3 percent. Hopefully, this will encourage the VAR to bring in new customers. Agreements between the ASP and the ISV must be tracked as well. Clearly, a whole level of billing resides above the plane of the customer, wherein the ASP must bill and collect revenue from its partners.

Everybody in the value chain wants a piece of the success. The mechanisms that track, record, and then pay out on revenue sharing are not very good. Micropayments could change this. ASPs are going to move toward usage-based or transaction-based pricing, and micropayments will be needed for shifting small transaction-type revenues from customers as well as between partners.

There are arrangements, everything from a micropayments-type arrangement for every transaction that is less than a penny, sharing between the content provider and the content hoster. A content provider could license the use or license the access through a service provider, and just take a fee per month or per year and let the service provider offer it on a per-use basis. There are also other revenue-sharing models that are a combination thereof—a fee plus a micropayment. Moreover, there are arguments for providing it free, with the approach of sharing on advertising revenue.

## Summary

SLAs are a critical part of any service provider's business, yet how they are measured is as important as what is measured. By using a managed network service with end-to-end measurements, that are independent of the router or switches, ASPs are able to provide the best overall SLA measurement capability.

A next-generation OSS should enable work to flow electronically across a service provider's organization, from end to end, yet still provide visibility to business processes and resource utilization. It also should enable the service provider to manage the end-to-end service delivery process that often involves more than one type of transaction across the organization, as well as with other service or network providers. Most important, the OSS should be available in a solution that can eliminate the complexity of dealing with a variety of systems. If a service provider cannot achieve this goal due to complex OSS requirements, the provider should carefully select best-of-breed vendors offering proven, seamlessly integrated solutions.

Over the last decade, the telecommunications network has been in transition. The old network was primarily designed for switched-voice traffic and was relatively simple. It was based on copper loops for subscriber access and a network of telephone exchanges to process calls. This network is evolving into one designed for integrated access, transport, and switching of voice, high-speed data, and video. The network will be based on a variety of complex technologies. Because of its complexity, each network element technology is accompanied by an EMS that harnesses the power of the technology while masking its complexity.

As an ASP, you are going to need to address your customers' various billing requirements. While most end users are billed on a per-month/user basis, it is not always a viable alternative. Businesses are expecting to get their specific needs met and do not want to have an ASP tell them what options to choose.

As you can see, there are ways in which ASPs can bill their customers based on an activity-usage model (per transaction/order processed/e-mail sent), a usage-based model (time spent on an application, use of disk space, etc.), or a customer preference-based model (power vs. regular users, office vs. home use, etc.).

Remember that it is all about customer preference. In this competitive market, ASPs have to be able to differentiate themselves. One way to do this is through pricing schemes and customer support.

As with any other business, the sole concept behind the ASP model is revenue; and at the heart of revenue is billing. As the ASP industry continues to grow in complexity, however, determining who is billed for what, and then incorporating it into one billing system, can be a problem.

ASPs typically face similar problems when it comes to billing and managing an ever-growing service portfolio, developing tailored agreements to meet their customers' specific needs, and forming and managing profitable partnerships.

## Solutions Fast Track

### The Effect of Outsourcing

- ☑ The service level agreement (SLA) allows the customer to set minimum (and maximum) limits to be met. There are three main areas in almost every SLA: Planning, Verification, and Troubleshooting.
- ☑ Frame Relay involves a number of system parameters that go beyond the standard parameters that can be monitored by the Simple Network Management Protocol (SNMP). Some of these elements cover the entire network, segmented networks, or even single circuits. The level at which an SLA can be defined depends entirely on the business need of the circuit.

### What Service Levels Should the Service Provider Consider?

- ☑ Most clients will want you to commit to a monthly guarantee of at least 99.5 (more often, 99.999) percent uptime. This guarantee generally includes all of the devices that are within your infrastructure, that connect to the local loop, or connect to the CPE. An uptime of 99.5 percent equals 3.6 total hours of downtime per month per site.
- ☑ Many of the largest companies guarantee a delay (round-trip) no greater than 300 milliseconds. You may be able to provide guarantees based on access line speeds, which can offer much lower delays for T1 and 64 kbps.
- ☑ Some service providers base effective throughput on the percentage of delivered frames based on a Committed Interface Rate (CIR) or frames that are labeled discard eligible (DE). Other providers base this calculation on the committed burst size rather than the excess burst size. You may be able to exclude configurations where the destination port is not configured to handle the bandwidth of the CIR.

- ☑ Response time can be whatever number of hours that you and the client agree upon. There is a pretty standard method that says that you will respond within four hours of reported outage. This also depends on the location of the service provider from the maintenance center. Usually this maintenance only covers CPE, as your facility will be handled on an internal basis.

## The Realities of Customer Compensation

- ☑ Many of your customers will want to know if you can find and fix issues (and potential issues) before they are affected. They will also most likely want to know if you will proactively fix issues, or wait for them to call and inform you. They will also wonder if you have the resources to meet the demand of the time to resolution or repair that is included within their SLA. In the customer's mind, compensation for downtime is not the correct answer, nor will it ever be. They just want you to take care of them, so that they in turn can take care of their clients.
- ☑ What will your clients look for in these reports on SLAs? Here are some things that your clients will ask you to do:
  - Continually check that the WAN is capable of handling the services that they are providing.
  - Verify that service levels are being maintained. This request may require your ability to show monitoring in real time.
  - If services are not being met, then there must be an immediate path to resolution. This may be entirely your responsibility.
- ☑ Many tools are available to monitor the systems in the data center environment. These tools are generally used to collect usage statistics and the percentage of uptime for devices. These packages will also inform a centralized management station of the number of outages, the length of these outages, the mean time between failures (MTBF), and the mean time to repair (MTTR).
- ☑ By making your model more customer oriented, you can offer SLAs for things such as: emergency response, response time guarantees, call center availability, and remote troubleshooting.

- ☑ As the corporate infrastructure has evolved, so have the dynamics of the corporate network. What you are more apt to find in these changing times is an internal staff that handles and maintains very little of the overall network, remaining entirely within their walls or boundaries. External staff is comprised of the outsourced applications and infrastructure support. When you combine these two teams, you can encompass the range of support, including intranet-based Enterprise Resource Planning (ERP), electronic mail (e-mail), messaging, scheduling, desktop support, operating systems, remote access, security, and other miscellaneous company needs.

## How Service Providers Have Responded

- ☑ With all of the mission-critical applications that are available, many service providers are now offering services that are more advanced than the typical “leased line” connectivity that had been their bread and butter for so long. Leased lines were the lifelines to companies that needed direct access to their sites, and to their applications.

## The Operation Support System Model

- ☑ The Operations Support System (OSS) model usually refers to a system (or systems) that can perform the management necessary to maintain and monitor your SLA requirements. This model takes the following items into account: performance management, inventory control, system engineering, design, and support.
- ☑ In order to truly understand OSSs, you must first become familiar with some of the fundamental systems that are involved. These systems handle the functions of ordering, service fulfillment (such as voice, data, and other IP-based services), inventory, circuit provisioning, and activation.
- ☑ Many of today’s OSS solutions are considered commercial off-the-shelf (COTS) packages. These applications are able to offer some out-of-the-box utilities and are intended to be modified to meet customer needs. This customization could allow your company to integrate management capabilities and enable your customers to take advantage of your services, thus adding efficiency.

## Broadband Access Changes the Market

- ☑ Broadband access has changed the way we do business, and how we live at home. At this moment in time, DSL and cable are surpassing every other method of access across the United States. This isn't to say that Frame or other connections are going to disappear; it is really saying that, like everything else, things change.
- ☑ Many of today's service providers are struggling with the deployment of these technologies. It's not because they don't have the bandwidth; it's because it is difficult to maintain and upgrade your infrastructure if you are unable to see your current copper allocation (for the local loop) and resource availability. One of the ways that a central office (CO) can handle these issues is to have an up-to-date, dynamic inventory of provisioning.
- ☑ In order for a service provider to incorporate DSL within its infrastructure, there is the need to integrate two components: a splitter and a DSL Access Multiplexer (DSLAM). A splitter distributes voice traffic to the Plain Old Telephone System (POTS) cloud, and data traffic to the DSLAM. A DSLAM is able to communicate with the DSL router that is located on the customer's premises.

## Quality of Service

- ☑ Quality of Service (QoS) is a measurement of the service value. Measurement of QoS is very subjective; it depends on the technology on which it is implemented to see if there are acceptable levels of performance.
- ☑ You will need to maintain a high level of QoS to maintain and attract new customers. Therefore, you should implement and manage your solution so that it is capable of meeting your customers' expectations. QoS will vary from customer to customer, so tailor your SLAs to reflect client needs; for example, a bank that may need to implement high-speed transport (ATM) and VPNs.

## Management Systems for Your ASP

- ☑ Many of today's service providers use (at least at some level) the Telecommunications Management Network (TMN) model. The TMN model provides the outline for attaining interconnectivity and communications across diverse platforms and environments.
- ☑ TMN was developed by the International Telecommunications Union (ITU) as a tool to help support, manage, and deploy services. TMN was originally based on the common management information service element (CMISE).
- ☑ The TMN model outlines what is necessary to make your network infrastructure flexible, scalable, manageable, and highly available. TMN defines standard ways of handling management tasks and communications across networks. TMN allows you to distribute the appropriate levels for growth, efficiency, and communication performance.

## What Tools Do You Need to Automate TMN?

- ☑ A multitude of tools are available to automate the task of building TMN agent or manager applications. You can deploy and tailor the TMN agent and manager toolkits to match your company's GDMO/ASN.1 MIB representations. These products should have the following features in order to take advantage of the TMN model and to most productively support a TMN infrastructure: automated prototyping, conformance to all TMN standards, dynamic information modeling, Management Information Base (MIB), platform-independent interfaces and tools, Q adaption capability or compatibility, and system management functions (SMFs).

## The ASP Transformation

- ☑ To transform from an ISP to an ASP, you will need a service management solution that is designed specifically to manage the unique functions and processes of ASPs with carrier-class reliability and scalability.
- ☑ The ultimate goal is to build a unified system that automatically and dynamically builds and provisions a packaged service in response to customer clicks on a service portal icon even across multiple data centers and service sources. Standards-based interfaces are beginning to make

that possible by allowing communications between provisioning systems and the applications.

## Pricing Models and Billing

- ☑ Usage-based billing is receiving a lot of attention; however, companies are struggling over what to measure, and how to measure it. If a company decides to measure usage, it must measure packets and relate the number of packets to some level of utilization. There are many methods of measurement. The interesting thing is, the way that the user data is gathered does not necessarily equate to how you present it back to the customer. If I say you use 75 units, how do I measure the units; is it the amount of bandwidth you use? The number of computer cycles you use? A formula that summarizes of all those? How much disk space you use?
- ☑ Pricing by transaction is gaining momentum. Still, defining a transaction and being able to capture the transactions for the billing system is no small task. Some applications could be open to pricing by the amount of data stored within; for example, the number of customers stored within an application for a dentist office.
- ☑ Threshold pricing is another possible variation; for example, users pay a flat fee for usage up to a certain threshold. Beyond that, they would pay a small fee per unit (CPU cycles) used.
- ☑ The most common pricing model today is to charge a flat fee per month, often on a per-license/per-user basis. For the larger applications such as ERP software, some pricing occurs per seat/per license within the software.
- ☑ As an ASP provider, you will face various billing issues that are likely to be among your greatest challenges. Regardless of what is offered, the billing systems must go through many changes before they can effectively meet your billing needs in this new ASP business model.
- ☑ Directory services are the way to manage an installation of numerous servers. Many applications, though, are not directory enabled. It took Bell Laboratories the better part of 100 years to get telephone systems into a format that was reliable to handle millions of customers uninterrupted. Software as we know it is going to have to go through a massive



transformation before the same can be said about software applications, especially in the ASP model.

## Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the “Ask the Author” form.

**Q:** How many Q interface types are there, and what are their functions?

**A:** There are two classes of Q interfaces: Q3 and Qx.

- **Q3 Interface** The Q3 interface is the lifeline to the operations system. Q3 is the only interface that QAs, MDs, or NEs may use to communicate directly with the OS. If a QA or NE does not use the Q3 interface, it cannot communicate directly with the OS; instead, it must communicate via an MD.
- **Qx Interface** The Qx interface always operates with a MD. It never takes the place of a Q3 interface. The MD can negotiate between local management information provided by a Qx interface and the OS information provided by a Q3 interface.

**Q:** Where can I find more information on billing?

**A:** More information is available at [ASPIndustry.org](http://ASPIndustry.org). This Web site contains up-to-date information that can help you understand what is happening in the market.

## Designing the Infrastructure

### Solutions in this chapter:

- Design Considerations
- Site Considerations
- Designing with the Hierarchy in Mind
- Frame Relay Internetwork Design Considerations
- Capacity Planning for Your Infrastructure
- Protocol Planning Concerns
- Addressing Considerations
- Application and Network Services
- Application-Aware Networking
- Scalability Considerations
- Multimedia Services
- Planning for the Future Growth of Your Company's Infrastructure
- High-Availability Design
- ☑ Summary
- ☑ Solutions Fast Track
- ☑ Frequently Asked Questions

## Introduction

You have decided that becoming an application service provider (ASP) is the right thing to do, so how do you go about setting it up? So far, the book has focused on the components of creating an ASP network infrastructure. This chapter will help you pull it all together.

In this chapter, we look at what it takes to design the infrastructure. This is what separates the network design and implementation engineers from the network administrators. If you design the network properly, there won't be peals of joy and adulation. However, if you design your infrastructure poorly, your company will not be able to retain its current customers and attract new ones. Because many of you already have a viable network infrastructure, I wrote this chapter as a high-level overview of some of the considerations that you may face. I ask you to forgive any dated material; by the time this book is finished and sent to press, there is bound to be the next "latest and greatest" toy available on the market.

So, why do we do this? Designing the network is a thankless job, and it is probably one of the most difficult and most overlooked areas in networking; it is also essential to the success of your business. I think that most of us do this because we enjoy the challenge, the thrill and exhilaration that comes from solving the impossible...or because we like to play with toys. Even with all of the new toys, though, there is no such thing as a perfect network, but here is how you go about trying to achieve one.

### Designing & Planning...

#### **End-to-End Network Services for a Data Center Environment**

Some of the challenges associated with creating an end-to-end ASP solution are protecting mission-critical programs and applications, creating security policies, being able to manage the environment, provide high availability, support for multimedia content (voice and video) and make the solution so that it scalable.

Always keep these issues in mind when applying for funding to build or expand your network. Remember that by addressing these concerns, you will add value to the company, so stress that whenever possible.

There is a great deal of detail with site preparation alone. In this chapter, there are some best practices on design and location of equipment. These are general guidelines and not a complete list, so if you are setting up a major data center, the designs in this chapter may not be large enough examples, as they are scaled to medium and large ASPs. This chapter mostly deals with Cisco Systems equipment and Microsoft Windows 2000 environments, but I do not mean to imply that these are the only vendors that you can use.

## Design Considerations

There is no such thing as a perfect deployment of an ASP. Some things will make your company successful, and some will inhibit your growth. Remember that you should keep the following issues in mind in order to deploy a dynamic, interactive, application solution while providing access to various dissimilar clients.

Your infrastructure must address the following issues:

- It must provide an interoperable and distributed network environment that is compatible with Internet standards.
- Your application must not be location dependent.
- You must provide a framework that will incorporate disparate network data and application types from your clients and internal application developers.
- It must provide centralized management and security functions.
- It must be scaleable as usage grows.

Other items that you need to address in the implementation of a packet-switching infrastructure include:

- Hierarchical internetwork design
- Topology design
- Broadcast issues
- Performance issues

**NOTE**

Part of this chapter focuses on general packet-switching considerations and Frame Relay internetworks. I selected Frame Relay because it presents a broad picture of design considerations for interconnection to packet-switching devices.

## Getting Started: The Design Process

An internetwork requires many layers of thought and design that encompasses everything from physical space to future network considerations. There are generally three components when designing a large internetwork: data center networks, wide area networks (WAN), and remote users (in this case, your external clients).

- **Data center networks** are generally comprised of locally housed equipment that will service your clients from a building, or set of buildings.
- **Wide area networks** are the connections between the data center and the customer.
- **Remote users** are your clients and telecommuter traffic that are subsets of your main clients.

Designing the network is a challenging task, but as I said earlier, you probably are doing this job because you like a challenge. You must take into account that each of the three components has its own distinct requirements. For example, an internetwork that is comprised of five meshed platforms can create all sorts of unpredictable problems, so attempting to create an even larger series of intermeshed networks that connect multiple customers who have their own network issues can be downright mind-boggling.

This is an age in which equipment is getting faster, sometimes by being more granular in the services that are offered, and other times, allowing more to be done within a single chassis. Infrastructure design is becoming more difficult due to ASPs moving toward more sophisticated environments that use multiple protocols, multiple media types, and allowing connections to domains outside your areas of influence because of customer requirements.

One of the greatest trade-offs in linking local area networks (LANs) and WANs into a packet-switching data network (PSDN) infrastructure is between

cost and performance. The ideal solution would optimize packet-based services, yet this optimization should not be interpreted to picking the mix of services that would represent the lowest possible tolls. Your customers are going to want speed, availability, and ease of use. To successfully implement a packet-service infrastructure, you should adhere to two basic rules:

- When designing and implementing a packet-switching solution, try to balance cost savings with your company's internal performance requirements and promises to its customers.
- Try to build a manageable solution that can scale when more links and services are required by your company's clientele.

## Data Center, WAN, and Remote Links Defined

The *data center* is a building or set of buildings that house the infrastructure of your network. Most data centers consist of many platforms that are joined together to form the backbone. These centers usually use LAN technologies such as Fast Ethernet, Gigabit Ethernet, and Asynchronous Transfer Mode (ATM) to power their network.

The *WAN* is the way to connect your multiple customers and buildings across geographically dispersed areas. When you use WAN technologies (such as fiber, satellite, or copper) to connect your customers, you may have to pay for bandwidth from an external telecommunications provider.

*Remote links* (users) are those clients who work on a per-use agreement, or their internal resources use the applications from disparate areas, such as a home office, and need to access the network in various ways. Some of the more common ways are over the WAN through a virtual private network (VPN) connection and through remote access servers (RAS).

### NOTE

There are several types of VPN solutions available. Depending on your needs and likes, multiple companies can provide clients for your type of network. Be sure to research whether these clients will support the types of traffic that will be running over these connections (for example AppleTalk, IPX/SPX, and in rare instances, SNA bridging). VPNs also allow for security by allowing encryption and authentication services.

## The Design Process—Getting Down to Business

Most design processes start with a good conceptual drawing, so let's begin there—with a pencil and paper (or if you are more inclined, Visio or any other publishing software with which you are comfortable. I like crayons, but they don't taste the same as the wrapper says they do) and some brainstorming. Try to start mapping out the proposed layout of the infrastructure. This is a multiple-step process, and these designs may never be fully finished, as they should be updated as the network grows and changes. Remember that one of the most important features of a network is scalability.

### NOTE

---

A conceptual drawing is a reference tool that can later be used as documentation. It will help when you bring in other people to support your infrastructure, and will assist in the explanation of how and why your network functions the way it does. It will also help in future growth areas, as you can see any potential problems that may crop up. A map of your infrastructure will also help in the training of your network team, so you may not have to take that weekend call while you are on vacation when someone decides to “modify” or “streamline” the network. Being able to create a good conceptual drawing is a good skill to have for our highly competitive market. A well documented network should always have the most current and correct details, such as IP addresses. This is critical, not just a good idea.

---

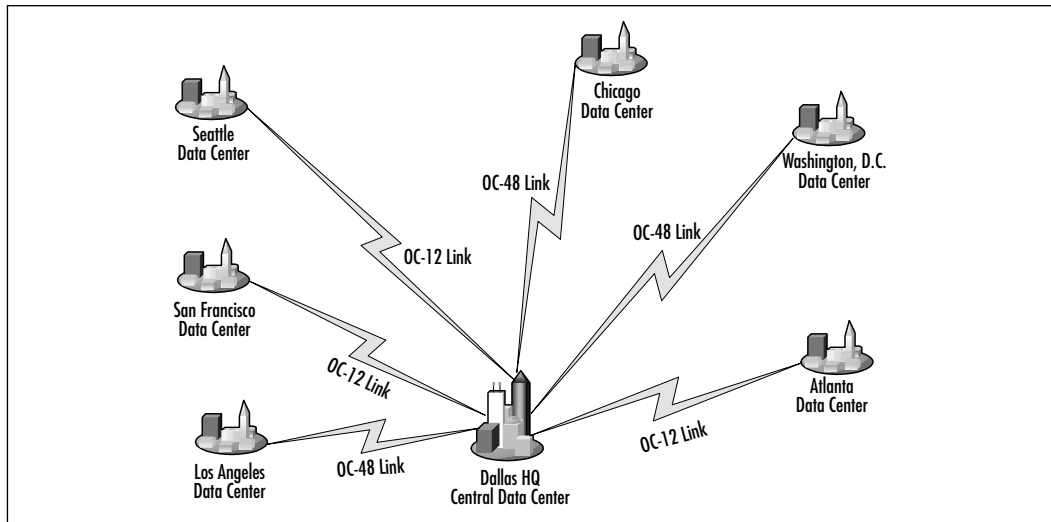
The first step of this process is diagramming the “design” from the proverbial “30,000-foot” view. Depending on the situation, the first drawing should be just the physical locations and possible future locations of your buildings if designing a dispersed data center or WAN model. This will help you provision appropriate WAN connections, as well as future circuit provisions (Figure 8.1).

### NOTE

---

Remember that provisioning bandwidth is an important issue and should be covered in detail. The design in Figure 8.1 may not be the most cost effective, but may be necessary due to network requirements. We discuss that later in this chapter.

---

**Figure 8.1** The 30,000-Foot View

## Site Considerations

Now that we have a conceptual drawing, we can delve into site considerations. This might be an area over which you have very little control, yet it has a high impact on your design. The following sections provide the basis of how to implement your design, while considering physical limitations.

## Physical Equipment Space

The next step in the design process should continue with an analysis of the physical space. The point here is that if there is insufficient space and/or resources (such as power, air-conditioning, etc.), there will be severe limitations on design and implementation. When you are building a new physical plant for your ASP, make sure that there is adequate space available and sufficient resources (power and cabling, as well as security) to suit your needs.

Here are some considerations for the physical plant implementation:

- When you can, scope out the physical layout of the room; there should be sufficient space to place the equipment to be installed.
- Find out whether the room is “fireproof” with typical methods to contain a fire, and if the room is climate controlled.
- Find out how secure the room is (combination lock, keyed, keyless entry, retinal scan deoxyribonucleic acid (DNA) analysis, and so forth).



- Decide on what types of floor racks and/or cabinets for build-out purposes (19-inch network racks, 23-inch telecommunications racks).
- Make sure to pay attention to weight, power, and environment concerns.
  - **Weight** This category will vary according to equipment and cabling. Try to scope out what type of equipment will go into each rack, and where it will be placed on the floor and within the rack itself. Depending on your wiring scheme, there may be an uneven distribution of weight, which may unbalance the racks you are installing. If at all possible, get the weight requirements from the vendor, and find a rack that is able to support what you are placing, plus approximately 20 percent of that weight for future equipment and swap out. In addition, take into account the flooring, and whether it can support the rack.
  - **Power** This is usually an easy piece to scope out. Based on the vendor specifications for the equipment, you should ask the following questions. Does the room have a dedicated power source? Does that source support the needs of the proposed equipment? Is there a backup power supply (uninterruptible power supply (UPS))? Does the room have a circuit breaker and a master shutoff switch? These are important questions for the present and future expandability. Note that if there are any major questions, call an electrician.
  - **Environment** This is always a little tricky, as this is the area with the most variance. You could write an entire book on this topic, but it probably would only be for civil engineers and architects. Some of the questions that you should ask when assessing the environment for the infrastructure include:
    - How is the room secured?
    - How are the climate control and ventilation handled (heat, ventilation, and air conditioning (HVAC))?
    - Is the room sealed, and is it set up for fire control?
    - Does the flooring prevent static buildup?
    - How is the cabling plant implemented?
    - Should you use ladder racks or raised flooring?

Let us discuss the security issue first; as there is usually a high cost associated with the equipment installed, you will want the room theft-proof. The other issue to security is that you don't want people to just come in and change the equipment or install things on a whim, so security is also able to provide restricted access to authorized personnel only.

Climate control topics can include many variables, such as temperature, humidity, and so forth. These factors are important; if the room is too hot, the equipment will not function properly if it overheats, and if it's too cold, the possibility of freezing and cracking equipment exists. If the room is too humid, the equipment could short due to water in the equipment; if it is too dry, you run the risk of static electricity shorting the equipment. Recall that the equipment will emit its own heat, so even though the room may feel chilled while the equipment is off, once the equipment is running, the temperature will average out. Make sure to check the heat ratings of the equipment before designing your facility and rack layout.

Fireproofing is more straightforward. For a controlled room, the walls should go all the way to the roof to create an enclosed/sealed environment. A main reason for a sealed environment is a Halon dispensation system or other viable fire deterrent system that smothers the fire by displacing the oxygen in the room. (Technically, this can also be used for security, but I am not sure that killing someone is the way to go.)

Prevent static buildup by installing antistatic tile flooring. Under no circumstances should you use carpeting, as this will actually increase the static buildup that is natural with movement (and don't wear corduroys; not only do they make bad fashion sense, they definitely add to your ability to create static). If you are building your infrastructure from scratch, you may want to consider installing raised flooring, as this will also help with cable management.

For cable plant implementation and management, you will want to hire a company that can install and certify their work with a registered communication distribution designer (RCDD). This will make insuring the room and equipment somewhat easier. In addition, they can usually come up with a best practice for you and your location.

## Configuring & Implementing...

### General Rules for the Care and Feeding of Networks

- **Weight** This is pretty self-explanatory; if the equipment weighs more than the rack can support, then your equipment will fall down, go boom.
- **Power** Without sufficient power, all you have is an expensive paperweight.
- **Environment** If the equipment is not properly raised, it goes bad, and a psychologist will say that it is the product of its environment. What this really means is that you need a clean and controlled environment to get the most out of your equipment.

## Network Equipment Basics

As we will discuss network hardware, and on which layers that it functions, I should probably give a quick overview of types of equipment so that you have a general understanding of some of the following terminology:

- A **hub** (sometimes referred to as a *concentrator* (I call it a Davidson Special)) is a Layer 1 network device that is used to connect multiple users to a single physical device that connects to the network. A hub also acts as a repeater, as it can regenerate a signal as it passes traffic.
- A **bridge** is a Layer 2 network device that segments the network within the same network and is independent of higher-layer protocols. Bridges generally will have fewer ports than switches and hubs have. A bridge is primarily designed to help separate collision domains, which will improve bandwidth utilization because you can cut down on retransmissions.
- A **switch** is a Layer 2 network device that provides network ports and separate collision domains. A switch is generally used in place of a hub in network designs, because it can connect many users to the network. Lots of devices are called switches, but switches offer higher speed, as they do not share bandwidth and broadcast domains as hubs do. A switch can be considered a bridge with many ports.

- **Routers** are Layer 3 network devices that connect separate networks and pass traffic between subnets. Routers are used to link disparate media types as well, thus expanding the scope of the network. Routers are also protocol dependent, so if you are running anything other than the standard stuff (IP for example), you may want to consult your vendor.

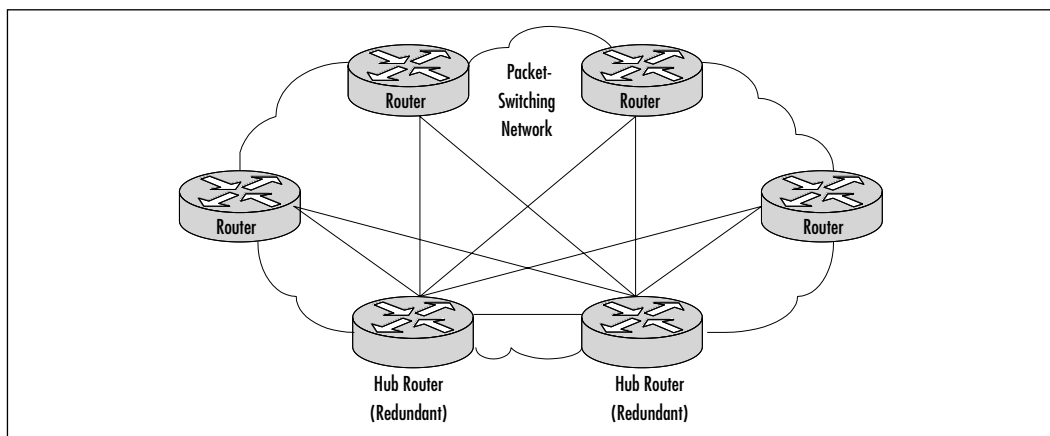
Since most network designs center on switches and routers, we will focus on their roles (later in this chapter) within the network.

## Designing with the Hierarchy in Mind

One of the most beneficial tasks that you can perform in the design of your network is to create a hierarchical internetwork design that will modularize the elements of a large internetwork into layers of internetworking. The key layers that will help you create these modules in this model are the Access, Distribution, and Core routing layers.

The hierarchical approach attempts to split networks into subnetworks, so that traffic and nodes can be more readily managed. Hierarchical designs assist in the scaling of internetworks, because new subnetworks and technologies can be integrated into the infrastructure without disrupting the existing backbone. This also makes the swapping out of equipment and upgrades much easier, because it is a modular environment. Figure 8.2 illustrates the basic approach to hierarchical design.

**Figure 8.2** Hierarchical Packet-Switched Interconnection



Some advantages of a hierarchical approach include:

- Inherent scalability
- Easier to manage
- Allows for the optimization of broadcast and multicast control traffic

Note that this three-tier model is defined by the Core, Distribution, and Access layers:

- The **Core layer** is where the backbone of the network is located and is the central point that data must traverse. This area should be designed for speed. The most important aspect of this layer is to pass information to the rest of the network. The core should have a meshed, redundant design for higher efficiency.
- The **Distribution layer** is where your border routers are located. Most of the routing decisions should be made at this level. This is the area where you would implement policies for the network.
- The **Access layer** is the customer's network. This area may allow you the least control because differing media and protocols may be used. This is usually the most over-subscribed part of the network.

## Scalability of Hierarchical Internetworks

Hierarchical internetworks are more scalable, because they allow you to grow your internetwork in a gradual way with the implementation of modules. This allows an infrastructure to grow in increments without running into the limitations that are normally associated with flat, nonhierarchical infrastructures. The drawback is that hierarchical internetworks require careful planning and implementation. There are many issues to consider when designing your network, including:

- The costs that are included in virtual circuits
- The complexity that will be inherent in a hierarchical design (particularly when integrated with a meshed topology)
- The need for additional hardware interfaces, which are necessary to separate the layers within your hierarchy
- The scalability of the software and routing protocols

To fully utilize a hierarchical design, you should create your hierarchy with your regional topologies in mind. Remember that the specifics of the design will depend on the services you implement, as well as your requirements for fault tolerance, cost, and overall performance. Always think, “How can I get the most out of this design, and what are the potential problems that could arise?”

## Manageability of Hierarchical Internetworks

There are management advantages inherent to hierarchical designs, such as the flexibility of your design, the ease of installing these modular segments into your network, and the management of fewer peers to your main convergence points.

- **Design flexibility** Designs that use the hierarchical approach will provide greater flexibility in the use of WAN circuit services. Leased lines can be implemented in the Core, Distribution, and Access layers of the internetwork.
- **Internetwork ease** By adopting a hierarchical design, you will reduce the overall complexity of an internetwork by being able to separate the components into smaller units. This will make troubleshooting easier, and provide protection against broadcast storms, routing loops, or other potential problems.
- **Hardware management** The complexity of individual router and switch configurations is greatly reduced, because each router has fewer peers with which they need to communicate.

## Optimization of Broadcast and Multicast Control Traffic

The effect of broadcast traffic in your internetworks (discussed in “Broadcast Issues” later in this chapter) requires that you implement smaller groups of routers and switches, which will make your network more efficient. I know that sounds odd, because you will probably need to add equipment, but you will see much more efficient usage of your resources.

A typical example of broadcast traffic is the routing updates that are broadcast between routers on a PSDN. A large amount of routers in any area or layer of the internetwork may result in bottlenecks because of broadcast replication. With a hierarchical scheme implemented, you can limit the level of broadcasting between these areas and from coming into your core. This requires hierarchical addressing

such as variable length subnet mask (VLSM), classless inter-domain routing (CIDR), and a routing protocol that can support these methods.

## Possible Types of Topology Design

Once you have established your internetwork scheme, you must design a way for handling interconnections among sites within the same region or area of administrative control. In designing regional WANs, whether you are using packet-switching services or point-to-point interconnections, three basic design approaches are common throughout the industry:

- Star topologies
- Fully meshed topologies
- Partially meshed topologies

In the following pages, I will try to help you understand these topologies and how you can use them to your advantage. Remember, though, that the discussions presented in this chapter address the application of these topologies specifically to packet-switching services.

### NOTE

---

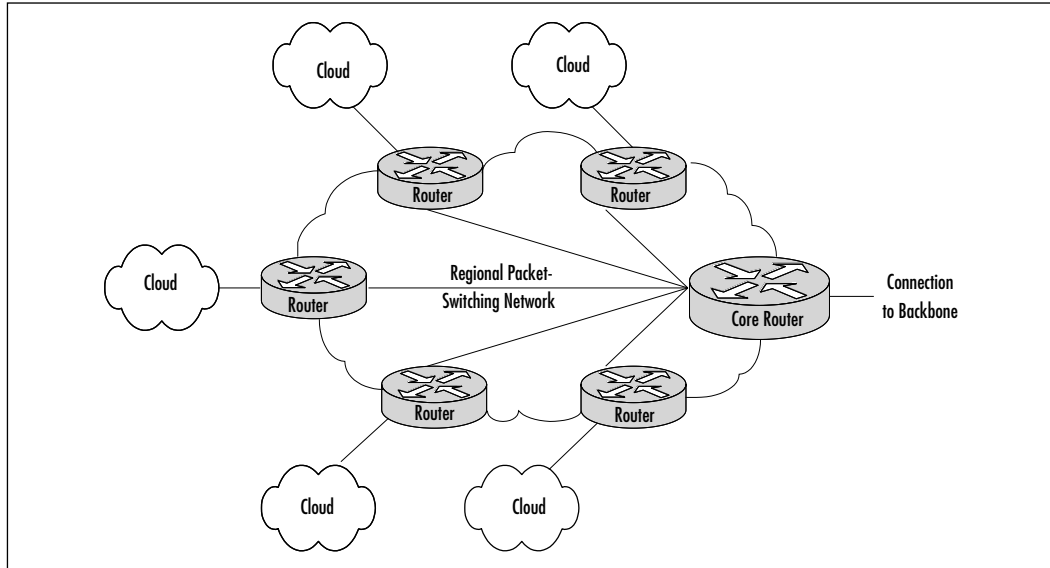
Illustrations in this chapter use lines to show the connection of specific routers on the PSDN network. These connections are considered virtual connections, as the circuits are mapped within the routers themselves. Normally, all physical connections are generally made to switches within the PSDN. Unless otherwise specified, the connecting lines represent virtual connections within the PSDN.

---

## Star Topologies

The star topology (also known as a hub and spoke) is a grouping of network devices that has a single internetworking hub, and provides connections for the external cloud networks to the backbone and access to each other, although only through the core router. Figure 8.3 illustrates a packet-switched star topology for a regional internetwork.

One of the main advantages of a star topology is that there is simplified management and minimized tariff costs or tolls. Whereas tolls aren't much of a factor

**Figure 8.3** Star Topologies for a Regional Internetwork

these days, they were in the past, and with some of the things that are happening politically, they could be again. However, there are significant disadvantages.

First, the core router is a single point of failure for the entire internetwork. Second, the core router can limit overall performance for access to backbone resources; the core may not be robust enough, or have enough bandwidth to handle all of the traffic from the external networks. Third, this topology is not very scalable, as there are generally only a certain number of ports on the core router.

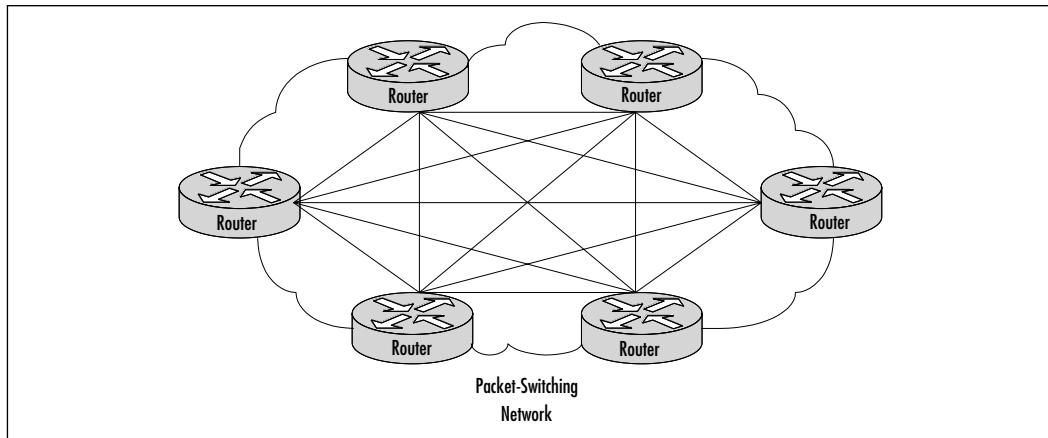
## Fully Meshed Topologies

In a fully meshed topology, each routing node on the edge of a packet-switching network has a direct path to every other node on the cloud. Figure 8.4 illustrates this type of arrangement.

One of the best reasons for creating a fully meshed environment is that it provides for a high level of redundancy. A fully meshed topology helps to facilitate the support of all routing protocols, but it is not tenable in large packet-switched internetworks.

Some of the main issues are due to the large number of virtual circuits that are required (one for every connection between routers). There are also problems associated with the large number of packet and broadcast replications necessary for routing protocols or application traffic, and the configuration complexity for routers in the absence of multicast support in nonbroadcast environments.



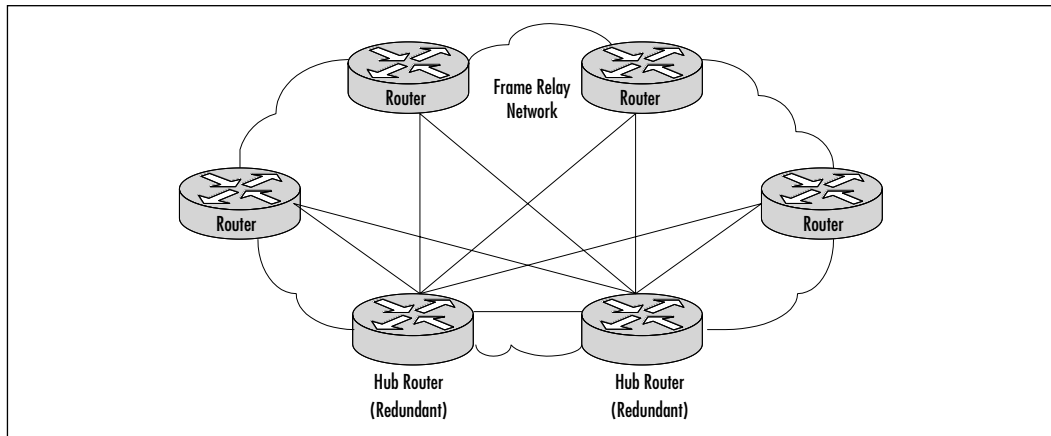
**Figure 8.4** Fully Meshed Topology

There is a middle ground, though; by combining fully meshed and star topologies into a partially meshed environment, you can improve fault tolerance without encountering the performance and management problems that are normally associated with a fully meshed internetwork. The following section discusses the partially meshed topology.

## Partially Meshed Topologies

As discussed earlier, a partially meshed topology reduces several of the problems that are inherent in star and fully meshed topologies. There is a reduction in the number of routers within a region that need direct connections to all other nodes in the region. Not all nodes need to be connected to all other nodes. For a non-meshed node to communicate with another nonmeshed node, it will send traffic through one of the hub routers. This is a lot like the star topology, but there is a redundant path available in the event the hub router becomes inoperable. Figure 8.5 illustrates such a situation.

There are many forms of partially meshed topologies. Generally, partially meshed implementations provide the optimum balance for regional topologies in terms of the number of virtual circuits, their ability to provide redundancy, and their overall performance. By providing a greater amount of connectivity and high availability, you will be able to use your bandwidth more effectively.

**Figure 8.5** Partially Meshed Topology

## Broadcast Issues

Broadcast traffic presents problems when it is introduced into a packet-service environment. Broadcasts are necessary for a node to reach multiple other node stations with a single packet when the sending node does not know the address of the intended recipient, or when the routing protocols need to send hello packets and other miscellaneous services.

As an example, the level of broadcast traffic that is generated in an Enhanced IGRP environment depends on the setting of the enhanced IGRP hello-timer interval. The size of the internetwork determines other issues. In a small network, the amount of broadcast traffic generated by Enhanced IGRP nodes might be higher than with comparable internal gateway routing protocols that run on the Internet.

However, for large-scale internetworks, Enhanced IGRP nodes generate substantially less broadcast traffic than RIP-based nodes, for example.

### NOTE

Usually, it is a good practice to manage packet replication when going over your design considerations. When integrating broadcast-type LANs (such as Ethernet) with nonbroadcast packet services (such as X.25), you should try to figure out where replication will cause bottlenecks within your network. With the multiple virtual circuits that are characteristic of connections to packet-switched environments, routers need to replicate broadcasts for each virtual circuit on a given physical link.

Within a highly meshed environment, the replicating broadcasts can be resource intensive in terms of increased required bandwidth and number of CPU cycles. Because of this, highly meshed networks are impractical for large packet-switching networks. However, circuit meshing is essential to enable fault tolerance. You really need to balance the trade-offs in performance with requirements for redundancy. Also remember that as you scale your network, there will be other issues that fall within the same vein; as you add routing nodes, you will want to add redundancy, which will add at least two paths to your core infrastructure.

## Performance Issues

When designing your WAN around a specific application service type, you should consider the characteristics of the virtual circuit. Sometimes the performance of a virtual circuit will depend on its capability to handle mixed-protocol traffic. Depending on how the traffic is queued and streamed from one node to the next, certain applications may require special handling. One solution might be to assign specific circuits to specific application and protocol types.

There are always going to be performance concerns for specific packet-switching services. That is why there is the ability to include Committed Information Rates (CIR) in Frame Relay internetworks and window size limitations in X.25 networks. (The CIR matches the maximum average rate per connection for a period of time.)

What is highly common within the ISP market is to sell guaranteed CIRs to your customers and give them the ability to “burst” (for a fee, of course) outside of the limits that they were given. As it is, a CIR is the minimum amount of bandwidth that your client is guaranteed at any point in time. The CIR is usually defined within the service level agreement (SLA) to which you and the customer agreed.

## Frame Relay Internetwork Design Considerations

A major concern when designing a Frame Relay implementation is scalability. As the number of remote clients and their links grows, your network must be able to grow to accommodate these growth spurts. The network must also provide a high level of performance, yet minimize support and management requirements.

Meeting all these objectives can be quite a feat. The following sections focus on some of the critical factors for Frame Relay internetworks, such as:

- Hierarchical design
- Regional topologies
- Broadcast issues
- Performance issues

The following are suggestions to provide a solid foundation for constructing scalable networks that can balance performance, fault tolerance, and cost. Again, I am only using Frame Relay as a template; it is not the only technology that you can use.

## Hierarchical Design for Frame Relay Internetworks

As discussed earlier in this chapter, the arguments supporting hierarchical design for packet-switching networks apply to hierarchical design for Frame Relay networks. Remember the three factors that lead us to recommend the implementation of a hierarchical design:

- Scalability
- Manageability
- Optimization of broadcast and multicast control traffic

One of the ways in which many Frame Relay vendors charge for services is by data link connection identifier (DLCI) numbers. These DLCI numbers identify a Frame Relay permanent virtual connection (PVC, also known as a permanent virtual circuit, which is the X.25 terminology). The DLCI number is locally significant, and defines the connection between Frame Relay elements. The number of Frame Relay PVCs within the network is highly dependent on what protocols are in use and actual traffic patterns.

To figure out how many DLCIs are going to be used within your environment, and how many will be mapped to your interfaces, depends on several factors that should be considered together:

- **What protocols are being routed?** Any protocol that is broadcast intensive will constrain the number of assignable DLCIs. For example, AppleTalk is a routed protocol characterized by high levels of broadcast overhead. Another example is Novell Internetwork Packet eXchange (IPX), which sends both routing and service updates, which results in

higher broadcast bandwidth overhead. In contrast, IGRP is less broadcast intensive, because it will send routing updates less often (by default, every 90 seconds). You can modify the timer for IGRP, so it can become broadcast intensive if timers are modified to send updates more frequently.

- **What are the levels of broadcast traffic?** Broadcasts, such as routing updates, are one of the most important considerations when determining the number of DLCIs that can be defined. The amount and type of broadcast traffic will be a factor in your ability to assign DLCIs within this general recommended range.
- **What is the speed of the connection?** Broadcast traffic levels are expected to be high, so you should consider faster links and DLCIs with higher CIR and higher burstable limits. You should also consider implementing fewer DLCIs.
- **Are there any static routes?** If static routing is implemented, you can use a greater amount of DLCIs per line, because a larger number of DLCIs will help to reduce the level of broadcasting.

To assist in your design considerations, here are two forms of hierarchical design that you can implement:

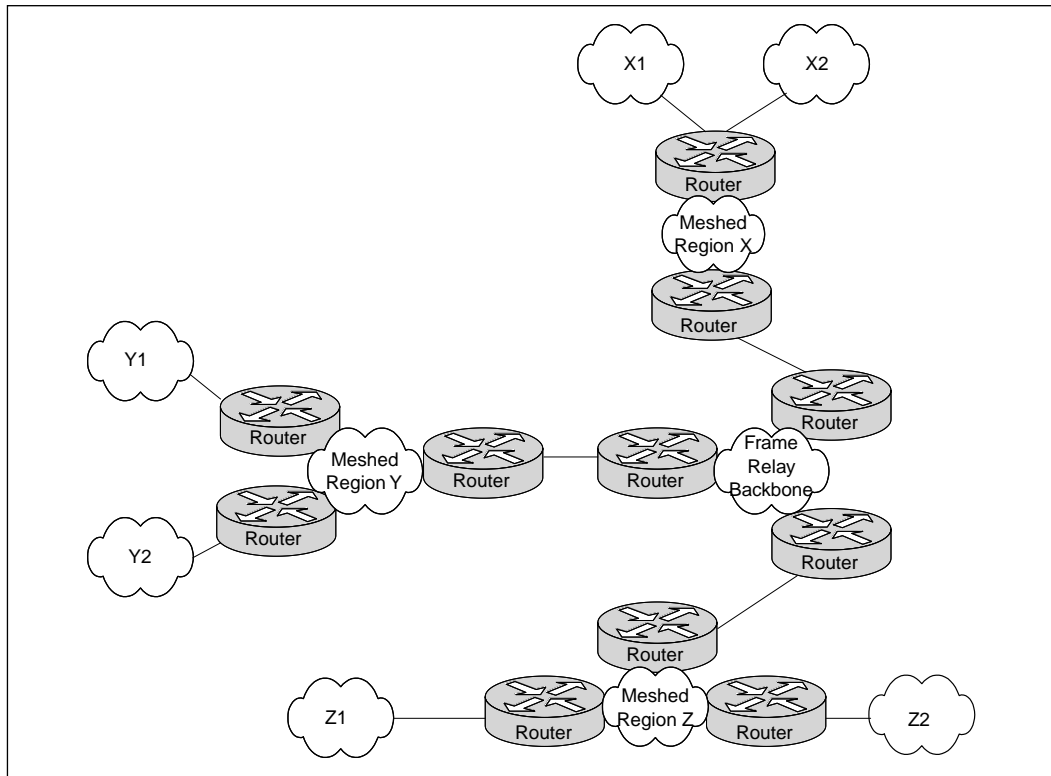
- The hierarchical meshed Frame Relay internetwork
- The hybrid meshed Frame Relay internetwork

These designs have their advantages and disadvantages, and are compared in the following sections.

## Hierarchical Meshed Frame Relay Internetworks

Implementing a hierarchical mesh for Frame Relay environments can assist you in avoiding implementing an excessively large number of DLCIs. This will allow for a more manageable, segmented environment. The hierarchical meshed environment features full meshing within the core PSDN and throughout the surrounding networks. Locating routers between network elements creates the hierarchy.

Figure 8.6 illustrates a simple hierarchical mesh. The internetwork shown illustrates a fully meshed backbone, with meshed regional internetworks and broadcast networks at the outer edges.

**Figure 8.6** Fully Meshed Hierarchical Frame Relay Environment

The advantage of the hierarchical mesh is that it scales well and helps to localize traffic. By placing routers between fully meshed portions of your network, you limit the number of DLCIs that need to be configured per physical interface, segment your internetwork, and make the network more manageable.

However, please remember these two issues when implementing a hierarchical mesh:

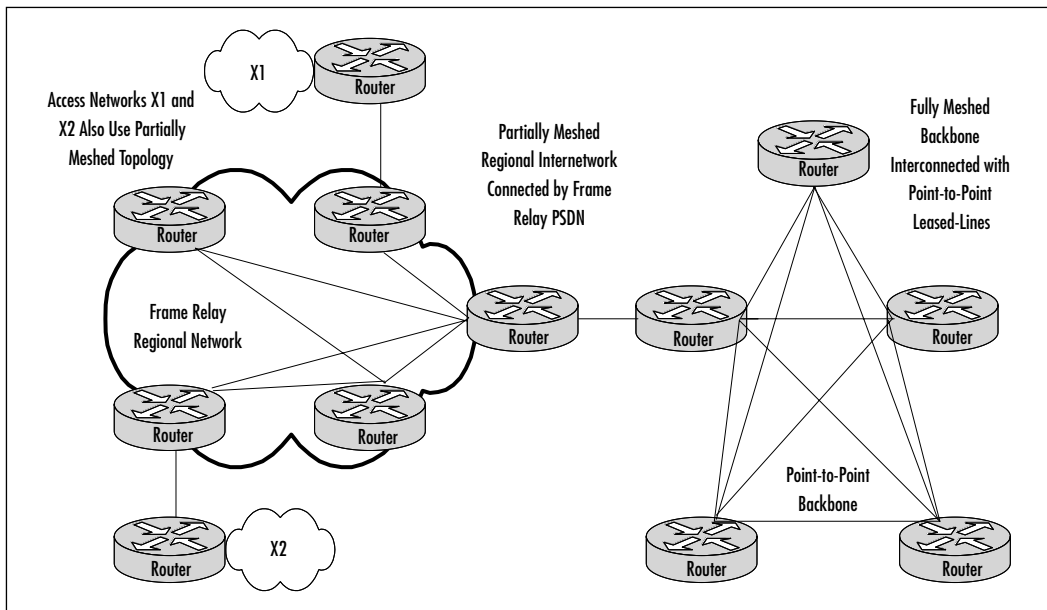
- Broadcast traffic and packet replication** In an environment that has a many routers with multiple DLCIs per interface, there will be excessive broadcast and packet replication, which can impair overall performance. Due to a high level of meshing throughout a network, there will be excessive broadcasts and packet replication that will be a significant resource threat. In the core, where throughput requirements are typically high, the prevention of bandwidth loss due to broadcast traffic and packet replication is particularly important.

- Increased costs associated with additional router interfaces**  
 When compared with a fully meshed topology, additional routers will be necessary to split the meshed core from the meshed edge networks. However, by implementing these routers, you are creating much larger networks that scale ad infinitum when compared to a fully meshed internetwork.

## Hybrid-Meshed Frame Relay Internetworks

The cost-effective and strategic significance of the core network often forces network designers to implement a hybrid-meshed network for their WAN internetworks. A hybrid-meshed network is composed of redundant, meshed lines in the WAN core, and partially (or fully) meshed Frame Relay PSDNs on the network edge. Routers separate the two networks. Figure 8.7 illustrates such a hybrid arrangement.

**Figure 8.7** Hybrid Hierarchical Frame Relay Internetwork



The hybrid hierarchical mesh designs can provide higher performance on the core because they can localize traffic and simplify the scaling of the network. Hybrid-meshed networks for Frame Relay can provide better traffic control in

the core and allow the backbone to be composed of dedicated links, which results in greater stability.

Some of the disadvantages of hybrid hierarchical meshes include the high costs associated with leased lines, and increased broadcast and packet replication traffic.

## Regional Topologies for Frame Relay Networks

There are generally three accepted designs that are relevant for a Frame Relay-based packet service regional network:

- Star topology
- Fully meshed topology
- Partially meshed topology

Each of these topologies is discussed in the following sections. Generally, I have emphasized partially meshed topologies as those that are integrated into a hierarchical environment; star and fully meshed topologies are discussed more for their structural context.

### Star Topologies

Star topology was addressed earlier in the section, “Possible Types of Topology Design.” Star topologies are attractive because they minimize the number of DLCIs that are required, which will result in a lower-cost solution. However, some inherent issues are associated with the star topology because bandwidth limitations. In an environment in which a backbone router is attached to a Frame Relay cloud at 768 Kbps, and the remote sites are attached at 256 Kbps, there will be some throttling of traffic coming off the core that is intended for remote sites.

A star topology does not offer the fault tolerance that is necessary for many networking situations. For example, if the link from the hub router to a specific cloud router is lost, all connectivity to that router is lost.

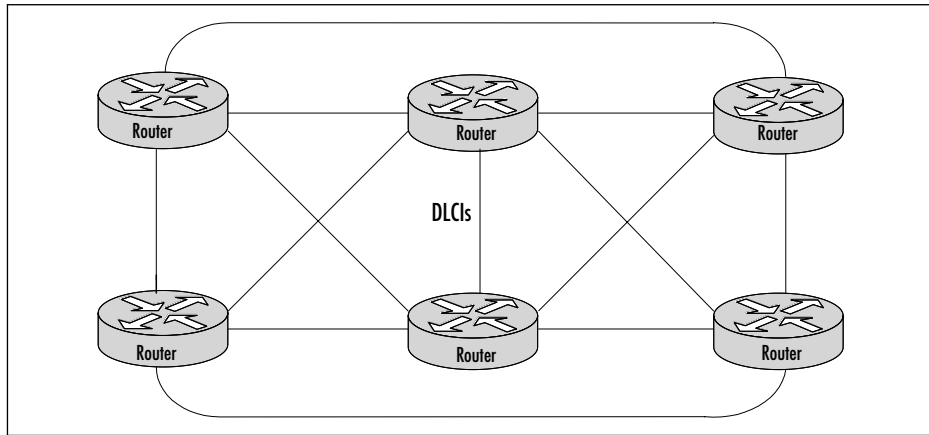
### Fully Meshed Topologies

A fully meshed topology requires that every routing node connected to a Frame Relay network is logically linked by an assigned DLCI to every other node on the cloud. This topology is not easy to manage, support, or even implement for larger Frame Relay networks for several reasons:



- Large, fully meshed Frame Relay networks require many DLCIs. There is a requirement for each logical link between nodes to have a DLCI. As shown in Figure 8.8, a fully connected topology requires the assignment of  $[x(x-1)]/2$  DLCIs, where  $x$  is the number of routers that will be directly connected.

**Figure 8.8** Fully Meshed Frame Relay Network

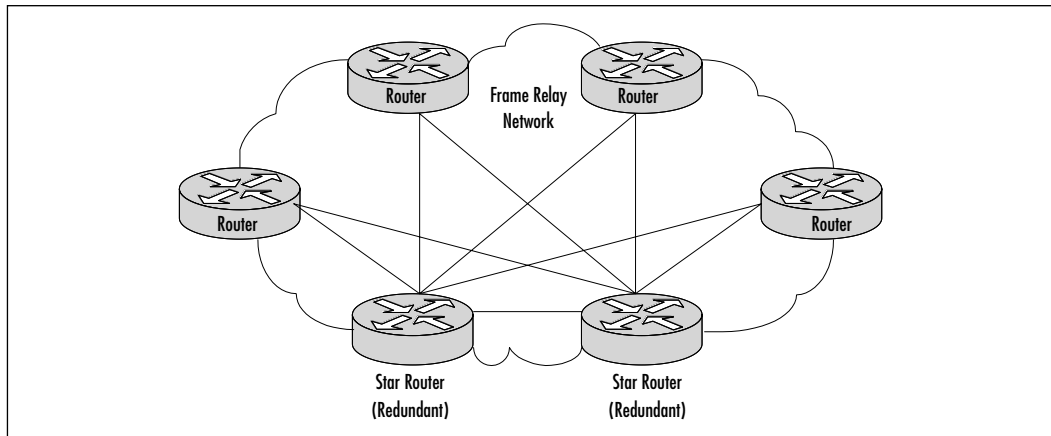


- Broadcast and replication traffic will clog the network in large, meshed Frame Relay topologies. Routers tend to use Frame Relay as a broadcast medium. Every time a router sends a multicast frame (such as a routing update, or spanning tree update), the router will copy the frame to each DLCI for that Frame Relay interface.

This makes fully meshed topologies highly nonscalable for all but relatively small Frame Relay networks.

## Partially Meshed Topologies

By combining the star topology and the fully meshed topology, you will be able to implement a partially meshed topology. Partially meshed topologies are usually recommended for Frame Relay networks that span regional environments as they can provide fault tolerance (through redundant star routers) and are less expensive to implement than a fully meshed environment. As a rule, you should implement at least minimum meshing to eliminate single point-of-failure issues. Virtual interfaces allow you to create networks using partially meshed Frame Relay designs, as shown in Figure 8.9.

**Figure 8.9** A Twin-Star Router, Partially Meshed Network

To create this type of network, you need to assign multiple virtual interfaces (these are considered logical addresses) to individual physical interfaces. In this manner, DLCIs can be grouped or separated to maximize their functionality. As an example, a small, fully meshed cloud of Frame Relay networked routers can travel over a group of four DLCIs that are clustered on a single virtual interface, but a fifth DLCI on a separate virtual interface can provide the connectivity to a completely separate network. This all happens over a single physical interface that is connected to the Frame Relay cloud.

## Broadcast Issues for Frame Relay Networks

Routers treat Frame Relay as a broadcast media, which means that each time the router sends a multicast frame (such as a routing update, or a spanning tree update), the router must replicate that frame to each DLCI that is associated with the Frame Relay physical interface. This broadcast and replication traffic results in substantial overhead for the router and for the physical interface itself.

Consider an IP RIP environment with multiple DLCIs configured for a single physical serial interface. Every time a RIP update is detected, which occurs every 60 seconds, the router must replicate it and send it down the virtual interface associated with each DLCI.

**NOTE**

There are several ways to reduce broadcast and replication traffic within your network. One of the most efficient is to implement some of the more efficient routing protocols, such as Enhanced IGRP, and to adjust timers on lower-speed Frame Relay services.

## Creating a Broadcast Queue for an Interface

When it comes to designing and implementing a very large Frame Relay network, you may come across performance issues when you have many DLCIs that terminate in a single router or that access a server that must replicate routing updates on each DLCI. What occurs is that these updates consume bandwidth and can cause noticeable latency in user traffic. These updates can also consume interface buffers, which will lead to packet loss, and therefore loss of user data and routing updates.

There is a way to avoid these problems, though. You can create a special broadcast queue on an interface. The broadcast queue can be managed independently of the normal interface queue, because it has its own buffers, and the size and service rates of traffic can be configured.

A broadcast queue has a maximum transmission rate (throughput) limit applied to the interface, which is measured in both bytes per second and packets per second. This queue is regulated to make certain that no more than the configured maximum bandwidth is provided. The broadcast queue also has priority when transmitting at a rate that is below the configured maximum to guarantee the minimum bandwidth allocation for the interface. These transmission rate limits are implemented to avoid flooding the interface with broadcasts and replication traffic.

## Committed Interface Rates

When you implement Frame Relay packet-switched networks, external service providers create a CIR that measures in bits per second. This is one of the main metrics. CIR is the maximum permitted traffic level that a carrier will allow on a specific DLCI across its packet-switching environment. The CIR can be anything up to the capacity of the physical media of the connecting link.

One of the drawbacks associated with the CIR is that there are relatively few ways to automatically prevent traffic on a line from exceeding the maximum bandwidth. Although Frame Relay uses the Forward Explicit Congestion Notification (FECN) and Backward Explicit Congestion Notification (BECN) protocols to control traffic in the Frame Relay network, there isn't a standardized mapping between the Frame Relay (link) level and most upper-layer protocols at this time.

What really happens when the specified CIR is exceeded depends on the types of applications that are running on your network. For instance, TCP/IP has something called a backoff algorithm that will see dropped packets as an indication that there is congestion, and sending hosts might reduce output. You should consider the line speeds and what applications are going to be running on your network.

There is a way to buffer Frame Relay traffic so that it can handle instances when traffic exceeds the CIR for a given DLCI. These buffers pool excess traffic to reduce packet loss, especially if you are using a vigorous transport protocol such as Transmission Control Protocol (TCP). Even with these buffers in place, overflows can occur. Remember that routers have the ability to prioritize traffic; Frame Relay switches do not.

You can also specify which packets have low priority or are not time sensitive, so that if there is a need to drop traffic, they will be the first to be discarded when there is congestion. The method that allows a Frame Relay switch to identify these packets is within the packet itself and is known as the discard eligibility (DE) bit.

The Frame Relay network must be able to interpret the DE bit, though, so that there is an action taken when the switch encounters a DE bit during heavy congestion times. Sometimes, networks will take no action when the DE bit is set; at other times, networks use the DE bit to determine which packets to discard. Probably the best way to implement the DE bit is to set it to determine which packets should be dropped first, but also which packets have lower time sensitivity. By doing this, you can define DE lists that will identify the types of packets that are eligible to be discarded, and you can also specify DE groups to identify the DLCI and interface that is being affected.

You can specify DE lists by the protocol and/or the interface used. Characteristics such as fragmentation of the packet, a specific TCP or User Datagram Protocol (UDP) port, an access list number, or a packet size can also be factors that you should consider as DE.

**NOTE**

To avoid packet loss, be sure that you understand what types of applications are going to be on your network and how dropped packets affect them. Try to implement unacknowledged application protocols (such as packetized voice and video) carefully; these protocols have a greater chance of buffer overflow. Dropped packets are okay in voice and video applications, whereas data applications such as ERP should not drop frames.

By defining separate virtual circuits for your different types of traffic, and specifying the queuing and an outbound traffic rate, you can give the customer a guaranteed bandwidth for each of your applications. By being able to specify different traffic rates for different virtual circuits, you can perform virtual time division multiplexing (TDM).

You can also ease congestion and data loss in the network by throttling the outbound traffic that comes from high-speed lines in central offices, and go to low-speed lines in remote locations. This enhanced queuing will also help to prevent congestion that causes data loss. This type of traffic shaping can be applied to both PVCs and SVCs.

## Capacity Planning for Your Infrastructure

Have you ever tried to install a server, and there just isn't a spare rack, or even an available network port? This is usually an issue for service providers who did not plan properly for explosive growth. A ripple effect of this is that the network probably is staggering with the amount of clients that were added. In this section, we discuss some best practices for implementing a capacity plan.

## Connection and Expansion

Capacity planning is an issue that you will need to address, along with the design and implementation procedure. If you have a general idea of where you stand for number of servers and expected growth, you can use those as a baseline for the capacity of your network. The reason that I say "baseline" is that many things can happen in the business world over the course of six months that might cause

your design to be underpowered and/or oversubscribed. Depending on the size of the current and expected network, there should be a padding area of approximately 10 to 15 percent for unexpected growth. You should reevaluate capacity after every new customer you install.

## Best Practices

One of the best practices for planning is to map out where the different customer areas are located, and what the server count is going to be. Once these figures are determined, decide if the servers need one data link or multiple connections.

With the number of connections decided, you now need to plan on subscription and maximum bandwidth provisions of the network. Sounds impressive, huh? It really isn't that difficult. What you are doing is calculating the aggregate average bandwidth of the network devices located on the segment. With these calculations in place, you can plan whether the segment is powerful enough to support the clients and resources on a given network.

So, how do you calculate the aggregate average bandwidth? The calculations are based on network topology, users' traffic patterns, and network connections. Ask questions such as, "What type of links should we use to connect the clients to the data center? What should the bandwidth requirement be for the backbone?"

You need to plan what type of link goes to each client, so that it has the proper bandwidth, yet does not allow for the monopolization of network resources, and/or completely shut down the network with oversubscription of a segment. Monopolization of a segment occurs when a user has the equal bandwidth of a resource (such as an application server), and the application takes all of the available bandwidth and maintains the trunk. This will not allow other clients to access the resource.

Oversubscription of a segment occurs when multiple customers use all of the application resource's bandwidth, and therefore, other clients are unable to access the resource. While the two symptoms just described result in the same conclusion, they are different. The segment in a monopolized environment runs consistently, whereas the segment that is oversubscribed may shut itself down because it cannot pass traffic due to multiple requests flooding the buffers on the switches and routers.

With all of this information, you should get an idea for the type and capacity of network equipment you need to deploy at each location. Further concerns for planning capacity are based on factors such as which protocol you use, the addressing schemes, the geography, and how these fit the topology of the network.

# Protocol Planning Concerns

The following section provides details on how to choose a protocol to best match your environment. With proper planning and implementation, any choice will work. By determining the physical layout of the network, you will be able to map the correct topology and form a logical addressing scheme that will grow as your network grows.

## Routing Protocols

Choosing routing protocols and their configuration are important parts of every network design. You must be prepared to spend a significant amount of time implementing your policies for the network to provide optimal performance. Routing protocols are a fundamental component of networking and creating a reachable network that can transfer data.

If designed properly, the network will build routing tables and maps that you can use to see adjacent routers and their status. There is also the ability to see network paths, congestion, and bandwidth of those links. This information helps in deciding on the optimal network paths.

The more complex routing protocols allow you to add secondary metrics. Some of those metrics include reliability, delay, load, and bandwidth. Using these metrics, the router can make these routing decisions dynamically.

The basic difference in the various routing protocols lies in the sophistication of their decision-making capabilities and metric support. This is one of the main factors to consider when choosing a protocol to match the characteristics of your network.

There are two types of routing protocols, internal and external. Internal protocols are those that you would implement within your network infrastructure, and are controlled completely within that domain. Conversely, external protocols work with external domains, such as the Internet or other ISP networks. These protocols are designed to protect your domain from external errors or misrepresentation. To read more about internal and external protocols, visit the Internet Engineering Task Force's site ([www.ietf.org](http://www.ietf.org)).

## Interior Gateway Protocols


How do you decide on which interior protocol to use? The following section explains some of the inherent differences that are implemented in each of the interior protocols. Some are self-explanatory, and some are just more complex.

I hope that this section will clarify some of the terms and allow you to make reasonable, well thought-out decisions.

Let's deal more with interior protocols, since those are the ones that you may have used less in your ISP environment. With the exception of OSPF and IS-IS, the interior routing protocols described are distance-vector protocols, and they use distance and next-hop data to make their routing and forwarding decisions.

Some distance-vector protocols are very simplistic and don't scale well in larger environments. One example is RIP, which uses hops (the number of connections between it and its destination) as its determining factor. The largest hop count before it disregards the packet is 15, making it one of the least scalable protocols. Another drawback to RIP is that it does not take into account varying available bandwidth. If, for example, you have a packet that needs to get from network A to network D, RIP will take the path with two hops, rather than the path with two hops but higher speed (Figure 8.10).

**Figure 8.10** RIP and Routing Decisions



If your network is fairly simple in terms of the topology and number of routers, a distance-vector protocol such as RIP or IGRP (discussed later in this chapter) could work fine. If you're running a multivendor network, RIP, RIPv2, IS-IS, and OSPF are common protocols across many vendors' router implementations.

This chapter lists what is available for use in a semichronological order, as this is probably how you will see them listed in other reference manuals. Some of the strengths and limitations for each protocol are also listed.



### *Routing Information Protocol*

The Routing Information Protocol (RIP) was derived from Xerox Corporation's XNS for IP networks. It supports IP and IPX networks.

Strengths:

- Still viable in networks that use a constant internal subnet
- Usable on most vendors' equipment
- Low cost (generally free from most vendors)

Weaknesses:

- Scalability is minimal (15-hop maximum)
- Path determined by hop count, and may take best path
- Broadcasts full routing table frequently, wasting bandwidth
- Cannot handle variable-length subnet masks (VLSM)

### *Interior Gateway Routing Protocol (Cisco Only)*

The Interior Gateway Routing Protocol (IGRP) was created by Cisco for Cisco IP and OSI networks. It supports IP and OSI networks.

Strengths:

- Uses multiple metrics in decision making
- Fast convergence

Weaknesses:

- Runs only on Cisco equipment
- Broadcasts routing table frequently, wasting bandwidth

### *Open Shortest Path First*

The Open Shortest Path First (OSPF) was developed by the IETF for IP networks. It supports IP networks.

Strengths:

- Usable on most vendors' equipment
- Only broadcasts routing table when changes are made

Weaknesses:

- Uses only bandwidth as a metric
- Restricts some topologies

### *Integrated Intermediate System-to-Intermediate System*

The Integrated Intermediate System-to-Intermediate System (IS-IS) was developed by IETF for OSI and IP networks. It supports IP and OSI networks.

Strengths:

- Usable on most vendors' equipment
- Only broadcasts routing table when changes are made
- Fast convergence

Weaknesses:

- Uses only bandwidth as a metric
- Restricts some topologies

### *RIPv2*

RIPv2 was developed by the IETF for IP networks. It supports IP and IPX networks.

Strengths:

- Added authentication and multicast ability to RIP
- Usable on most vendors' equipment

Weaknesses:

- Scalability is minimal (15-hop maximum)
- Only uses bandwidth as metric

### *Enhanced Interior Gateway Routing Protocol (Cisco Only)*

The Enhanced Interior Gateway Routing Protocol (EIGRP) (Cisco only) was designed by Cisco for multiprotocol Cisco networks. It supports IP, IPX, and AppleTalk networks.

Strengths:

- Incremental updates to reduce broadcast traffic
- Supports VLSMs
- Uses multiple metrics
- Fast convergence
- Retains backward compatibility with IGRP

Weakness:

- Works only with Cisco equipment

## External Protocols

There is really only one external protocol to talk about, the Border Gateway Protocol (BGP). There have been entire books written on BGP, and several companies make their living working with BGP exclusively. These companies are highly paid, because you do not want to create incorrect BGP parameters and advertisements. These common mistakes will definitely adversely affect your network.

I will not go into any detail here, because if you have BGP and it is up and running, you most assuredly do not want to mess with it.

## Choosing the Right Interior Protocol

Using the previous references that explained the differences inherent with the interior protocols, it is time to address which networking protocol will be best for your network. There are several considerations to take into account for your choices. You want the protocol to add functionality, you want it to be scalable, you want it to easily adapt to changes that will be implemented, you want it to be manageable, and you want it to be cost effective.

Scalability is an issue that will keep cropping up. You want your design to meet current and future growth. What happens if you decide to implement RIP, and then exceed the hop count within your own network? Preplanning will save you some massive restructuring headaches later in the development of the infrastructure.

Adaptability to new technologies is essential. With the world moving toward higher and higher bandwidth usage, you must take into account that what you design must be readily adaptable for future technologies (e.g., Voice over IP (VoIP) and video traffic) that will possibly be implemented. This is an area where adding value to your network will save your company money in the long run.

Another key concern is the manageability of the network. You are going to get into an area where there is more network traffic than bandwidth, and you are going to want to manage traffic and to some extent monitor the usage of the links that are on your network. With that criteria in mind, plan on protocols that will be more beneficial to you and the management of the network.

Finally, one of the toughest areas when choosing a routing protocol is cost effectiveness. Some of the protocols can only run on Cisco equipment because they are proprietary. While RIP will run on most equipment, and is very cost efficient, the network would have to be scaled down so that it could be properly implemented. A protocol such as EIGRP would be fantastic in a larger-scale environment, but it only works with Cisco equipment, so you need to keep that in mind if you have legacy equipment you are planning to phase out.

## Route Selection

So why are we talking about route selection? It seems somewhat silly if there is only one route to the destination. The question is, what happens if that route should fail? What about using other routes to allow more traffic and less congestion? These reasons are why most networks are designed with multiple routes (redundancy and load balancing), so there is always an alternate connection in case of failure or to alleviate traffic issues. Routing protocols use metrics to select the best route based on weighted decisions, from groups of existing routes.

Metrics are values that can be assigned and weighted to make decisions on routing paths within a network. Metrics are assigned characteristics or a set of characteristics on each link/route of a network. When traffic is passed along a link, the network equipment makes a choice on how to route the traffic by calculating values of the metrics and assigning the traffic to the selected path.

Metrics are handled differently depending on the routing protocol. Most of the routing protocols can use multiple paths if they are equal cost. Some protocols can use paths that have routes that are not cost equivalents. By implementing multiple paths, you can use load balancing to improve bandwidth allocation.

With a multiple path design, there are two widely used ways for packet distribution: *per-packet load balancing* and *per-destination load balancing*.

Per-packet load balancing uses all possible routes in proportion to the route metrics. What this means is that if all routes are equal cost, it will cycle through them in a “round robin” selection scheme, where one packet is sent to each possible path that is available. Routers are defaulted to this method when fast switching is disabled.

Per-destination load balancing uses routes based on what the destination. Each destination is assigned an available route and maintains that route for future use. This way, traffic tends to arrive in the proper order. Routers default to this method when fast switching is enabled.

## NOTE

---

TCP can accommodate out-of-order packets, and there will always be some on multipath networks. Excessive out-of-order packets can cause performance issues, so don't go overboard on the redundant load balancing. Make sure to have at least one extra path between segments and NO MORE if redundancy is an issue.

---

Preplanning these rollouts will save aggravation and long nights of troubleshooting. Sometimes it is best to sit down, create a matrix with these considerations, and cross-reference what will be best for your network. Try to use a single interior routing protocol for your network. Sometimes this is not possible, so in those instances, plan for the best protocol mesh and implement as cleanly as possible.

The best advice that I offer is, “keep it simple.” Complexity is not a good thing in network design. You want a stable, simple network. Now, on to the next decision-making process, addressing.

## Addressing Considerations

The addressing scheme is dependent on several variables. Since most companies are now creating firewalls and implementing Network Address Translation (NAT) and/or Port Address Translation (PAT) to their public address space to conserve addresses, you have to figure out how to implement your addresses in a logical manner. Since NAT is so popular, private addressing is almost completely up to you. The three private address spaces are:

- **Class A** 10.x.x.x
- **Class B** 172.16.x.x
- **Class C** 192.168.x.x

These are private addresses, and cannot be routed on the Internet. Using subnet masks will allow you to configure Network and Host IDs to suit your needs.

Since addressing is important, and time consuming, you want to plan allocation of addresses as neatly as possible so you have to readdress your network as little as possible. Plan on how you want to slice up the company, whether it is by site, location, department, telecommuter, and so forth. If you are using a public address scheme, conservation is key. If you are using a private addressing scheme, you should focus on layout. Deployment should be in a logical and readily understood manner for public and private addressing. You want to keep similar users as well and similar resources together.

Set aside address space for current and potential (have I said those words before?) users. If you want to keep all of the servers in one address space, set aside enough for future growth. Generally, you should group users with similar needs in similar locations (with the advent of VLANs, this is not the major concern it was in the past). Remember that when you add subnets, you need to place routers between subnets that cannot speak to each other.

Since the network doesn't physically exist at this point in time, be sure to revise all of your ideas on paper. This will help for the actual implementation. Another way to help segment traffic and design your addressing scheme is by topology.

## Topology

The topology of a network is defined by sets of routers and the networks to which they connect. Routing protocols can also establish a logical topology depending on implementation.

TCP/IP requires the creation of a hierarchical topology that establishes a Core layer, a Distribution layer, and an Access layer. For example, OSPF and IS-IS protocols use a hierarchical design. The hierarchical topology takes precedence over any topology created through address segmentation. Therefore, if you choose a hierarchical routing protocol, you should also create the addressing topology to reflect the hierarchy. If you decide to use a flat routing protocol, the addressing will create the topology.

There are two regularly accepted ways to assign addresses in a hierarchical network. The easiest way is to assign a unique network address to all network areas (including the core). A more complex way is to assign ranges of addresses to each area. Areas should be comprised of contiguous addresses for networks and

hosts. Areas should also include all the router interfaces on any of the included networks. By doing so, each area maintains its own topology database, because all interfaces run a separate copy of the basic routing algorithm.

Flat networks were originally campus networks that consisted of a single LAN to which new users were added. This LAN was a logical or physical cable into which the network devices connected. In the case of Ethernet, all the devices shared the half-duplex 10 Mbps available. The LAN was considered a collision domain, because all packets were visible to all devices on the LAN; therefore, they were free to collide, given the carrier sense multi-access with collision detection (CSMA/CD) scheme used by Ethernet.

A bridge was inserted when the collision domain of the LAN became congested. This allowed a segmentation of traffic into several collision domains, because a bridge is a store-and-forward packet switch. With this ability to cut collision traffic, network throughput increased. The drawback is that bridges pass all traffic, including, flood broadcasts, multicasts, and unknown unicasts, to all segments.

All the bridged segments in the campus together form a single broadcast domain. The Spanning Tree Protocol (STP) was developed to prevent loops in the network and to route around failed connections.

## Configuring & Implementing...

### The STP Broadcast Domain

There are some issues with the STP broadcast domain. It has a high time threshold for convergence, typically 40 to 50 seconds. It allows for nonoptimized paths to exist. Redundant links carry no data because they are blocked. Broadcast storms affect the whole domain, and each network host must process all traffic. Security is limited within the domain, and troubleshooting problems is time consuming. New versions and options can reduce convergence time.

Broadcast traffic sets a practical limit to the size of the broadcast domain. Managing and troubleshooting a bridged campus becomes harder as the number of users increases because it adds to the broadcast domain. One misconfigured or malfunctioning workstation can disable an entire broadcast domain for an extended period of time, as it is generally hard to locate.

**NOTE**

The 80/20 traffic rule has been steadily changing due to the rise of intranets and distributed applications. With new and existing applications moving toward a distributed applications and storage model, which are accessed through Web retrieval, the traffic pattern is going toward the 20/80 model, where only 20 percent of traffic is local to the workgroup LAN, and 80 percent of the traffic is destined for a nonlocal domain.

## Application and Network Services

Knowing the various network applications that will be on the wire is essential to the planning and design of your new network. You need to take into account all of the different types of protocols, ports, and bandwidth that will be used on the wire when running different programs. For instance, standard file and print traffic has far less overhead than database and backup applications do. Each program behaves in its own way, and you need to anticipate the behavior in order to design the network properly.

Make a list of the applications that will be running on the wire, and remember, be thorough! Make sure that you take into account every possible traffic generator (whether it is an application or a network device) that you can think of, no matter how trivial, and add it to the pile of your bandwidth calculations. Make sure also to give each application its proper “weight” in your calculations.

For instance, it is highly doubtful that people will be printing 100 percent of the time on the network, so make sure that you apply the proper percentage bandwidth subscription and utilization in your calculations as to how much time you will see the traffic.

Although this seems like a tedious task, it is well worth the exercise. Remember, the customer will be extremely unhappy with poor performance, and in all likelihood, you, the provider, will be the one blamed if the bandwidth is choked on the first day of the new network rollout. Try to overestimate the amount of bandwidth required at the worst times of traffic by about 20 percent. That will always lead to a safe estimate of bandwidth.



## Designing the Data Center Network

When designing the data center, you should build the network as a modular building block using multilayer switching. This way, you can segment the traffic so that it goes over specific bandwidths. For example, the Gigabit Ethernet trunk carries the server-to-server traffic, and the Fast EtherChannel trunk carries the backbone traffic, so all server-to-server traffic is kept off the backbone, which has performance and security advantages.

Data centers should have Virtual Router Redundancy Protocol (VRRP) redundancy between the multilayer switches, with access lists used to control access policy to the data center. With this setup, the core switches are separate from the distribution switches, thereby cutting down on broadcast domains and allowing for better throughput.

Note that when using the Hot Standby Router Protocol (HSRP) (Cisco specific) or Virtual Router Redundancy Protocol (VRRP), which can also add redundancy, you should consider implementing Fast EtherChannel so you can scale bandwidth from Fast Ethernet, and from Gigabit Ethernet to Gigabit EtherChannel. You should also consider using adaptive load balancing for high availability.

## Terminal Data Centers

One of the benefits of using the Windows 2000 software platform is the built-in capability to use Terminal Server on the network. Terminal Server is the Microsoft version of thin-client technology that allows the workstations to receive a video snapshot of the desktop running on another machine, and control that machine remotely.

The benefit of this is so that bandwidth-heavy applications such as database and other query-based applications need not send the queries back and forth over the entire network on the “slow” WAN links. Rather, they have access to the “client” machine via high-speed switched links, and users on the network will be accessing the database applications as if they are right next to the database servers on the switch.

This technology has become increasingly popular over the last couple of years, as bandwidth has become a premium on the WAN links. Microsoft has built this technology into Windows 2000 so that any Win2K server can activate this remote control capability.

# Application-Aware Networking

ASPs who want to deploy their applications need to realize that their success of mission-critical applications over both the internal LAN and clientele WAN is achieved by defining network policies, which assist in the apportioning of network resources with business objectives. These policies are then enforced with the creation of a Quality of Service (QoS) method for their application and bandwidth considerations. Without these QoS controls in place, nonvital applications and services can quickly overwhelm network resources by taking the available resources away from those applications that you are selling and causing poor customer satisfaction with your product.

When taking into account design considerations, remember that the ability to provide an end-to-end application that has prioritization within a LAN will need specific features that are dependent on where the networking device is located.

In a service provider environment, Access layer devices should deliver the following capabilities:

- **Traffic detection** The ability to identify types of applications.
- **Admission control** The ability to restrict application traffic.
- **Traffic classification** The ability to classify traffic, based on application type, destination, or physical port.

## Traffic Detection and Classification

Traffic detection uses Layer 4 IP UDP and TCP port numbers. A switch that is able to examine a packet and identify the UDP and TCP port numbers can use this information to identify the application that is using that packet. A switch can then compare a packet that is being used to run a nonmission-critical application or service, such as email, against a packet that is running an application that is considered mission critical and classify it appropriately.

## Admission Control

Admission control is provided by a mechanism that can reject or remove applications based on user-defined policies. For example, a client can define a policy to temporarily stop the transmission of email packets, so that the mission-critical applications can use the necessary resources.

Admission control can also be used to provide the following benefits:

- During network congestion, you can specify that a specific application (e.g., video or multimedia traffic, which is usually considered bandwidth intensive) be dropped for the congestion period.
- You can bar specific applications from entering the network, which would save network bandwidth.

## Traffic Classification

Traffic classification uses traffic detection to identify what applications are running. These applications can then be classified based on their user-defined policy. This classification will provide a tag that identifies the priority of the packet. This classification information is carried within the precedence bits of the IP packet.

The classification technique that uses the IP precedence fields is commonly referred to as Type of Service (ToS). If you are going to use either Cisco Systems Inter-Switch Link (ISL) or 802.1Q, this is usually referred to as Class of Service (CoS).

When using the IP precedence (which uses Layer 3) as the classification tag, you can deliver a QoS identifier that is independent of the transmission technology that is allowing for the mapping of QoS across disparate media types such as Ethernet and ATM.

ISL and 802.1Q use Layer 2 technologies that are more suitable to campus deployment across Ethernet networks that aren't totally IP based. Unlike Layer 2 classification (ISL and 802.1Q), Layer 3 classification (ToS) has the ability to provide consistent end-to-end classification for a packet as it negotiates through LAN/WAN and intranet/Internet borders.

In the ideal environment, packet classification should only happen once. Both ISL and 802.1Q have a three-bit field that can provide up to eight priority levels, which helps in delivering consistent priority-level maps between technologies.

Traffic classification provides the following benefits to an ASP:

- It can provide a tag that will identify the priority assigned to each packet. This gives devices that are downstream the information they need to prioritize packet traffic.
- It provides the mechanism that can reclassify packets based on policy. As an example, when a client's workstation applies a high priority to multimedia applications using traffic classification and detection, the

application, might be reclassified to a lower priority, depending on the implemented network policy.

In the Distribution and Core layers of a network, the devices should support the following features:

- **Congestion avoidance** The ability to drop certain packets based on their classification or policy during periods of network congestion.
- **Scheduling** The ability for the network equipment to show traffic transmission preference, based on classification of the application or the policy of the network.

## Congestion Avoidance

By implementing Weighted Random Early Detection (WRED), you can add congestion avoidance in your infrastructure. Congestion avoidance can take advantage of TCP windowing to slow lower-priority application traffic before network congestion can affect the higher-priority applications. WRED is based on the output queues of the network device.

When network traffic increases, the chances of network congestion increases, because the buffers begin to fill. Eventually, the buffer will overflow and packets will be dropped in an unrestrained fashion. TCP traffic tends to increase its rate of data transmission; until packets are dropped or the maximum TCP transmit size is reached. When a data transmission mismatch happens (such as packets from a Gigabit Ethernet port destined for a FastEthernet port) the device's buffers will begin to fill.

By dropping packets, random early detection (RED) helps to prevent a buffer overflow from randomly dropping packets, which can adversely affect application performance. WRED enhances RED by specifying which packets are dropped when they reach a buffer's threshold.

There are generally three classes to packet classification: premium, standard, and DE. Mission-critical applications are normally mapped to the premium class, while all other traffic would be mapped to the standard threshold. If the buffer reaches 60 percent of full capacity, packets that belong in the standard class would begin to be randomly dropped. As the buffer continues to fill, the drop rate for the standard class applications will also increase. The premium service may be configured to not drop those packets until the buffer has reached 90 percent of its full capacity.

Since it is harder to reach the higher thresholds (unless the network becomes overly congested due to a large amount of higher priority applications), higher-priority applications will continue to have end-to-end connectivity and performance. Consequently, traffic that is considered mission-critical is unaffected by applications with a lower priority.

## Scheduling

In a switch, there is a switching fabric with a finite amount of backplane and time to transmit packets out a given interface. By implementing weighted round robin (WRR) in your network devices, you can provide preferential treatment to packets based on their priority. With this method, you could allow a given interface to use 70 percent of its transmit capacity for mission-critical or delay-sensitive applications, and the other 30 percent could be slotted for applications that are less critical or delay sensitive.

## Scalability Considerations

We now need to discuss the issue of scalability. When you implement your design, there are going to be times when you will experience explosive growth. (I believe that marketing people call it “Hyper growth.”) So, how do you keep up with the Jones? Scalability. Scale the bandwidth. Scale the equipment. In the following section, we discuss the scaling of bandwidth and some of the concerns that you will face.

## Scaling Bandwidth

Bandwidth in the multilayer model can be scaled in many ways. Ethernet can be upgraded to Fast Ethernet, and Fast Ethernet can be combined into Fast EtherChannel or Gigabit Ethernet or Gigabit EtherChannel. Access switches can be partitioned into multiple VLANs with multiple trunks, and use ISL or 802.1q in the VLANs to combine the different trunks. Fast EtherChannel provides more efficient utilization of bandwidth by multiplexing multiple VLANs over one trunk. To scale bandwidth within ATM backbones, you must add more OC-3 or OC-12 trunks all the way up to OC-192.

## Scaling Considerations

The great thing about designing the network with multilayer switching is that it’s extremely scalable. Routing is able to scale because it is distributed, and therefore

it is easy to add pieces that point to other pieces. The backbone performance scales when you add more connections and/or switches.

Since the network is compartmentalized, it is also scalable from a management and administration perspective. When issues crop up in the network, you can pin the problem down to one of the layers, and troubleshoot it down from there.

When designing your network, avoid creating STP loops in the backbone. STP takes 40 to 50 seconds to converge and does not allow for load balancing across multiple paths. When using ATM for your backbone, use PNNI to handle load balancing.

You should always try to use high-level routing protocols such as OSPF, IS-IS, and Enhanced IGRP (Cisco only), which allow for path determination and load balancing. OSPF and IS-IS operating costs at the core will rise linearly as the number of switches in the Distribution layer increases. What happens is that OSPF elects a router and a backup router, which will connect with all of the other routers in the Distribution layer.

If multiple VLANs or ELANs are created in the backbone, a primary and a backup router are elected for each. Remember that with OSPF routing traffic, CPU overhead increases as the number of VLANs or ELANs increases on the backbone. Therefore, try to keep the number of VLANs or ELANs on a trunk to a minimum.

The following are some suggestions for best practices:

- Remember that OSPF needs summarization to allow it to scale. It is a common practice in large campuses to make each building an OSPF area, and make the routers area border routers (ABRs).
- Try to make all of the subnets to a single customer from a contiguous block of addresses. This will allow you to use a single summary advertisement on the ABRs. By doing so, you will reduce the amount of routing information traffic and increase the stability of the routing table.
- Enhanced IGRP can be configured in roughly the same way; however, here are some exceptions to the rule. Protocols such as Novell Server Advertisement Protocol (SAP), Novell Routing Information Protocol (RIP), and AppleTalk Routing Table Maintenance Protocol (RTMP) have an overhead that increases exponentially as you add connections. These protocols should be the exception rather than the rule in an ASP environment.

# Multimedia Services

According to a study by the Telecommunications Industry Association, the multimedia application market (such as video on demand, VoIP, etc.) is expected to reach \$16 billion in 2001. Some of the needs that are pushing the growth of these services and applications are distance learning, virtual workplace, audio and videoconferencing, streaming media applications (video on demand), data storage (SAN/NAS), and financial applications (such as ERP and CRM).

Many of the new multimedia applications that customers want, require IP multicast for proper operation. Any network communication that needs to transmit information to multiple clients can benefit from the efficiency of multicast technologies.

As an example, applications that involve one-to-many or many-to-many communications include:

- Database replication
- Dissemination of news feeds and stock quotes
- Software downloads
- Video and audio broadcast
- Videoconferencing and collaboration
- Web site caching

To get reasonable full-motion full-screen viewing requires approximately 1.5 Mbps of bandwidth. In a unicast environment, a separate video stream is sent to the network for each client who wants to view the information (this uses  $X \star 1.5$  Mbps of link bandwidth, where  $X =$  number of viewers). With a 100 Mbps FastEthernet interface on the server, 60 or 70 video streams will saturate the interface.

Even with a Gigabit Ethernet interface on a high-end server, the practical limit would be from 250 to 300 1.5 Mbps video streams. In as such, the server's interface can be a bottleneck by limiting the number of unicast video streams per server.

As you can see, this replicated unicast transmission can consume a lot of bandwidth and other network resources, and is therefore a limitation. With clients separated from the server by two Distribution and/or Core switch hops and two Access layer switch hops, a single multi-unicast stream would consume 300 Mbps

of Distribution and/or Core switch bandwidth, and 300 Mbps of Access layer switch bandwidth.

Even if video stream bandwidth were scaled back to 100 Kbps (kilobits per second, which can provide an acceptable quality of video in smaller windows areas), the multi-unicast traffic would still consume 20 Mbps of Distribution/Core switch and Access layer switch bandwidth.

Even in a multicast environment, a video-streaming server will still have to transmit a single video stream for each multicast group, irrespective of the number of end clients who will view it. The video stream is replicated as required by the network's multicast Distribution or Core switch and Access layer switches to allow a subjective number of clients to subscribe to the multicast address, and therefore receive the broadcast. In the Distribution and Core switch network, replication occurs only at branches of the distribution tree. What this means is that essentially, all of the replication occurs at the last switch hop before the end user. Within the network, the ability to have multicast transmission offers greater efficiency, as it consumes only 1/nth of the bandwidth of the multi-unicast solution.

When there are many end users for a replicated transmission, multicast technology can make a tremendous difference in both server load and network load, even in a small network with a limited number of Distribution and Core switches and Access layer switch hops.

## IP Multicast

Applications that use IP multicast are a rapidly growing piece of company networks. Applications such like Microsoft NetShow and NetMeeting are being used more commonly to do voice and video streaming. There are several considerations when using IP multicast:

- Protocol independent multicast (PIM) routing, in either dense mode or sparse mode
- Clients/servers join multicast groups with Internet Group Management Protocol (IGMP)
- Multicast tree pruning with Cisco Group Multicast Protocol (CGMP) or IGMP snooping
- Switch and router multicast performance
- Multicast policy



The most common routing protocol for multicast is PIM. PIM is broken up into two parts: PIM sparse mode, and PIM dense mode. Sparse-mode operation is used with applications such as NetMeeting, whereas dense-mode operation is used for an application such as IPTV. PIM is being used in the Internet and in corporate intranets.

A strong point of PIM is that it works with various unicast routing protocols such as OSPF and Enhanced IGRP. PIM routers are also compatible with the Distance Vector Multicast Routing Protocol (DVMRP). DVMRP is an older multicast routing protocol that was used somewhat extensively on the Internet multicast backbone (MBONE). It is expected that PIM will replace DVMRP over time.

PIM works by building multicast trees that minimize traffic on the network. This is important for applications such as real-time video, which uses large amounts of bandwidth. PIM is usually configured in sparse-dense mode, and automatically uses either sparse or dense mode depending on the application.

Multicast clients and servers that wish to join or advertise multicast groups use IGMP. The router in this type of environment makes multicast available on subnets with configured (open to receive) clients, but blocks the traffic if there are no clients open. CGMP allows multicast pruning on a Catalyst switch. A Cisco router sends out a CGMP message to advertise all MAC addresses that have joined the multicast group.

Cisco Catalyst switches receive the CGMP message and forward traffic only to ports with those MAC addresses in the forwarding table. This blocks multicast traffic from ports that don't have members connected to them. The Catalyst series has an architecture that forwards multicast streams to one port, multiple ports, or all ports (there is no performance penalty). Catalyst switches can support many multicast groups concurrently.

Multicast policies can be implemented by placing multicast servers in a data center behind a Catalyst switch. The multilayer switch acts as a multicast firewall that controls access to multicast traffic. To segregate multicast traffic, create a separate multicast VLAN or subnet on the backbone. In addition, you have to create a PIM rendezvous point, which is the root of the multicast tree.

## Virtual LANs and Emulated LANs

*Virtual LANs* (VLANs) were developed to enable Layer 2 switching across the campus. A VLAN is a way to create an extended logical network that is independent of the physical network layout. A VLAN functions as a separate broadcast

domain, and is similar to an extended bridged network. STP is generally implemented between the switches in a VLAN.

Another technology developed to enable campuswide VLANs is VLAN *trunking*. Trunking allows traffic from several logical Layer 2 networks to be multiplexed (combined). Creating a VLAN trunk between a Layer 2 switch and a router allows the router to connect to several networks with a single physical interface.

Inter-Switch Link (ISL), 802.10, and 802.1q are VLAN tagging protocols that were created for VLAN trunking. A VLAN tag is a number that is placed in the header of frames that go between two devices. The tag number value allows the data from different VLANs to be multiplexed and demultiplexed.

ATM LANE is the technology that permits multiple logical LANs to exist over a single switched ATM network. ATM Emulated LANs (ELANs) use a similar tagging method as packet-based technologies, so ISL, 802.10, and 802.1q are compatible with Ethernet VLANs. LANE clients (LECs) connect Ethernet VLANs across the ATM backbone. To make ATM LANE work like Ethernet, you also need, a LANE configuration server (LECS), LANE server (LES), and broadcast and unknown server (BUS). With these implemented, ATM LANE will emulate the Ethernet broadcast protocol over ATM.

## NOTE

---

Ethernet-connected hosts and servers in one VLAN cannot talk to Ethernet-connected hosts and servers in a different VLAN without a Layer 3 device in between.

---

## Policy in the Core

With routing done in the Distribution layer, it is possible to implement the backbone as a single logical network or multiple logical networks. VLANs can be used to create separate logical networks that can be used for multiple purposes. For example, a VLAN could be created for traffic management. Policies could be implemented for each core VLAN and applied with access lists on the router.

You could partition the core by protocol. What you would need to do is create one VLAN for each server, based on the protocols they use, (IP, IPX, etc.). These partitions can become complete physical separations on multiple core switches.

Try to keep the backbone topology simple. Try keeping the number of VLANs (or ELANs) small so that they are easily managed.

## WAN Link Considerations

When creating a WAN network, the need for QoS comes into play. Any time there is a potential bottleneck in the network, queuing techniques should be applied so that delay- and drop-sensitive traffic such as voice and real-time video pass through with the least interference. This is typical at the WAN edge router, where all data traffic destined for other networks is aggregated into slower-speed links. Whatever the queuing mechanism you use, it will likely classify data, voice, and video packets so that they are allowed proper throughput.

Low-speed links need special consideration. For example, if voice traffic is sent out a connection, and a data packet is sent at the same time, the data packet could possibly be in the order of 1500 bytes long. That size will take more time to be clocked out a slow-speed interface than a fast LAN interface. Many delay-sensitive applications will need prioritization. Fortunately, both point-to-point and Frame Relay networks can support fragmentation techniques that will allow smaller packets and therefore keep latency low.

Using ACLs also allows you to keep certain types of traffic, such as unnecessary routing protocol traffic, off slow network links.

## Routing and Scalability

As discussed earlier in the chapter, a router provides connectivity between networks and broadcast domains. Routers forward packets based on network addresses rather than Media Access Control (MAC) addresses (which is how Layer 2 works). Internetworks are generally more scalable than flat-bridged networks, because routers summarize status by network number. Routers use protocols such as OSPF and Enhanced IGRP to exchange network status information.

### NOTE

When compared with STP, routing protocols have improved on the following issues. Convergence is achieved in a more acceptable timeframe. It allows for more load balancing and optimization of links by implementing metrics. Finally, it is more scalable because it maintains status and routing tables.

# Planning for the Future Growth of Your Company's Infrastructure

Okay, so you have secured funding with your stellar speech that made the chief financial officer (CFO) pull out the checkbook and hand you a blank check, Now what? A run for political office? A screen test in Hollywood? Nah, it's time to purchase networking equipment (and a small villa in the Swiss Alps).

If possible, err on the side of building out too much. Whereas this might be a cost concern, think about the loss of money due to downtime or insufficient resources. Also, there is the issue of future technologies that may be able to add value to the network. Raise these points in allocation meetings, and discuss why more, in these instances, is necessary.

## Even More Network Scalability

Okay, you've designed this network and have taken into account that there would be more people added and more bandwidth being used for applications, so what happens when it's max'ed out? Can you expand on your existing design? Is your résumé printed and ready to go?

Here is where your design can be put to the test. Remember that scalability is dependent on what you have installed in the way of hardware, and what you are using at the software level (routing protocols). Scalability is usually limited by two factors: technical issues and operational issues. Technical issues with scaling are mainly about finding the right mix of routing protocols and network equipment. What you want are protocols that scale well with the addition of more network equipment. Operational issues, on the other hand, are mainly concerned with large areas and protocols that aren't based on the hierarchical design.

Remember that when designing your network, choosing the right equipment is key. Three resources must be taken into account for your decisions: central processing unit (CPU), memory, and bandwidth. The CPU utilization is dependent on protocols. Some of the protocols use the speed of the processor in their routing metrics so that they can choose the best path. Other protocols use the CPU to help with convergence (which is fairly processor intensive). A suggestion is to keep areas small and use route summarization when using link-state protocols. This reduces the convergence issues by keeping the number of routes that need to be recalculated to a minimum. Routing protocols use memory to store topology information and routing tables. Summarization eases the usage of memory for the same reasons as the CPU. Finally, there is bandwidth that, believe

it or not, is dependent on the protocol. With bandwidth, there are three items that you need to take into account:

- When the routing tables are sent?
- What those routing tables are sending?
- Where the information being sent to?

Distance routing protocols such as RIP, IGRP, SAP, and RTMP broadcast their complete routing tables on a periodic schedule. These updates will occur whether or not there have been any changes to the network. These replications happen anywhere from every ten seconds to every three minutes (sometimes this is dependent on what you set for the variable). These advertisements use up bandwidth, and if failures occur within the network, they may take a long time to come to convergence.

Link-state protocols such as OSPF and IS-IS were designed to improve on the limitations the distance vector routing protocols such as slow convergence and unnecessary usage of bandwidth. There are caveats to running these protocols, though; they require more CPU and memory usage. Enhanced IGRP is an advanced distance vector protocol that tries to be the best of both worlds. It does not suffer from standard distance vector issues, and only updates when there is a change in the network.

## Layer 2 Switching

Layer 2 switching is hardware-based bridging. In particular, the frame forwarding is handled by hardware, usually application-specific integrated circuits (ASICs). This layer is handled by most of the major vendors on the market. Let's face it; if they couldn't do this, there would be no reason for them to do the higher stuff. There are exceptions, of course, but they are rare.

## Layer 3 Switching

Layer 3 switching is hardware-based routing. The packet forwarding is handled by hardware, usually ASICs. Depending on the protocols, interfaces, and features supported, Layer 3 switches can be used in place of routers in a campus design (for this reason, I will sometimes refer to a router as a Layer 3 switch). Layer 3 switches that support standards-based packet header rewrite and time-to-live (TTL) decrements are called packet-by-packet Layer 3 switches.

High-performance packet-by-packet Layer 3 switching is achieved in different ways. The Cisco Gigabit Switch Router (GSR) series achieves wire-speed Layer 3 switching with a method called *crossbar switch matrix*. The Catalyst series of multilayer switches performs Layer 3 switching with ASICs that are located in the Supervisor Engine. Regardless of the underlying technology or vendor, packet-by-packet Layer 3 switching works like a router to external networks.

Layer 3 switching on the Catalyst series of switches combines multiprotocol routing with hardware-based Layer 3 switching. The Route Switch Module (RSM) and Multi-Layer Switch Feature Card (MSFC) are IOS-based routers with the same Reduced Instruction Set Computing (RISC) processor engine as the Cisco router family. The Layer 3 switching is also done with ASICs.

## Layer 4 Switching

Layer 4 switching is hardware-based routing that considers the application. Many vendors' routers have the ability to control traffic based on Layer 4 information using extended access lists and some form of accounting. In TCP or UDP traffic flow, a port number in the packet header is encoded for each application.

Many of today's vendors' switches can be configured to operate as a Layer 3 or Layer 4 switch. When operating as a Layer 3 switch, these devices will cache flows based on destination IP address. When operating as a Layer 4 switch, they will cache traffic based on source address, destination address, source port, and destination port.

Because of the ability to perform Layer 3 or Layer 4 switching in hardware, there is usually no performance difference between the two modes. Choose Layer 4 switching if you want your policy to dictate control of traffic by application, or if you require accounting of traffic by application.

## Bridged Protocol Needs

The great thing about the multilayer design is that addressing and routers are not dependent on media. The principles are the same whether the implementation occurs on FDDI, token ring, Ethernet, or ATM. This is not always true in the case of bridged protocols such as NetBIOS and Systems Network Architecture (SNA), which depend on the media type.

Cisco has implemented Data-Link Switching Plus (DLSw+) in their systems, which is an updated version of standard DLSw. This allows SNA frames from native SNA clients, which are then encapsulated in TCP/IP by a router. A second router de-encapsulates the SNA traffic. Using DLSw+ will allow you to use

multiple media types; for example, you can translate the traffic out to a token ring-attached front-end processor (FEP) at a centralized area on the network. Multilayer switches can be attached to different media types with versatile interface processor (VIP) cards and port adapters (PA).

## Bridging in the Multilayer Model

When using nonrouted protocols, such as NetBIOS, bridging must be configured. The Layer 3 switch handles bridging between VLANs on the Access layer and the Core. Remember that if you are using VLANs and are running spanning tree, the Layer 3 switch cannot be configured with a bridge group. The reason is that allowing bridging on the Layer 3 switch collapses all the spanning trees from the VLANs into a single spanning tree and a single root bridge.

## Security in the Multilayer Model

Security in the data center can be handled in several ways. A common security measure is to use access control lists (ACLs). Multilayer switching supports ACLs with little or no performance degradation. The best place to implement the ACL is at the Distribution layer, because at the Core and Access layers, you want high-speed switching; in addition, all traffic must pass through the Distribution layer. The great thing about ACLs is that they can be used to control the network by restricting access to the switches.

You could also implement additional security by using TACACS+ and RADIUS, which will provide centralized access control to switches. The Cisco software itself will also provide security, as it can assign multiple levels of authorization by password. This is much like using root-level or administrator-level access where people who manage the network can be assigned a password that will allow them access to certain sets of commands.

Using Layer 2 switches at the Access layer and in the data centers also has security benefits. When using bridges or other shared media networking equipment, all traffic is visible to all other connected clients on the local network. This could allow a user to capture clear-text passwords or files with a sniffer program. By implementing switches, packets are only visible to the sender and receiver.

Security on the WAN is usually taken care of with firewalls, such as a Cisco PIX. A firewall is implemented at your border to the customer, where routers are attached between outside connections and the servers in the ASP.

# High-Availability Design

Have you ever had a network connection just drop? This is usually due to either a hardware failure or the network connection going down. Wherever users could lose their connections to the ASP, say, in the event of a power failure or if WAN links to the ASP go down, are known as points of failure.

Technologies were designed to deal with these points of failure. The most common features that should be incorporated into most designs is high availability.

## NOTE

---

There are instances in which high availability is both unnecessary and costly.

---

## High Availability

Availability is the measurement of the uptime of database servers, mainframe applications, email, World Wide Web, multimedia, VoIP, and ERP (Enterprise Resource Planning). It also takes into account the network resources that are used to reliably transport those applications .

Network availability uses metrics that include the amount of redundancy that is built into the infrastructure and network device hardware. Most vendors have added technology to network devices that help with their resiliency. These additions include dual power supplies, hot-swappable modules, and redundant switching fabrics within a chassis. This can help to resolve issues that are isolated to one device, but it does not address the need for network availability as a systemwide entity.

## Things to Consider When Implementing High-Availability

Clearly, for high availability to be effective, it needs to be an end-to-end network solution. Well-designed network architectures, with specific fail-over and load-balancing capabilities at key places in the infrastructure to ensure network resiliency and stability:



- **Access layer** This is the initial connection into the network for the client.
  - There should be multiple paths (also known as dual homing) into the Distribution and Core layers of the network.
  - There should be load-balanced connectivity for optimal bandwidth utilization.
  - There should be fast convergence at Layer 2 and Layer 3 to recover from network failures so that there is minimal disruption and downtime.
- **Distribution layer** These are the aggregation points for the connectivity between Access layers, server farms, and WAN/LAN.
  - Load balancing between connections should be implemented for optimal bandwidth utilization.
  - There should be fast convergence, hopefully at Layer 3, so that your network can recover from failures with minimal disruption and downtime.
  - There should be redundant Layer 3 paths to help optimize network Layer 3 convergence.
- **Core layer** These are the convergence points for the Distribution layer.
  - Load balancing between connections should be implemented for optimal bandwidth utilization.
  - There should be fast convergence, hopefully at Layer 2, so that your network can recover from failures with minimal disruption and downtime.
  - There should be fail-over Layer 2 paths to help optimize network Layer 2 convergence associated with spanning tree.

When designing redundancy into networks, your objective should be to build an infrastructure that can deal with component, link, power, and any other types of failures. The network should be able to converge around these failures, fixing itself with little or no intervention and with minimal disruption or outages of service. Whereas you want it to be a sophisticated solution, it should remain simple enough that minimal configuration, monitoring, and management is necessary. A byproduct of the components being able to provide this level of redundancy may also be able to provide additional benefits to the network, such as load balancing.

If you can implement the redundant links that connect the servers to a pair of Catalyst multilayer switches in the data center, fail-over at the router (or Layer 3) can be achieved with Cisco System's Hot Standby Router Protocol (HSRP) or the generic VRRP (Virtual Router Redundancy Protocol). The data center switches should provide HSRP gateway routers for all servers on the segment. Fast fail-over at Layer 2 is achieved by using Cisco's Uplink Fast feature. With UplinkFast, fail-over takes about three seconds for convergence from the primary link to the backup link, as opposed to conventional STP, where convergence would take 40 to 50 seconds. Again many of the current vendors that are out there have some sort of similar features that are capable of doing similar things using VRRP.

Redundancy in the core can be achieved by installing two or more switches (or multilayer switches) in the backbone. Redundant links from the Distribution layer can provide fail-over and load balancing over multiple paths to the customer, depending on the routing.

The network should be designed so that it can notify network operations personnel if there are failures, and be able to provide enough detail of the events that led up to the failure so that you can isolate and fix the issues. This information is also useful for the prevention and forecasting of future outages. Cisco Systems provides tools and management applications such as Cisco Resource Manager (CRM), Network Time Protocol (NTP), Syslog, debug, and various other tools that are embedded within the products to handle the monitoring and management of your network infrastructure.

## NOTE

---

Cisco IOS software supports load balancing over up to six equal-cost paths for IP and over many paths for other protocols.

---

## Summary

With all these factors taken into consideration, you can probably understand why this area of networking is a science all to itself (there may be some black magic involved in there as well). With a little planning and a lot of foresight, your networks should provide stability and efficiency for you and your company.

We began the chapter by sitting down and drawing the network at a concept level, trying to keep things at the 30,000-foot view so that we are able to encompass future growth issues. Remember that the network must start somewhere, and this is always a good place to begin. We talked about the data center model, and how it should relate to the overall picture. You must remember to consider things such as customer access and satisfaction if you want to build your network correctly.

In the site considerations portion, we discussed things that should be taken into account for the physical design and layout of the network. Things such as environment, electricity, and weight concerns will affect the growth of the network, so positioning of the equipment is a very important part of the design. We also read a small overview of networking terminology in the area called networking equipment basics; this is not very in depth, but it helps to understand the process.

Capacity planning tries to help you with growth of that can be planned for, and some for which you cannot plan. The general idea is to think big, and plan your network accordingly. We reviewed some best practices that should be implemented on the network; again, it is not a complete list, and your needs may vary.

Routing protocols and how they relate to the network are a major concern to the design of a network. I have tried to stress that at all times in this chapter, from the selection of the interior protocols and how they are affected by convergence, to choosing the correct protocol. We also talked a little about redundancy by using route selection, and how it allows for bandwidth dedication.

We discussed address considerations and how they can affect all areas of the network. This can be combined with topology to create stable and efficient networks that are very secure.

IP Multicast is a growing part of the new network and must be taken into account for design considerations. With the use of video to the desktop and other corporate meeting software, you need to be aware of the impact that this will have on the network.

We discussed VLANs, ELANs, and policy in the core; these are ways to segment the traffic that exists on the network to improve efficiency and stability and allow greater security.

We touched on the Layer 3 and Layer 4 switch model, and where you would implement it, as well as how it may be best utilized.

In the WAN link considerations section, we discussed QoS and how it affects the implementation of the WAN router and bandwidth provisioning.

Planning for future growth and network scalability revisited. (This is where I went more in depth on the different layers of multilayer switching, and how they could be used on the network.)

Security in the multilayer model can be handled in various ways. Most were covered in other areas of the chapter, but this is the area that explains access control lists and their ability to help with security and bandwidth concerns.

Reliability and redundancy were also covered throughout the chapter, but this section discusses where and when to deploy VRRP (or HSRP).

## Solutions Fast Track

### Design Considerations

- ☑ There are generally three components when designing a large internet-network: data center networks, wide area networks (WAN), and remote users (in this case, your external clients).
- ☑ The *data center* is a building or set of buildings that house the infrastructure of your network.

### Site Considerations

- ☑ When you are building a new physical plant for your ASP, make sure that there is adequate space available and sufficient resources (power and cabling, as well as security) to suit your needs.
- ☑ *Routers* are Layer 3 network devices that connect separate networks and pass traffic between subnets.

## Designing with the Hierarchy in Mind

- ☑ One of the most beneficial tasks that you can perform in the design of your network is to create a hierarchical internetwork design that will modularize the elements of a large internetwork into layers of internetworking.
- ☑ Hierarchical internetworks are more scalable, because they allow you to grow your internetwork in a gradual way with the implementation of modules.
- ☑ The effect of broadcast traffic in your internetworks requires that you implement smaller groups of routers and switches, which will make your network more efficient.

## Frame Relay Internetwork Design Considerations

- ☑ A major concern when designing a Frame Relay implementation is scalability. As the number of remote clients and their links grows, your network must be able to grow to accommodate these growth spurts.
- ☑ Implementing a hierarchical mesh for Frame Relay environments can assist you in avoiding implementing an excessively large number of DLCIs.
- ☑ The cost-effective and strategic significance of the core network often forces network designers to implement a hybrid-meshed network for their WAN internetworks.

## Capacity Planning for Your Infrastructure

- ☑ If you have a general idea of where you stand for number of servers and expected growth, you can use those as a baseline for the capacity of your network.
- ☑ One of the best practices for planning is to map out where the different customer areas are located, and what the server count is going to be. Once these figures are determined, decide if the servers need one data link or multiple connections.

## Protocol Planning Concerns

- ☑ By determining the physical layout of the network, you will be able to map the correct topology and form a logical addressing scheme that will grow as your network grows.
- ☑ If your network is fairly simple in terms of the topology and number of routers, a distance-vector protocol such as RIP or IGRP (discussed later in this chapter) could work fine. If you're running a multivendor network, RIP, RIPv2, IS-IS, and OSPF are common protocols across many vendors' router implementations.

## Addressing Considerations

- ☑ The topology of a network is defined by sets of routers and the networks to which they connect. Routing protocols can also establish a logical topology depending on implementation.
- ☑ Broadcast traffic sets a practical limit to the size of the broadcast domain. Managing and troubleshooting a bridged campus becomes harder as the number of users increases because it adds to the broadcast domain.

## Application and Network Services

- ☑ When designing the data center, you should build the network as a modular building block using multilayer switching.
- ☑ Note that when using the Hot Standby Router Protocol (HSRP) (Cisco specific) or Virtual Router Redundancy Protocol (VRRP), which can also add redundancy, you should consider implementing Fast EtherChannel so you can scale bandwidth from Fast Ethernet, and from Gigabit Ethernet to Gigabit EtherChannel.

## Application-Aware Networking

- ☑ ASPs who want to deploy their applications need to realize that their success of mission-critical applications over both the internal LAN and clientele WAN is achieved by defining network policies, which assist in the apportioning of network resources with business objectives.

- ☑ Admission control is provided by a mechanism that can reject or remove applications based on user-defined policies. For example, a client can define a policy to temporarily stop the transmission of email packets, so that the mission-critical applications can use the necessary resources.

## Scalability Considerations

- ☑ Fast EtherChannel provides more efficient utilization of bandwidth by multiplexing multiple VLANs over one trunk.
- ☑ When designing your network, avoid creating STP loops in the backbone. STP takes 40 to 50 seconds to converge and does not allow for load balancing across multiple paths. When using ATM for your backbone, use PNNI to handle load balancing.

## Multimedia Services

- ☑ According to a study by the Telecommunications Industry Association, the multimedia application market (such as video on demand, VoIP, etc.) is expected to reach \$16 billion in 2001.
- ☑ Many of the new multimedia applications that customers want, require IP multicast for proper operation. Any network communication that needs to transmit information to multiple clients can benefit from the efficiency of multicast technologies.

## Planning for the Future Growth of Your Company's Infrastructure

- ☑ Distance routing protocols such as RIP, IGRP, SAP, and RTMP broadcast their complete routing tables on a periodic schedule. These updates will occur whether or not there have been any changes to the network.
- ☑ Cisco has implemented Data-Link Switching Plus (DLSw+) in their systems, which is an updated version of standard DLSw. This allows SNA frames from native SNA clients, which are then encapsulated in TCP/IP by a router.

## High-Availability Design

- ☑ Availability is the measurement of the uptime of database servers, main-frame applications, email, World Wide Web, multimedia, VoIP, and ERP (Enterprise Resource Planning).
- ☑ The network should be designed so that it can notify network operations personnel if there are failures, and be able to provide enough detail of the events that led up to the failure so that you can isolate and fix the issues.

## Frequently Asked Questions

The following Frequently Asked Questions, answered by the authors of this book, are designed to both measure your understanding of the concepts presented in this chapter and to assist you with real-life implementation of these concepts. To have your questions about this chapter answered by the author, browse to [www.syngress.com/solutions](http://www.syngress.com/solutions) and click on the “Ask the Author” form.

**Q:** How many DLCIs can be configured per serial port?

**A:** It varies depending on the traffic level. You can use approximately 1000 per interface, but in common use, 200–300 is a typical maximum. If you broadcast on the DLCIs, 30 to 50 are more realistic numbers, due to CPU overhead in generating broadcasts. There are no hard-and-fast specific numbers, as overhead will vary by configuration. However, on low-end boxes, the architecture is bound by the available I/O memory.

**Q:** I have implemented a star topology. My edge nodes are able to communicate with the core (hub) router, but not with other edge nodes. What could be the problem?

**A:** This problem commonly arises when you run into spanning tree issues. When a router learns a route to a node, it will block broadcasts from continuing on to other routes on that same main interface. These paths may actually be the correct way to get to another router, but the interface is blocked to prevent routing loops. One of the drawbacks to a star topology (sometimes known as a hub and spoke) is that you are generally going to create subinterfaces on



your ports. When a route is learned on the interface, it will block that main interface for broadcasts that came from within its subinterface.

For example, let's say that you have a serial interface that is connected in a point-to-multipoint configuration. If the subinterface s0/0.1 wants to contact a router that is connected to the same main interface on the core router, once a broadcast message is sent, the core router will block that main interface from rebroadcasting that route. What this means is that s0/0.2 would not receive this broadcast.

**Q:** I have agreed to a Committed Interface Rate (CIR) with a client. How do I charge them if they “burst” over this agreement?

**A:** This was covered in Chapter 7, but most ISP/ASP companies charge what is called the 95 percentile. This is where traffic is measured in timed intervals. You will be charging the customer for whatever their usage is for 95 percent of the time.

## Sample Configuration for an Application Service Provider Network

### Solutions in this chapter:

- The Test Network
- Configuration with Cisco Systems Commands and References

# Introduction

This appendix contains a sample network and configurations to give you a feel of what is involved in the setup of an application service provider (ASP). We used Cisco Systems ([www.cisco.com](http://www.cisco.com)) equipment, as they have the largest market share for networking equipment in this arena.

Many people consider Cisco Systems equipment to be an enterprise-based network component, but they have a proven record and above-average end-to-end solutions. As stated earlier, they have the largest market share within the service provider space. However, several large providers use other equipment vendors such as the following:

- **Juniper Networks** ([www.juniper.net](http://www.juniper.net)) makes some of the fastest and most efficient performing network devices available today. In fact, they have taken a large share of the core market away from Cisco. Their products are mainly design for the core and are capable of delivering high performance and throughput. Their M class of core routers is rated among the best in the business, and they are trying to expand out of the core market into voice, data, distribution, and access. They have a solid command-line interface (CLI) that spans all of their platforms.
- **Extreme Networks** ([www.extremenetworks.com](http://www.extremenetworks.com)) has extremely (no pun intended) fast internetworking equipment that can be implemented from the core to the Access layer. Extreme Networks Equipment is considered very cost conscious and is able to give a good return on investment. Their CLI is similar to the feel of Cisco's CLI, and as such, it is very easy to port your Cisco knowledge to this platform. One of their largest clients is the United States Pentagon.
- **Foundry Networks** ([www.foundrynetworks.com](http://www.foundrynetworks.com)) is in the same category as Extreme networks. They, like Juniper, offer a consistent command-line interface across the breadth of their equipment. They have been extensively used in several large networks and ISPs such as Mindspring and America OnLine.
- **Nortel Networks** ([www.nortelnetworks.com](http://www.nortelnetworks.com)) provides high-speed optical network devices that can be implemented in the core. They are considered one of the pioneers of the optical market (along with Fore/Marconi). AT&T Latin America currently is installing their equipment within their core to provide a high-speed infrastructure and more services.

These are not the only vendors in this category; there are many others (too numerous to mention really) from which to choose. You should research what functions and abilities you are looking for, and then design the network with that equipment.

## The Test Network

The following is an implementation plan that we put together to assist with some basics of design and implementation for an ASP network. These configurations have many of the commands that you will see if you are using Cisco Systems for your network. These are not a comprehensive list, but they are a good general overview. I also did not include every piece of equipment that is shown in the figures, but in the figures I have highlighted those commands that would make them work within the infrastructure. The following figures will give you an overview of what we are talking about in the rest of this appendix. There is the logical “30,000 foot view,” the access, the distribution (Internet), and the core (head-end).

## The Logical Network Overview

A network, in its most basic form can be considered something akin to a complex plumbing and electrical system. The reason that I say this is, like a complex plumbing job, you want to design your network to allow information to flow from one point to another with as little impediment as possible. Again, on the most basic level a plumber tries to implement your plumbing so that there is good flow, with no trouble areas.

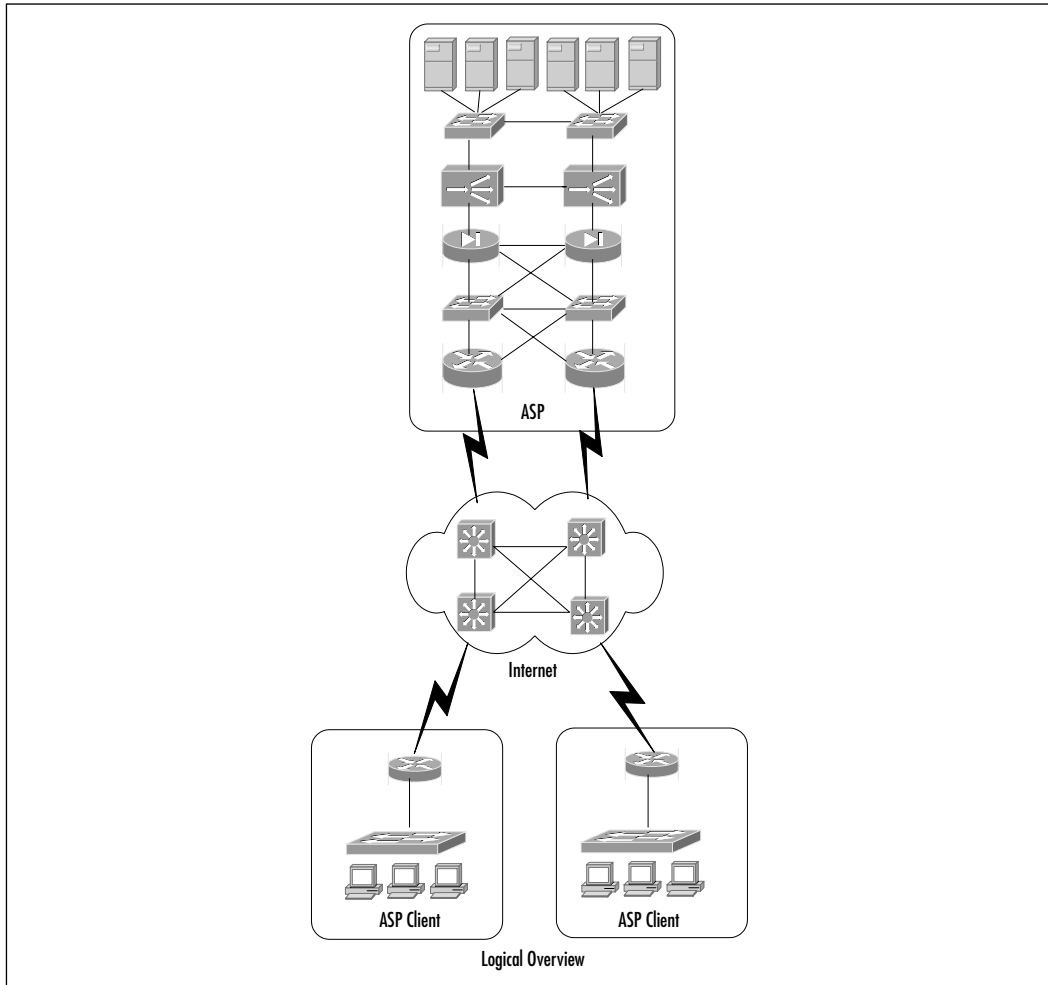
When you draw up a logical network diagram, you should look for potential issues before you get too far into the implementation. Figure A.1 is a basic overview of the network that I talk about in this configuration appendix.

As you can see, several types of equipment are installed within this infrastructure. This is only a logical view, so it is simplified as to what equipment is used, where it is located, and how you get from your content from the ASP to the client.

## The Access Layer

The *Access layer* is one of the areas over which you will normally have little control. This area is usually located at the client site, and therefore is out of your area of influence. Figure A.2 has a switch that is connected to a cache engine and the client access links. When applications or content are requested, the traffic will flow to the switch, and then either accesses the cache engine, or goes out to the Internet and pulls the information back to the client and cache engine.

**Figure A.1** The Logical Drawing of the Test Network from a “30,000 Foot View”

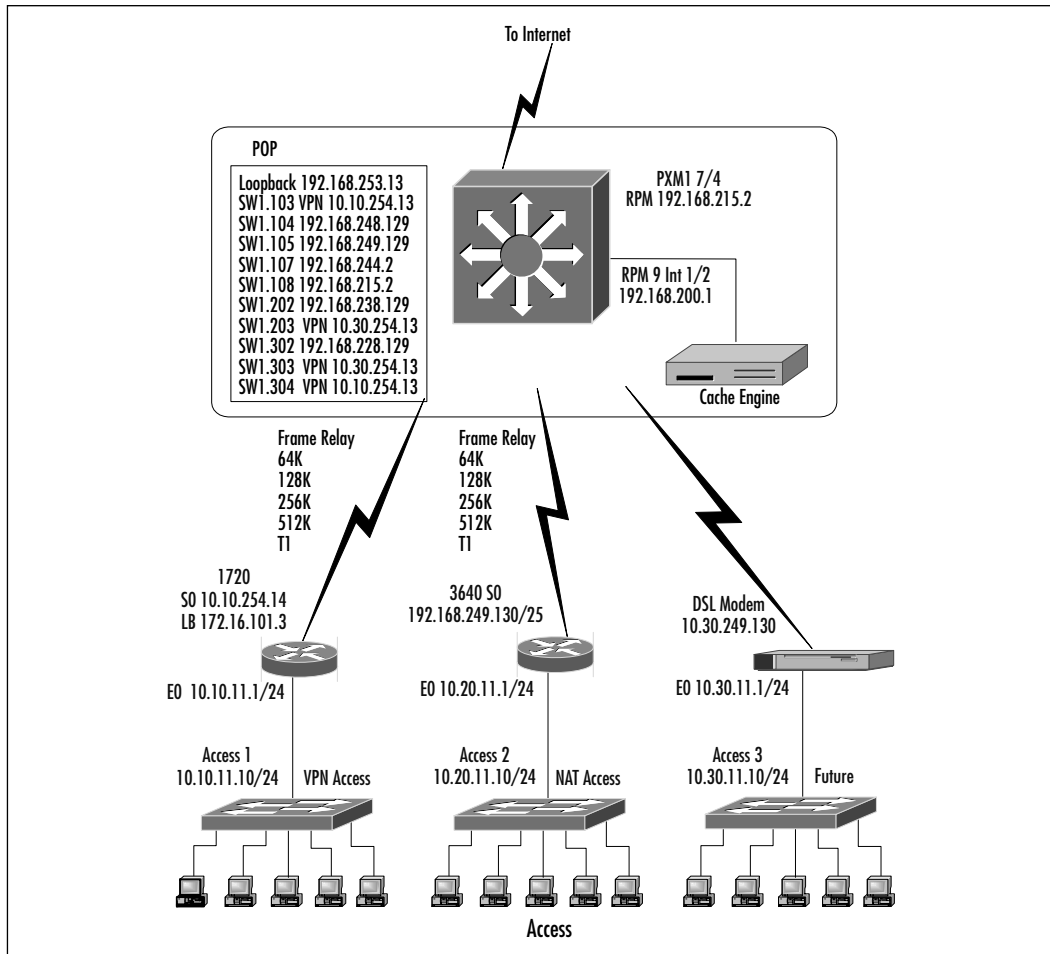


As you can see, the Access layer is comprised of clients that are located in topologically diverse areas. These Clients are then connected to switches and routers (layer 2 and 3) which are then connected to a Point-of-Presence (POP). This POP is then connected to the distribution (or Internet) layer

## The Distribution Layer

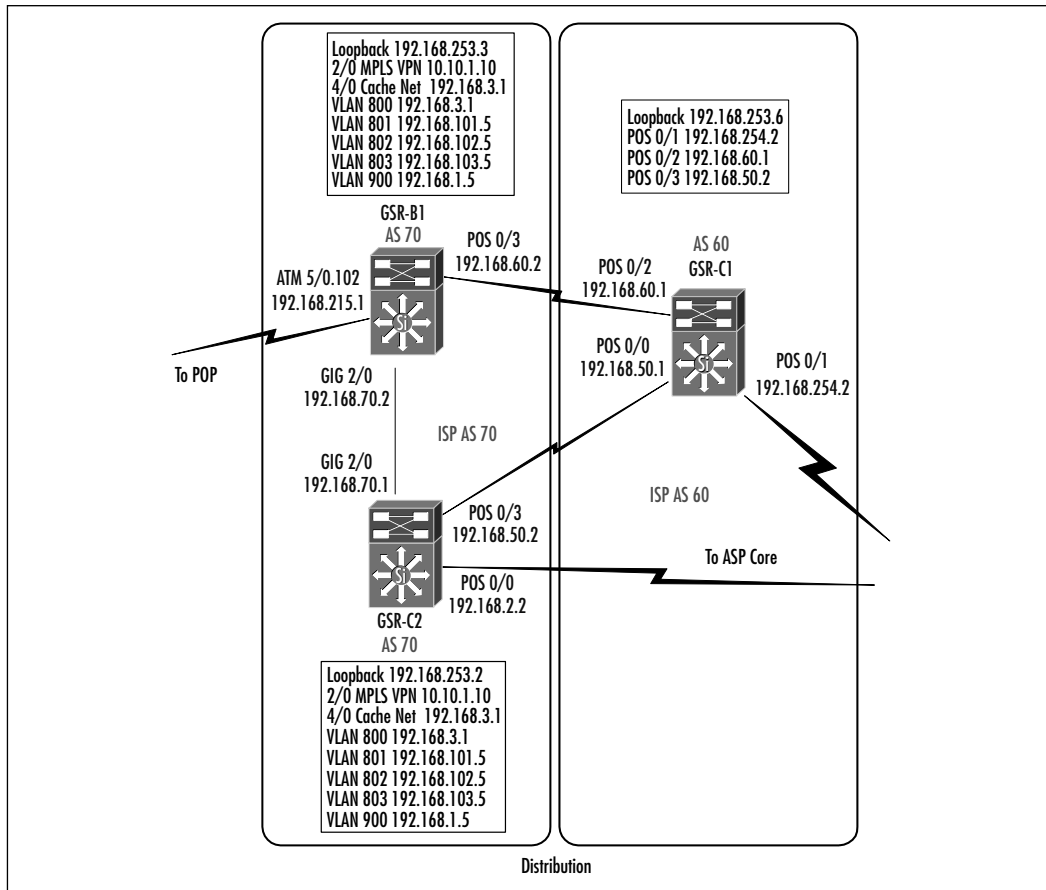
The *Distribution layer*, also known as the *Internet layer*, is the area that your application or contact must traverse to get to your clients. This area may or may not be

Figure A.2 The Access Layer of the ASP Test Network



under your influence or control. This is the area that most of your customers may know very little about, and you may need to contact the providers that are between you and your customers. Figure A.3 contains multiple autonomous systems (ASs) through which content must pass.

The Distribution layer is then connected to your ASP. Depending on the method of accessing your system, you may need to create VPN tunnels, or some other form of secure transfer transport.

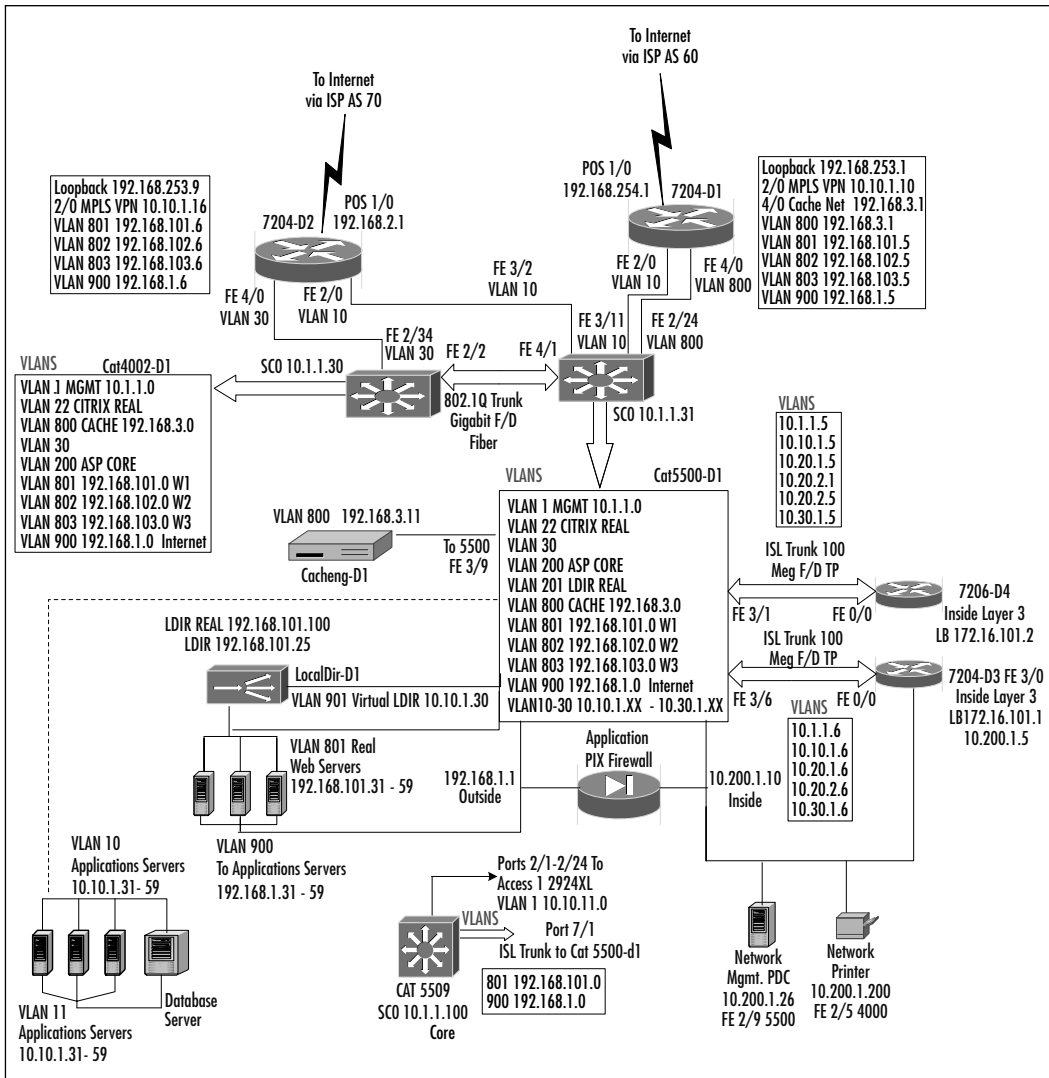
**Figure A.3** The Distribution (Internet) Layer

## The Core Layer

The *Core layer*, also known as the *Head-End layer*, is the area over which you should have the most control. This area is where your services and applications are stored and controlled. This is usually a very complex area (as you will see in Figure A.4), and requires a lot of design and discussion as to what needs to be deployed to make your ASP successful.

As you can see, this is a very complex area and will require a lot of thought before you get to the implementation. This is a sample network, so your network may be different.

Figure A.4 The Core (Head-End) Layer



## Configuration with Cisco Systems Commands and References

The following configurations are from our test network. I have tried to pick some of the more common or complex commands, as well as throw in a little information to some of the more basic commands. The configurations that I



have included are as complete as I could make them. Generally, I will only explain the command once within this appendix.

## Configuration for a Cisco Systems 7200 Router That Is Located within the Core Layer

The *service password-encryption* command tells the IOS software to encrypt passwords, such as CHAP secrets, and similar information, which are saved in the configuration file. This prevents people who are viewing the configuration from reading passwords; for example, if someone was to happen to look at the screen over your shoulder when you are looking at the configuration.

```
ASP1-DFT-7200-D1#show running-configuration
```

```
Building configuration...
```

```
Current configuration:
```

```
!
version 12.1
service timestamps debug uptime
service timestamps log uptime
service password-encryption
-----
```

```
ASP1-DFT-7200-D1(config)#service password-encryption
```

The algorithm that is implemented by *service password-encryption* is a simple Vigenere cipher. It can be cracked in a short amount of time by any competent cryptographer. The algorithm was not created to protect configuration files against serious analysis, and should not be used as the only security on the router. Cisco configuration files that contain encrypted passwords should therefore be treated as clear text if someone really wants to get past them.

This encryption does not apply to passwords that are implemented with the *enable secret* command, but it does work with passwords that are created with the *enable password* command.

```
-----
hostname ASP1-DFT-7200-D1
-----
```

```
ASP1-DFT-7200-D1(config)#hostname ASP1-DFT-7200-D1
```

This command sets the host name for the router. It is entered in global configuration mode and is used to set the system name that appears in the prompt. The prompt itself can be changed with the *prompt* command.

```
-----
boot system slot0:c7200-jk2o3s-mz_121-1_E.bin
-----
```

This is the image from which the router will boot, and the location in which it is stored. In this instance, the location of the file is in memory that is located in slot 0 and called c7200-jk2o3s-mz\_121-1\_E.bin. This binary file is stored in memory and decompressed when it is run.

```
-----
enable secret 5 $1$ShLc$HBF2vRWSEkd/GqQCI2.Ni0
enable password 7 08004257061700573305150B242E
-----
```

```
ASP1-DFT-7200-D1(config)#enable secret ThatsRight!
```

```
ASP1-DFT-7200-D1(config)#enable password Anyone Anyone
```

The *enable secret* command uses Message Digest version 5 (MD-5) for password encryption hashing. This algorithm is highly secure; in fact, it is considered nonreversible as far as anybody at Cisco knows. It is still possible to bypass this password by using a *dictionary attack* (a dictionary attack is when a hacker or cracker has a computer application that will try every word in a dictionary or any other list of possible passwords). You must keep your configuration files out of the hands of people whom you do not trust. You can find more information about password encryption on Cisco's Web site at [www.cisco.com/warp/public/701/64.html](http://www.cisco.com/warp/public/701/64.html).

```
-----
class-map match-all ASP1_4
  description Identify File Transfer Protocol Traffic for ASP1
  match protocol ftp
  match source-address mac 0090.278A.EAB5
```

```

class-map match-all ASP2_4
  description Identify File Transfer Protocol Traffic for ASP2
  match protocol ftp
class-map match-all ASP2_3
  description Identify Joint Photographic Experts Group Traffic for ASP2
  match protocol http mime jpeg
class-map match-all ASP1_1
  description Identify Citrix for ASP1
  match protocol citrix
class-map match-all ASP2_2
  description Identify Web 1 Class for ASP2
  match protocol http
class-map match-all ASP1_2
  description Identify Web 1 Class for ASP1
  match protocol http
class-map match-all ASP2_1
  description Identify Citrix Class for ASP2
  match protocol citrix
class-map match-all ASP1_3
  description Identify Joint Photographic Experts Group Traffic for ASP1
  match protocol http mime jpeg
class-map match-any OverHead_08
  description Identify all Overhead Protocols that need Bandwidth
  match protocol bgp
  match protocol arp
  match protocol dns
  match protocol dhcp
  match protocol tftp
  match protocol telnet
  match protocol icmp
!
```

-----

Two commands are usually implemented here, the *class-map match-any* command and the *class-map match-all* command. The match-any and match-all options

are able to determine how packets are evaluated when they meet multiple match criteria. Traffic must either meet all of the match criteria (match-all), or one of the match criteria (match-any) to be considered a part of that traffic class definition.

The following example shows you how to configure traffic classes with the *class-map match-all* command.

```
ASP1-DFT-7200-D1 (config) #class-map match-all ASP1_4
```

```
ASP1-DFT-7200-D1 (config-cmap) #description Identify File Transfer  
Protocol Traffic for ASP1
```

```
ASP1-DFT-7200-D1 (config-cmap) #match protocol ftp
```

```
ASP1-DFT-7200-D1 (config-cmap) #match source-address mac 0090.278A.EAB5
```

If a packet arrives on a router that is intended for class ASP1\_4, a filter may be configured on the interface, and the packet will then be evaluated to see if it matches the FTP protocol, and the source address of 0090.278a.eab5. If all of these match criteria are met, and the packet matches traffic class ASP1\_4, it will be filtered and classified as such.

The following example shows you how to configure traffic classes with the *class-map match-any* command.

```
ASP1-DFT-7200-D1 (config) #class-map match-any Overhead_08
```

```
ASP1-DFT-7200-D1 (config-cmap) #description Identify all Overhead  
Protocols that need Bandwidth
```

```
ASP1-DFT-7200-D1 (config-cmap) #match protocol bgp
```

```
ASP1-DFT-7200-D1 (config-cmap) #match protocol arp
```

```
ASP1-DFT-7200-D1 (config-cmap) #match protocol dns
```

```
ASP1-DFT-7200-D1 (config-cmap) #match protocol dhcp
```

```
ASP1-DFT-7200-D1 (config-cmap) #match protocol tftp
```

```
ASP1-DFT-7200-D1(config-cmap)#match protocol telnet
```

```
ASP1-DFT-7200-D1(config-cmap)#match protocol icmp
```

For traffic to be classified as `Overhead_08`, the criteria for the packets are evaluated in order until a successful match is located. The packet is first evaluated to see whether the BGP protocol can be used as a match. If BGP is a match, then the packet is classified as traffic class `Overhead_08`. If BGP is not a successful match, then the ARP protocol will be evaluated to see if it is a match—and so on, and so forth.

## Configuring & Implementing...

### The Difference between Match-All and Match-Any

Remember that the major difference between the two class-maps are that the `class-map match-all` command needs to have the entire match conditions met in order for the packet to be considered a member of the specified traffic class. In contrast, only one match must be met for the packet in the `class-map match-any` command to be defined as a member of the traffic class.

When a successful match happens, the packet will then be defined as a member of traffic class `Overhead_08`. If the packet does not match any of the specified conditions, the packet will then be classified as a member of the default class.

```
-----
policy-map POS_1/0
description Bandwidth Allocation for POS 1/0
class OverHead_08
bandwidth 2000
random-detect
police 2000000 10000 50000 conform-action set-dscp- transmit 8
exceed-action set-dscp-transmit 8
class ASP1_1
```

```
bandwidth 2000
random-detect
  police 2000000 10000 50000 conform-action set-dscp-transmit 18
  exceed-action set-dscp-transmit 22
class ASP1_2
bandwidth 1000
random-detect
  police 1000000 10000 50000 conform-action set-dscp-transmit 26
  exceed-action set-dscp-transmit 30
class ASP1_3
bandwidth 1000
random-detect
  police 1000000 10000 50000 conform-action set-dscp-transmit 34
  exceed-action set-dscp-transmit 38
class ASP1_4
bandwidth 1000
random-detect
  police 1000000 10000 50000 conform-action set-dscp-transmit 10
  exceed-action set-dscp-transmit 14
class ASP2_1
bandwidth 2000
random-detect
  police 2000000 10000 50000 conform-action set-dscp-transmit 18
  exceed-action set-dscp-transmit 22
class ASP2_2
bandwidth 1000
random-detect
  police 1000000 10000 50000 conform-action set-dscp-transmit 26
  exceed-action set-dscp-transmit 30
class ASP2_3
bandwidth 1000
random-detect
  police 1000000 10000 50000 conform-action set-dscp-transmit 34
  exceed-action set-dscp-transmit 38
```

```

class ASP2_4
  bandwidth 1000
  random-detect
    police 1000000 10000 50000 conform-action set-dscp-transmit 10
    exceed-action set-dscp-transmit 14
class class-default
  bandwidth 1000
  random-detect
    police 10000000 10000 20000 conform-action set-dscp-transmit 0
    exceed-action drop

```

-----

*PolicyMaps* create or modify a policy map. These maps can be attached to one or more interfaces to specify a service policy; use the *policy-map* global configuration command.

```
ASP1-DFT-7200-D1(config)# policy-map POS_1/0
```

```
ASP1-DFT-7200-D1(config-pmap)#description Bandwidth Allocation
for POS 1/0
```

```
ASP1-DFT-7200-D1(config-pmap)#class Overhead_08
```

```
ASP1-DFT-7200-D1(config-pmap-c)#bandwidth 2000 (Note: this is in kbps)
```

```
ASP1-DFT-7200-D1(config-pmap-c)#random-detect
```

```
ASP1-DFT-7200-D1(config-pmap-c)#police 2000000 10000 50000
conform-action set-dscp-transmit 8 exceed-action set-dscp-transmit 8
```

- **police** Watch and match traffic. Related to the *rate-limit* command.
- **2000000** Average rate in bits per second.
- **10000** Normal burst size in bytes.
- **50000** Excess burst size in bytes. (Note: In IOS release 12.1(5)T and later, the excess burst-size does not have to be specified unless the

*violate-action* option is also specified. In IOS releases 12.0(5)XE through 12.1(1)E, the excess burst size has to be specified.)

- ***conform-action*** Action to take on packets that conform to the rate limit.
- ***set-dscp-transmit*** Sets the differentiated services code point (DSCP) value and transmits the packet.
- ***exceed-action*** Action to take on packets that exceed the rate limit. (***violate-action***—Action to take on packets that violate the normal and maximum burst sizes.) (Note: This option is not available in IOS releases 12.0 XE or 12.1 E.)
- ***set-dscp-transmit*** Sets the DSCP value and transmits the packet.

```
-----  
ip subnet-zero
```

```
ASP1-DFT-7200-D1 (config)#ip subnet-zero
```

By entering the global configuration command *ip subnet-zero*, the subnet zero restriction is lifted and the zero subnet address can then be assigned to an interface, giving you more address space. However, it also makes troubleshooting more difficult.

Note: Prior to IOS version 12.0, Cisco routers didn't allow an IP address belonging to subnet zero to be configured on an interface, by default.

```
-----  
ip wccp web-cache  
-----
```

```
ASP1-DFT-7200-D1 (config)#ip wccp web-cache
```

This enables the Web Cache Communication Protocol (WCCP). WCCP allows you to use the Cisco cache engine to handle Web traffic. These cache engines help to reduce transmission costs and download time. The router will send a user's request to a cache engine; if the cache has a copy of the page in storage, it will send it to the user. Otherwise, the cache engine will retrieve the requested page and store a copy of that page and content, and then forward the page to the user.



```

-----
ip tftp source-interface Loopback1
ip domain-name dft.exn.com
ip name-server 192.168.1.11
-----

```

```
ASP1-DFT-7200-D1(config)#ip tftp source-interface Loopback1
```

This allows you to select the interface address that will be used as the source address for TFTP connections. A loopback interface is a software-based connection that can be configured for testing your router as well as an interface.

```
ASP1-DFT-7200-D1(config)#ip domain name dft.exn.com
```

You can specify the Domain Name System (DNS) to automatically determine host-name-to-address mappings. The drawback to this command is that if you mistype a command, the router will perform a domain name lookup for the item that you typed.

```
ASP1-DFT-7200-D1(config)#ip name-server 192.168.1.11
```

You can specify the name server to automatically determine host-name-to-address mappings.

```

-----
ip vrf ip-mpls1
rd 10.10.254.13:5
route-target export 10.10.254.13:5
route-target import 10.10.254.13:5
-----

```

```
ASP1-DFT-7200-D1(config)#ip vrf ip-mpls1
```

Enters VPN forwarding routing (VRF) configuration mode, and defines the VPN routing instance by assigning a VRF name.

```
ASP1-DFT-7200-D1(config-vrf)#rd 10.10.254.13:5
```

Creates routing and forwarding tables with the route distinguisher (RD).

```
ASP1-DFT-7200-D1(config-vrf)#route-target export 10.10.254.13:5
```

Creates a list of export route target communities for the specified VRF.

```
ASP1-DFT-7200-D1(config-vrf)#route-target export 10.10.254.13:5
```

Creates a list of import route target communities for the specified VRF.

```
-----
ip vrf lab1-access1
  rd 65535:1
  route-target export 65535:1
  route-target export 70:1
  route-target import 70:1
  route-target import 65535:1
ip cef
ip inspect name ASP1 realaudio timeout 30
ip inspect name ASP1 ftp timeout 3600
ip inspect name ASP1 smtp timeout 3600
ip inspect name ASP1 udp timeout 15
ip inspect name ASP1 tcp timeout 3600
ip inspect name ASP1 http
ip audit notify log
ip audit po max-events 100
mpls traffic-eng tunnels
frame-relay switching
mls rp ip
-----
```

```
ASP1-DFT-7200-D1(config)#ip cef
```

This command enables Cisco express forwarding (CEF). CEF is designed to accommodate changing network dynamics and traffic that results from increased numbers over a short period of time. These patterns are usually associated with Web-based applications and interactive applications.

```
ASP1-DFT-7200-D1(config)#ip inspect name ASP1 realaudio timeout 30
```

Use the *ip inspect name in* global configuration command to define a set of inspection rules to which packet traffic must adhere.

```
ASP1-DFT-7200-D1(config)#ip audit notify log
```

Use the *ip audit notify log* command in global configuration mode to specify the method of event notification, so that you can view these notifications and tweak your network for better efficiency.

```
ASP1-DFT-7200-D1(config)#ip audit po max-events 100
```

Use the *ip audit po local* command in global configuration mode to specify the local post office parameters that should be used when sending event notifications to your network administrator.

```
ASP1-DFT-7200-D1(config)#mpls traffic-eng tunnels
```

The *mpls traffic-eng tunnels* command enables multiprotocol label switching (MPLS) traffic engineering tunnel signaling on a device.

```
ASP1-DFT-7200-D1(config)#frame-relay switching
```

Enables Frame-Relay switching.

```
ASP1-DFT-7200-D1(config)#mls rp ip
```

Globally enables IP multilayer switching (MLS) on the router.

```
-----
cns event-service server
-----
```

Cisco Networking Services Management Server provides infrastructure elements that can enable end-to-end management of your network.

```
-----
interface Loopback1
ip address 192.168.253.1 255.255.255.255
ip wccp web-cache redirect out
ip router isis
-----
```

```
ASP1-DFT-7200-D1(config)#interface loopback 1
```

This command creates loopback interface 1.

```
ASP1-DFT-7200-D1(config-if)#ip address 198.168.253.1 255.255.255.255
```

This command configures an IP address for the interface.

```
ASP1-DFT-7200-D1(config-if)#ip wccp web-cache redirect out
```

This command configures an interface to enable a router to verify that the appropriate packets are being redirected to the cache engine.

```
ASP1-DFT-7200-D1(config-if)#ip router isis
```

This enables the Intermediate System-to-Intermediate System (IS-IS) routing protocol on the interface. This command also identifies the area in which the router will work, while letting the router know that it will be routing dynamically rather than statically.

```
-----  
interface FastEthernet0/0  
  no ip address  
  no ip redirects  
  ip nbar protocol-discovery  
  full-duplex  
  mls rp vtp-domain EXN_ASP_LAB  
  mls rp ip  
  mls rp ipx  
-----
```

```
ASP1-DFT-7200-D1(config)#interface FastEthernet 0/0
```

This command enables interface configuration mode for FastEthernet slot/port.

```
ASP1-DFT-7 200-D1(config-if)#no ip address
```

This is the default setting for the interface.

```
ASP1-DFT-7200-D1(config-if)#no ip redirects
```

This is the default setting for the interface.

```
ASP1-DFT-7200-D1(config-if)#full-duplex
```

Enables full-duplex on the interface. This will allow the interface to send and receive data traffic at the same time.

```
ASP1-DFT-7200-D1(config-if)#mls rp vtp-domain EXN_ASP_LAB
```

Configures virtual local area network (VLAN) Trunking Protocol (VTP) domain. VTP allows you to make configuration changes centrally on a single

network device, and have those changes automatically communicated to all the other devices within the domain.

```
ASP1-DFT-7200-D1(config-if)#mls rp ipx
```

This command enables Internetwork Packet eXchange (IPX) multilayer switching on the router interface.

```
-----  
interface FastEthernet0/0.1  
no ip redirects  
-----
```

```
ASP1-DFT-7200-D1(config)#interface FastEthernet 0/0.1
```

Creates, enables, and enters configuration mode for a subinterface on a FastEthernet slot/port.

```
-----  
interface FastEthernet0/0.2  
encapsulation isl 900  
ip address 192.168.1.5 255.255.255.0  
no ip redirects  
ip wccp web-cache redirect out  
ip nbar protocol-discovery  
ip router isis  
tag-switching ip  
mls rp management-interface  
mls rp ip  
mls rp ipx  
standby 2 priority 100 preempt delay 120  
standby 2 ip 192.168.1.2  
standby 2 track POS1/0  
-----
```

```
ASP1-DFT-7200-D1(config)#interface fastethernet 0/0.2
```

Creates, enables, and enters configuration mode for a subinterface on a FastEthernet slot/port.

```
ASP1-DFT-7200-D1(config-if)#encapsulation isl 900
```

Creates inter-switch link (ISL) VLAN encapsulation on the interface. ISL is a Cisco-specific VLAN encapsulation method.

```
ASP1-DFT-7200-D1(config-if)#ip nbar protocol-discovery
```

Enables Network-Based Application Recognition Protocol-Discovery (NBAR). NBAR dynamically recognizes applications and employs network services to attain end-to-end availability, performance, and security.

```
ASP1-DFT-7200-D1(config-if)#tag-switching ip
```

Enables packet forwarding to go across cell-based devices that are connected to the interface. Tag switching was created to resolve the challenges that face an evolving Internet and high-speed data communications in general. Tag switching uses two main components: forwarding and control. Forwarding uses the tag information that is carried by packets, and tag-forwarding information, which is handled by a tag switch that executes packet forwarding. Control is in charge of retaining the correct tag-forwarding information for a group of connected tag switches.

```
ASP1-DFT-7200-D1(config-if)#mls rp management-interface
```

This command specifies an interface as the management interface for MLS.

```
ASP1-DFT-7200-D1(config-if)#standby 2 priority 100 preempt delay 120
```

Configures HSRP priority and sets the preempt delay.

```
ASP1-DFT-7200-D1(config-if)#standby 2 ip 192.168.1.2
```

Sets the IP address for the standby unit.

```
ASP1-DFT-7200-D1(config-if)#standby 2 track POS1/0
```

Configures the interface so that the HSRP priority can change based on the availability of other interfaces.

```
-----  
interface FastEthernet0/0.801  
  encapsulation isl 801  
  ip address 192.168.101.5 255.255.255.0
```

```
no ip redirects
ip wccp web-cache redirect out
ip nbar protocol-discovery
ip router isis
tag-switching ip
mls rp ip
standby 101 priority 100 preempt delay 120
standby 101 ip 192.168.101.1
standby 101 track POS1/0
!
interface FastEthernet0/0.802
encapsulation isl 802
ip address 192.168.102.5 255.255.255.0
no ip redirects
ip wccp web-cache redirect out
ip nbar protocol-discovery
ip router isis
tag-switching ip
mls rp ip
standby 102 priority 50
standby 102 ip 192.168.102.1
standby 102 track POS1/0
!
interface FastEthernet0/0.803
encapsulation isl 803
ip address 192.168.103.5 255.255.255.0
ip helper-address 192.168.1.11
no ip redirects
ip wccp web-cache redirect out
ip router isis
tag-switching ip
mls rp ip
standby 103 priority 100
standby 103 ip 192.168.103.1
```

```

!
interface POS1/0
  ip address 192.168.254.1 255.255.255.0
  ip wccp web-cache redirect out
  no keepalive
  tag-switching mtu 1500
  tag-switching ip
  clock source internal

```

```

-----
ASP1-DFT-7200-D1(config-if)#no keepalive

```

The *keepalive* command specifies how many seconds of inactivity will elapse before it sends a transmission to another router.

```

ASP1-DFT-7200-D1(config-if)#tag-switching mtu 1500

```

This command sets the maximum transmission unit (MTU) for tag-switching packets to 1500 on this interface.

```

ASP1-DFT-7200-D1(config-if)#clock source internal

```

This command specifies that the interface will clock its data from its internal clock.

```

-----
interface FastEthernet2/0
  ip vrf forwarding lab1-access1
  ip address 10.10.1.10 255.255.255.0
  no ip redirects
  ip wccp web-cache redirect out
  ip nbar protocol-discovery
  no ip route-cache cef
  shutdown
  full-duplex
  tag-switching ip
  standby 11 preempt
!
interface Serial3/0

```



```

no ip address
shutdown
framing c-bit
cablelength 10
dsu bandwidth 44210

```

-----

```

ASP1-DFT-7200-D1(config-if)#framing c-bit

```

This specifies that the C-bit framing will be used as the framing type for this interface. This command frees up the C bits so that other traffic types can use them.

```

ASP1-DFT-7200-D1(config-if)#cablelength 10

```

This command specifies the distance of the cable from the interface processor to the network equipment.

```

ASP1-DFT-7200-D1(config-if)#dsu bandwidth 44210

```

This command specifies the maximum allowable bandwidth used by the port adapter. Maximum bandwidth is 22 kbps to 44736 kbps. The default varies for different port adapters.

```

-----
interface FastEthernet4/0
description CacheEngine Network
ip address 192.168.3.1 255.255.255.0
ip wccp web-cache redirect out
full-duplex
tag-switching ip
!
router isis
redistribute connected
net 49.0001.0000.0000.00d1.00

```

-----

```

ASP1-DFT-7200-D1(config-router)#redistribute connected

```

This command redistributes routes from one routing domain into another routing domain. The connected switch is the source protocol from which routes are being redistributed.

```
ASP1-DFT-7200-D1(config-if)#net 49.0001.0000.00d1.00
```

This command is used to configure an IS-IS network entity title (NET) for the routing process.

```
-----  
router rip  
  version 2  
-----
```

```
ASP1-DFT-7200-D1(config)#router rip
```

This enables RIP (Routing Information Protocol) for routing between network devices. RIP uses hop count as a routing metric.

```
ASP1-DFT-7200-D1(config-router)#version 2
```

This command enables RIP version 2. RIP v2 allows the router to pass subnet information.

```
-----  
address-family ipv4 vrf lab1-access1  
  version 2  
  network 10.0.0.0  
  no auto-summary  
  exit-address-family  
-----
```

```
ASP1-DFT-7200-D1(config-router)#address-family ipv4 vrf lab1-access1
```

To enter the address family submode for configuring routing protocols such as BGP, RIP, and static routing.

```
ASP1-DFT-7200-D1(config-router-af)#version 2
```

Listen for and use RIP v2 on this address family.

```
ASP1-DFT-7200-D1(config-router-af)#network 10.0.0.0
```

Sets the default network to 10.0.0.0 for this address family.

```
ASP1-DFT-7200-D1(config-router-af)#no auto-summary
```

Turns off VLSM (the default). This makes the router act classful for address allocation and subnetting.

```
ASP1-DFT-7200-D1(config-router-af)#exit-address-family
```

This command exits the address-family submode.

```
-----
router bgp 65535
no bgp default ipv4-unicast
network 192.168.1.0
network 192.168.101.0
network 192.168.102.0
network 192.168.253.1
network 192.168.254.0
neighbor 192.168.253.5 remote-as 70
neighbor 192.168.253.5 ebgp-multihop 255
neighbor 192.168.253.5 update-source Loopback1
neighbor 192.168.253.5 activate
neighbor 192.168.253.5 send-community both
neighbor 192.168.253.6 remote-as 60
neighbor 192.168.253.6 ebgp-multihop 255
neighbor 192.168.253.6 update-source Loopback1
neighbor 192.168.253.6 activate
default-information originate
-----
```

```
ASP1-DFT-7200-D1(config)#router bgp 65535
```

This command enables BGP (Border Gateway Protocol) on the router, and places the router in an AS group (65535).

```
ASP1-DFT-7200-D1(config-router)#no bgp default ipv4-unicast
```

When you use *neighbor remote-as*, routing information for IPv4 is advertised by default when you configure a BGP routing session. To remove these advertisements, you need to enter the *no bgp default ipv4-unicast* command.

```
ASP1-DFT-7200-D1(config-router)#network 192.168.1.0
```

This command is used to specify which networks are to be advertised by BGP.

```
ASP1-DFT-7200-D1(config-router)#neighbor 192.168.253.5 remote-as 70
```

This command adds an entry to the BGP neighbor table.

```
ASP1-DFT-7200-D1(config-router)#neighbor 192.168.253.5 ebgp-multihop 255
```

Attempts and accepts BGP connections to external peers that reside on networks that are not directly connected.

```
ASP1-DFT-7200-D1(config-router)#neighbor 192.168.253.5 update-source
Loopback1
```

This command allows internal BGP sessions to use any operational interface for TCP connections.

```
ASP1-DFT-7200-D1(config-router)#neighbor 192.168.253.5 activate
```

This command enables the exchange of information with a BGP neighboring router.

```
ASP1-DFT-7200-D1(config-router)#neighbor 192.168.253.5 send-community
both
```

This command specifies the “communities” attribute that is sent to a BGP neighbor.

```
ASP1-DFT-7200-D1(config-router)#default-information originate
```

This sets the originate network 0.0.0.0 into BGP.

```
-----
address-family ipv4 vrf lab1-access1
 redistribute rip metric 1
 neighbor 192.168.253.5 remote-as 70
 neighbor 192.168.253.5 ebgp-multihop 255
 neighbor 192.168.253.5 activate
 neighbor 192.168.253.5 send-community both
 no auto-summary
```

```
no synchronization
network 10.10.1.0 mask 255.255.255.0
exit-address-family
-----
```

```
ASP1-DFT-7200-D1(config-router-af)#redistribute rip metric 1
```

This redistributes RIP advertisements with a metric of 1.

```
ASP1-DFT-7200-D1(config-router-af)#no synchronization
```

This command disables synchronization, so that you carry fewer routes in your IGP and allow BGP to converge more quickly.

```
-----
address-family ipv4 vrf ip-mp1s1
  redistribute connected
  redistribute static
  redistribute rip metric 1
  default-information originate
  no auto-summary
  no synchronization
  exit-address-family
!
address-family vpnv4
  neighbor 192.168.253.5 activate
  neighbor 192.168.253.5 send-community both
  neighbor 192.168.253.6 activate
  neighbor 192.168.253.6 send-community both
  default-information originate
  network 10.10.1.0
  exit-address-family
!
ip nat pool ASP-1 192.168.2.5 192.168.2.10 netmask 255.255.255.0
ip nat inside source route-map internet_out pool ASP-1 overload
ip classless
ip route 0.0.0.0 0.0.0.0 192.168.254.2
```

```
ip route 192.168.253.6 255.255.255.255 POS1/0
no ip http server
ip bgp-community new-format
```

```
-----
ASP1-DFT-7200-D1(config-router-af)#address-family vpnv4
```

This command tells BGP that it should use standard VPNv4 address prefixes.

```
ASP1-DFT-7200-D1(config)#ip nat pool ASP-1 192.168.2.5 192.168.2.10
netmask 255.255.255.0
```

This command creates and groups a pool of network addresses for the router to use in its Network Address Translation (NAT) process.

```
ASP1-DFT-7200-D1(config)#ip nat inside source route-map internet_out
pool ASP-1 overload
```

This command will translate the inside interface packets from addresses that match those on the access list. These addresses are then allocated from the named pool that was created in the command above. The **overload** command (optional) enables port translation for UDP and TCP.

```
ASP1-DFT-7200-D1(config)#ip classless
```

This command enables classless routing behavior, which selects a best route for packets destined for networks unknown by the router. This is on by default.

```
ASP1-DFT-7200-D1(config)#ip route 0.0.0.0 0.0.0.0 192.168.254.2
```

This command enables a default route for IP-based traffic, and sets up a best route for packets destined for networks unknown by the router.

```
ASP1-DFT-7200-D1(config)#route 192.168.253.6 255.255.255.255 POS1/0
```

Creates a static mapping to POS1/0.

```
ASP1-DFT-7200-D1(config)#ip bgp-community new-format
```

This command configures the new community format, wherein the community number is displayed in the short form.

```
-----
map-class frame-relay 3600
logging source-interface Loopback1
```

```

logging 192.168.1.11
access-list 105 deny tcp any any
access-list 105 permit udp any any eq snmp
access-list 105 permit udp any any eq snmptrap
access-list 105 permit icmp any any echo-reply
access-list 105 deny udp any any
access-list 120 permit ip 10.0.0.0 0.255.255.255 any
access-list 120 permit ip 192.168.1.0 0.0.0.255 any
access-list 120 permit ip 192.168.3.0 0.0.0.255 any
route-map internet_out permit 10
  match ip address 120

```

```

-----
ASP1-DFT-7200-D1(config)#map-class frame-relay 3600

```

Specifies Frame-Relay map class name, and enters map class configuration mode.

```

ASP1-DFT-7200-D1(config-map-class)#logging source-interface Loopback1

```

Sets the source for logging to the loopback interface.

```

ASP1-DFT-7200-D1(config-map-class)#logging 192.168.1.11

```

Logs information to 192.168.1.11.

```

ASP1-DFT-7200-D1(config)#access list 105 deny tcp any any

```

Creates an access list that denies all TCP packets from any to any.

```

ASP1-DFT-7200-D1(config)#route-map internet_out permit 10

```

Route maps are used to control and modify routing information. It can also define the conditions by which routes are redistributed between routing domains.

```

ASP1-DFT-7200-D1(config)#match ip address 120

```

The *match* command specifies conditions that must correspond in order for the packet to be processed.

```

-----
snmp-server engineID local 00000009020000D0BC326400
snmp-server community public RO

```

```
snmp-server community private RW
```

```
-----
ASP1-DFT-7200-D1(config)#snmp-server engineID local
00000009020000D0BC326400
```

Specifies the local copy of SNMP on the router.

```
ASP1-DFT-7200-D1(config)#snmp-server community public RO
```

Allows for read-only access. Only authorized management stations are able to retrieve MIB objects.

```
ASP1-DFT-7200-D1(config)#snmp-server community private RW
```

Allows for read-write access. Authorized management stations are able to retrieve and modify MIB objects.

```
-----
line con 0
  exec-timeout 0 0
  transport input none
line aux 0
line vty 0 4
  password 7 08004257061700573305150B242E
  login
  transport input lat pad v120 mop telnet rlogin udptn nasi
line vty 5 15
  login
  transport input lat pad v120 mop telnet rlogin udptn nasi
!
end
```

## Configuration for a Cisco Systems Gigabit Switch Router Router That Is Located within the Distribution Layer

The following is the configuration for a Cisco Systems gigabit switch router (GSR) that is located in the Distribution layer.



```

ASP1-DFT-GSR-B1#show running-configuration
Using 7792 out of 520184 bytes
!
! Last configuration change at 03:34:08 PST Tue Dec 19 2000
! NVRAM config last updated at 06:20:57 PST Mon Feb 5 2001
!
version 12.0
no service pad
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
service password-encryption
!
hostname ASP1-DFT-GSR-B1
!
boot system slot0:gsr-p-mz_120-9_S.bin
enable secret 5 $1$ShLc$HBF2vRWSEkd/GqQCI2.Ni0
enable password 7 08004257061700573305150B242E
!
clock timezone PST -8
clock summer-time PDT recurring
clock calendar-valid
-----

```

```
ASP1-DFT-GSR-B1(config)#clock timezone PST -8
```

This sets the system clock time zone to Pacific Standard Time (−8 from Greenwich Mean Time (GMT) or Zulu Time).

```
ASP1-DFT-GSR-B1(config)#clock summer-time PDT recurring
```

This sets the system clock to acknowledge daylight-savings time.

```
ASP1-DFT-GSR-B1(config)#clock calendar-valid
```

This command is used to configure a router as a time source for a network based on its calendar.

```

-----
class-map match-all test

```

```

!
!
policy-map test
!
ip subnet-zero
ip cef accounting non-recursive
ip domain-name dft.exn.com
ip name-server 192.168.1.11
clns routing
-----

```

```
ASP1-DFT-GSR-B1(config)#ip cef accounting non-recursive
```

This command enables accounting through nonrecursive prefixes. For prefixes that are directly connected to their next hops, it enables the collection of the number of packets and bytes express forwarded through a prefix.

```
ASP1-DFT-GSR-B1(config)#clns routing
```

This command enables Connectionless Network Services (CLNS) routing.

```

-----
interface Loopback1
ip address 192.168.253.3 255.255.255.255
 ip directed-broadcast
 ip router isis
-----

```

```
ASP1-DFT-GSR-B1(config-int)#ip directed broadcast
```

The default setting for routers is to forward directed broadcasts. You can disable this with the *no ip directed broadcast* command.

```

-----
interface POS0/0
 ip address 192.168.250.129 255.255.255.128
 no ip directed-broadcast
 rate-limit output dscp 8 15000000 10000 20000 conform-action transmit
 exceed-action transmit
-----

```

```

rate-limit output dscp 10 15000000 10000 20000 conform-action transmit
    exceed-action transmit
rate-limit output dscp 14 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 18 15000000 10000 20000 conform-action transmit
    exceed-action transmit
rate-limit output dscp 22 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 26 15000000 10000 20000 conform-action transmit
    exceed-action transmit
rate-limit output dscp 30 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 34 15000000 10000 20000 conform-action transmit
    exceed-action transmit
rate-limit output dscp 38 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 0 150000000 75000 75000 conform-action transmit
    exceed-action drop
no ip mroute-cache
tag-switching ip
crc 16
clock source internal

```

```

-----
ASP1-DFT-GSR-B1(config-int)#rate-limit output dscp 8 15000000 10000
20000 conform-action transmit exceed action transmit

```

This command is very similar to the *police* command. This command applies this Committed Access Rate (CAR) policy to packets sent on this interface, and what actions are taken if those limits are exceeded.

```

ASP1-DFT-GSR-B1(config-int)#no ip mroute-cache

```

This command configures IP multicast fast switching or multicast distributed switching (MDS) on the interface.

```

ASP1-DFT-GSR-B1(config-int)#crc 16

```

This command enables you to set the length of the cyclic redundancy check (CRC) on a fast serial interface processor (FSIP) or HSSI interface processor (HIP) on a Cisco router.

```
-----  
interface POS0/1  
  no ip address  
  no ip directed-broadcast  
  no ip mroute-cache  
  no keepalive  
  shutdown  
  crc 16  
  no cdp enable  
-----
```

```
ASP1-DFT-GSR-B1 (config-int) #no cdp enable
```

Cisco Discover Protocol (CDP) is enabled by default. If you do not want to use the CDP device discovery capability, you would use the *no cdp enable* command.

```
-----  
interface POS0/2  
  no ip address  
  no ip directed-broadcast  
  no ip mroute-cache  
  no keepalive  
  shutdown  
  crc 16  
  no cdp enable  
-----
```

```
ASP1-DFT-GSR-B1 (config-int) #shutdown
```

This shuts the port down. Shutdown is the default for all interfaces. If you would like to use the interface, remember to type *no shutdown* when you are ready to use it. (Note: If you cut and paste a configuration to the router, the interfaces will come up in *shutdown* mode.)

```

-----
interface POS0/3
 ip address 192.168.60.2 255.255.255.0
 no ip directed-broadcast
 rate-limit output dscp 8 5000000 10000 20000 conform-action transmit
   exceed-action transmit
 rate-limit output dscp 10 5000000 10000 20000 conform-action transmit
   exceed-action transmit
 rate-limit output dscp 14 5000000 10000 20000 conform-action transmit
   exceed-action drop
 rate-limit output dscp 18 5000000 10000 20000 conform-action transmit
   exceed-action transmit
 rate-limit output dscp 22 5000000 10000 20000 conform-action transmit
   exceed-action drop
 rate-limit output dscp 26 5000000 10000 20000 conform-action transmit
   exceed-action transmit
 rate-limit output dscp 30 5000000 10000 20000 conform-action transmit
   exceed-action drop
 rate-limit output dscp 34 5000000 10000 20000 conform-action transmit
   exceed-action transmit
 rate-limit output dscp 38 5000000 10000 20000 conform-action transmit
   exceed-action drop
 rate-limit output dscp 0 100000000 50000 50000 conform-action transmit
   exceed-action drop
 no ip mroute-cache
 no keepalive
 crc 16
!
interface GigabitEthernet1/0
 no ip address
 no ip directed-broadcast
 no ip mroute-cache
 shutdown
!
```

```
interface GigabitEthernet2/0
  ip address 192.168.70.2 255.255.255.0
  ip directed-broadcast
  ip router isis
  rate-limit output dscp 8 15000000 10000 20000 conform-action transmit
    exceed-action transmit
  rate-limit output dscp 10 15000000 10000 20000 conform-action transmit
    exceed-action transmit
  rate-limit output dscp 14 5000000 10000 20000 conform-action transmit
    exceed-action drop
  rate-limit output dscp 18 15000000 10000 20000 conform-action transmit
    exceed-action transmit
  rate-limit output dscp 22 5000000 10000 20000 conform-action transmit
    exceed-action drop
  rate-limit output dscp 26 15000000 10000 20000 conform-action transmit
    exceed-action transmit
  rate-limit output dscp 30 5000000 10000 20000 conform-action transmit
    exceed-action drop
  rate-limit output dscp 34 15000000 10000 20000 conform-action transmit
    exceed-action transmit
  rate-limit output dscp 38 5000000 10000 20000 conform-action transmit
    exceed-action drop
  rate-limit output dscp 0 150000000 75000 75000 conform-action transmit
    exceed-action drop
  no ip mroute-cache
  tag-switching ip
!
interface POS3/0
  no ip address
  no ip directed-broadcast
  shutdown
  crc 16
!
interface POS3/1
```

```

no ip address
no ip directed-broadcast
shutdown
crc 16
!
interface POS3/2
no ip address
no ip directed-broadcast
shutdown
crc 16
!
interface POS3/3
no ip address
no ip directed-broadcast
no keepalive
shutdown
crc 16
!
interface ATM5/0
no ip address
no ip directed-broadcast
no ip mroute-cache
no atm ilmi-keepalive

```

-----

```

ASP1-DFT-GSR-B1(config-int)#no atm ilmi-keepalive

```

This command disables Integrated Local Management Interface (ILMI) connectivity procedures for this interface.

```

-----
interface ATM5/0.102 point-to-point
ip address 192.168.215.1 255.255.255.0
no ip directed-broadcast
rate-limit output dscp 8 5000000 10000 20000 conform-action transmit
exceed-action drop

```

```

rate-limit output dscp 10 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 14 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 18 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 22 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 26 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 30 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 34 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 38 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 0 150000000 75000 75000 conform-action transmit
    exceed-action drop
no ip mroute-cache
atm pvc 1 1 1 aal5snap 155000 145000 256 random-detect
tag-switching ip

```

-----

```
ASP1-DFT-GSR-B1(config)#interface ATM5/0.102 point-to-point
```

This command creates a point-to-point subinterface on the ATM port adapter.

```
ASP1-DFT-GSR-B1(config)#atm pvc 1 1 1 aal5snap 155000 145000 256
    random-detect
```

This command creates a permanent virtual circuit (PVC) between ATM switches. This command is comprised of a VPI/VCI pair, a virtual channel (VC), and has an encapsulation method.

-----

```
interface ATM5/1
no ip address
```



```

no ip directed-broadcast
shutdown
no atm ilmi-keepalive
class-int dscp8
map-group MGX-B1
service-policy output test

```

-----

```

ASP1-DFT-GSR-B1(config-int)#class-int dscp8

```

This command allows you to assign a VC class to an ATM main interface or subinterface.

```

ASP1-DFT-GSR-B1(config-int)#map-group MGX-B1

```

This command allows you to associate an ATM map list to an interface or subinterface for either a PVC or switched virtual connection (SVC).

```

ASP1-DFT-GSR-B1(config-int)#service-policy output test

```

This command allows you to use a service policy as a QoS policy within a policy map (this is also referred to as a *hierarchical* service policy).

-----

```

interface ATM5/2
no ip address
no ip directed-broadcast
shutdown
no atm ilmi-keepalive
!
interface ATM5/3
no ip address
no ip directed-broadcast
shutdown
no atm ilmi-keepalive
!
interface Ethernet0
no ip address
no ip directed-broadcast

```

```

no ip route-cache cef
no ip mroute-cache
shutdown
no cdp enable
!
router ospf 99
  redistribute isis level-1-2 subnets
  network 192.168.215.0 0.0.0.255 area 0

```

```

-----
ASP1-DFT-GSR-B1(config)#router ospf 99

```

This command enables Open Shortest Path First (OSPF) and creates a process ID (99).

```

ASP1-DFT-GSR-B1(config-router)#redistribute isis level-1-2 subnets

```

This command redistributes IS-IS level-1 and level-2 traffic into OSPF.

```

ASP1-DFT-GSR-B1(config-router)#network 192.168.215.0 0.0.0.255 area 0

```

This command assigns that network to area 0.

```

-----
router isis
  redistribute ospf 99 metric 1 metric-type internal level-1-2
  net 49.0001.0000.0000.00b2.00
  metric-style transition

```

```

-----
ASP1-DFT-GSR-B1(config-router)#redistribute ospf 99 metric 1 metric-type
  internal level-1-2

```

This command redistributes OSPF into IS-IS.

```

ASP1-DFT-GSR-B1(config-router)#metric-style transition

```

This command allows you to configure a router to be able to generate and accept both old-style and new-style TLVs (TLV stands for type, length, and value).

```

-----
router bgp 70

```

```

no synchronization
network 192.168.60.0
network 192.168.70.0
network 192.168.80.0
redistribute connected
redistribute static
redistribute isis level-2
redistribute ospf 99 metric 1
neighbor 192.168.253.2 remote-as 70
neighbor 192.168.253.2 update-source Loopback1
neighbor 192.168.253.6 remote-as 60
neighbor 192.168.253.6 ebgp-multihop 255
neighbor 192.168.253.6 update-source Loopback1
neighbor 192.168.253.9 remote-as 70
neighbor 192.168.253.9 update-source Loopback1
neighbor 192.168.253.13 remote-as 70
neighbor 192.168.253.13 update-source Loopback1
default-information originate
no auto-summary

```

```
-----
```

```
ASP1-DFT-GSR-B1(config-router)#redistribute isis level-2
```

This command redistributes IS-IS level-2 into BGP.

```
ASP1-DFT-GSR-B1(config-router)#redistribute ospf 99 metric 1
```

This command redistributes OSPF 99 into BGP with a metric of 1.

```
-----
```

```

ip classless
ip route 0.0.0.0 0.0.0.0 POS0/0
ip route 192.168.250.0 255.255.255.0 POS0/0
ip route 192.168.253.6 255.255.255.255 GigabitEthernet1/0
!
!
map-list MGX-B1

```

```

ip 192.168.248.2 atm-vc 1 broadcast
snmp-server engineID local 00000009020000D0FF644820
snmp-server community public RO
snmp-server community private RW

```

```

-----
ASP1-DFT-GSR-B1(config)#map-list MGX-B1

```

This command allows you to define an ATM map statement for either a PVC or SVC.

```

ASP1-DFT-GSR-B1(config)#ip 192.168.248.2 atm-vc 1 broadcast

```

This command creates a logical circuit to ensure that there is reliable communication between two network devices. A virtual channel (VC) is defined by a VPI/VCI pair, and can be either permanent or switched.

```

-----
!
!
line con 0
  exec-timeout 0 0
  transport input none
line aux 0
line vty 0 4
  exec-timeout 39 0
  password 7 08004257061700573305150B242E
  login
!
ntp update-calendar
ntp server 192.168.78.1
ntp server 192.168.216.2
ntp server 192.168.67.1
end

```

```

-----
ASP1-DFT-GSR-B1(config)#ntp update-calendar

```

This command will allow the router to periodically update the calendar from Network Time Protocol (NTP).

```
ASP1-DFT-GSR-B1 (config) #ntp server 192.168.78.1
```

This command enables you to allow the system clock to be synchronized by a time-server that is located on your network.

## Configuration for a Second Cisco Systems Gigabit Switch Router Router That Is Located within the Distribution Layer

The following is the configuration for a second Cisco Systems gigabit switch router (GSR) that is located within the Distribution layer.

```
ASP1-DFT-GSR-C2#show running-configuration
```

```
Building configuration...
```

```
Current configuration:
```

```
!
version 12.0
no service pad
service timestamps debug datetime msec localtime show-timezone
service timestamps log datetime msec localtime show-timezone
service password-encryption
!
hostname ASP1-DFT-GSR-C2
!
boot system slot0:gsr-p-mz_120-9_S.bin
enable secret 5 $1$ShLc$HBF2vRWSEkd/GqQCI2.Ni0
enable password 7 08004257061700573305150B242E
!
clock timezone PST -8
clock summer-time PDT recurring
!
!
!
```

```
!  
!  
!  
!  
ip subnet-zero  
ip domain-name dft.exn.com  
ip name-server 192.168.1.11  
clns routing  
!  
!  
interface Loopback0  
  no ip address  
  no ip directed-broadcast  
  shutdown  
!  
interface Loopback1  
  ip address 192.168.253.2 255.255.255.255  
  ip directed-broadcast  
  ip router isis  
!  
interface POS0/0  
  ip address 192.168.2.2 255.255.255.0  
  no ip directed-broadcast  
  ip router isis  
  rate-limit output dscp 8 5000000 10000 20000 conform-action transmit  
    exceed-action transmit  
  rate-limit output dscp 10 5000000 10000 20000 conform-action transmit  
    exceed-action transmit  
  rate-limit output dscp 14 5000000 10000 20000 conform-action transmit  
    exceed-action drop  
  rate-limit output dscp 18 5000000 10000 20000 conform-action transmit  
    exceed-action transmit  
  rate-limit output dscp 22 5000000 10000 20000 conform-action transmit  
    exceed-action drop
```

```
rate-limit output dscp 26 5000000 10000 20000 conform-action transmit
    exceed-action transmit
rate-limit output dscp 30 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 34 5000000 10000 20000 conform-action transmit
    exceed-action transmit
rate-limit output dscp 38 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 0 100000000 50000 50000 conform-action transmit
    exceed-action drop
no keepalive
tag-switching ip
crc 16
!
interface POS0/1
no ip address
no ip directed-broadcast
rate-limit output dscp 8 5000000 10000 20000 conform-action transmit
    exceed-action transmit
rate-limit output dscp 10 5000000 10000 20000 conform-action transmit
    exceed-action transmit
rate-limit output dscp 14 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 18 5000000 10000 20000 conform-action transmit
    exceed-action transmit
rate-limit output dscp 22 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 26 5000000 10000 20000 conform-action transmit
    exceed-action transmit
rate-limit output dscp 30 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 34 5000000 10000 20000 conform-action transmit
    exceed-action transmit
rate-limit output dscp 38 5000000 10000 20000 conform-action transmit
```

```
    exceed-action drop
rate-limit output dscp 0 10000000 5000 5000 conform-action transmit
    exceed-action drop
rate-limit output dscp 0 10000000 5000 75000 conform-action transmit
    exceed-action drop
shutdown
tag-switching ip
crc 16
!
interface POS0/2
no ip address
no ip directed-broadcast
shutdown
crc 16
!
interface POS0/3
ip address 192.168.50.2 255.255.255.0
no ip directed-broadcast
rate-limit output dscp 8 1500000 10000 20000 conform-action transmit
    exceed-action transmit
rate-limit output dscp 10 1500000 10000 20000 conform-action transmit
    exceed-action transmit
rate-limit output dscp 14 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 18 1500000 10000 20000 conform-action transmit
    exceed-action transmit
rate-limit output dscp 22 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 26 1500000 10000 20000 conform-action transmit
    exceed-action transmit
rate-limit output dscp 30 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 34 1500000 10000 20000 conform-action transmit
    exceed-action transmit
```



```
rate-limit output dscp 38 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 0 50000000 25000 25000 conform-action transmit
    exceed-action drop
no keepalive
tag-switching ip
crc 16
!
interface GigabitEthernet1/0
no ip address
no ip directed-broadcast
shutdown
tag-switching ip
!
interface GigabitEthernet2/0
ip address 192.168.70.1 255.255.255.0
no ip directed-broadcast
ip router isis
rate-limit output dscp 8 15000000 10000 20000 conform-action transmit
    exceed-action transmit
rate-limit output dscp 10 15000000 10000 20000 conform-action transmit
    exceed-action transmit
rate-limit output dscp 14 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 18 15000000 10000 20000 conform-action transmit
    exceed-action transmit
rate-limit output dscp 22 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 26 15000000 10000 20000 conform-action transmit
    exceed-action transmit
rate-limit output dscp 30 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 34 15000000 10000 20000 conform-action transmit
    exceed-action transmit
```

```

rate-limit output dscp 38 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 0 150000000 75000 75000 conform-action transmit
    exceed-action drop
loopback internal
tag-switching ip
tx-cos new

```

```

-----
ASP1-DFT-GSR-C2(config-int)#tx-cos new

```

This command associates a class of service (CoS) queue group name with the transmit queues for this interface.

```

-----
interface ATM5/0
no ip address
no ip directed-broadcast
no ip mroute-cache
no atm ilmi-keepalive
!
interface ATM5/0.102 point-to-point
ip address 192.168.215.1 255.255.255.0
no ip directed-broadcast
rate-limit output dscp 8 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 10 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 14 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 18 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 22 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 26 5000000 10000 20000 conform-action transmit
    exceed-action drop

```

```
rate-limit output dscp 30 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 34 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 38 5000000 10000 20000 conform-action transmit
    exceed-action drop
rate-limit output dscp 0 150000000 75000 75000 conform-action transmit
    exceed-action drop
no ip mroute-cache
atm pvc 1 1 1 aal5snap 155000 145000 256 random-detect
tag-switching ip
!
interface ATM5/1
no ip address
no ip directed-broadcast
shutdown
no atm ilmi-keepalive
!
interface ATM5/2
no ip address
no ip directed-broadcast
shutdown
no atm ilmi-keepalive
!
interface ATM5/3
no ip address
no ip directed-broadcast
shutdown
no atm ilmi-keepalive
!
interface Ethernet0
no ip address
no ip directed-broadcast
no ip route-cache cef
shutdown
```

```
!  
router ospf 99  
  redistribute isis level-1-2 subnets  
  network 192.168.215.0 0.0.0.255 area 0.0.0.0  
!  
router isis  
  redistribute ospf 99 metric 1 metric-type internal level-1-2  
  net 49.0001.0000.0000.00c2.00  
  metric-style transition  
  mpls traffic-eng router-id Loopback1  
-----  
ASP1-DFT-GSR-C2(config-router)#mpls traffic-eng router-id Loopback1
```

This command is used to specify the traffic engineering router identifier for the node to be the address that is associated with the given interface.

```
-----  
router bgp 70  
  no synchronization  
  network 192.168.2.0  
  network 192.168.50.0  
  network 192.168.60.0  
  network 192.168.70.0  
  redistribute connected  
  redistribute static  
  neighbor 192.168.253.3 remote-as 70  
  neighbor 192.168.253.3 ebgp-multihop 5  
  neighbor 192.168.253.3 update-source Loopback1  
  neighbor 192.168.253.6 remote-as 60  
  neighbor 192.168.253.6 ebgp-multihop 255  
  neighbor 192.168.253.6 update-source Loopback1  
  neighbor 192.168.253.9 remote-as 70  
  neighbor 192.168.253.9 ebgp-multihop 5  
  neighbor 192.168.253.9 update-source Loopback1  
  neighbor 192.168.253.13 remote-as 70
```

```

neighbor 192.168.253.13 ebgp-multihop 255
neighbor 192.168.253.13 update-source Loopback1
maximum-paths 2
default-information originate
default-metric 1
no auto-summary

```

```

-----
ASP1-DFT-GSR-C2(config-router)#maximum-paths 2

```

This command is used to improve convergence for routing protocols.

```

ASP1-DFT-GSR-C2(config-router)#default metric 1

```

This command may be used to configure the value for the INTER\_AS metric attribute. The same metric value will then be applied to all BGP updates originating from this router. The default action is to not include an INTER\_AS metric in BGP updates.

```

-----
ip classless
ip route 0.0.0.0 0.0.0.0 192.168.80.2
ip route 192.168.253.6 255.255.255.255 GigabitEthernet1/0
ip route 192.168.253.6 255.255.255.255 POS0/3
!
!
cos-queue-group TEST
precedence 4 random-detect-label 4
random-detect-label 3 2000 3000 10
exponential-weighting-constant 14
queue 3 2000
snmp-server engineID local 00000009020000D0FF642820
snmp-server community public RO
snmp-server community private RW

```

```

-----
ASP1-DFT-GSR-C2(config)#cos-queue-group TEST

```

This command will create a queue group template and enter COS queue group configuration mode.

```
ASP1-DFT-GSR-C2(config-cos-que)#precedence 4 random-detect-label 4
```

This command maps packets that have a particular IP precedence to a random early detection (RED) profile.

```
ASP1-DFT-GSR-C2(config-cos-que)#random-detect-label 3 2000 3000 10
```

This configuration command is used to configure the packet drop characteristics for this group.

```
ASP1-DFT-GSR-C2(config-cos-que)#exponential-weighting-constant 14
```

This command sets the weight that is used to calculate the average queue depth for this group.

```
ASP1-DFT-GSR-C2(config-cos-que)#queue 3 2000
```

This configuration command is used to configure the packet drop characteristics for this group.

```
-----  
line con 0  
  transport input none  
line aux 0  
line vty 0 4  
  password 7 071F20191B1E161713  
  login  
!  
ntp clock-period 17180028  
ntp server 192.168.78.2  
ntp server 192.168.55.1  
ntp server 192.168.216.2  
end
```

## Configuration for a Third Cisco Systems Gigabit Switch Router That Is Located within the Distribution Layer

The following is the configuration for a third Cisco Systems gigabit switch router (GSR) that is located within the Distribution layer.

```
ASP-DFT-GSR-C1#show running-configuration
```

```
Building configuration...
```

```
Current configuration:
```

```
!  
version 12.0  
no service pad  
service timestamps debug datetime msec localtime show-timezone  
service timestamps log datetime msec localtime show-timezone  
service password-encryption  
!  
hostname ASP-DFT-GSR-C1  
!  
boot system slot0:gsr-p-mz_120-9_S.bin  
enable secret 5 $1$WjMw$c7ve2/9hSad2Dh8QpvXcT0  
enable password 7 1209044247  
!  
clock timezone PST -8  
clock summer-time PDT recurring  
!  
!  
!  
!  
class-map match-all TEST  
!  
!  
!  
ip subnet-zero
```

```
ip domain-name dft.exn.com
ip name-server 192.168.1.11
mpls traffic-eng tunnels
!
!
interface Loopback1
 ip address 192.168.253.6 255.255.255.255
 no ip directed-broadcast
!
interface POS0/0
 ip address 192.168.50.1 255.255.255.0
 no ip directed-broadcast
 rate-limit output dscp 10 5000000 10000 20000 conform-action transmit
   exceed-action drop
 rate-limit output dscp 20 5000000 10000 20000 conform-action transmit
   exceed-action drop
 rate-limit output dscp 30 5000000 10000 20000 conform-action transmit
   exceed-action drop
 no ip mroute-cache
 no keepalive
 crc 16
 clock source internal
!
interface POS0/1
 ip address 192.168.254.2 255.255.255.0
 no ip directed-broadcast
 no ip mroute-cache
 no keepalive
 tag-switching ip
 crc 16
!
interface POS0/2
 ip address 192.168.60.1 255.255.255.0
 no ip directed-broadcast
 no ip route-cache cef
```



```
no ip route-cache
no ip mroute-cache
no keepalive
shutdown
crc 16
clock source internal
!
interface POS0/3
no ip address
no ip directed-broadcast
no ip route-cache cef
no ip mroute-cache
shutdown
crc 16
!
interface GigabitEthernet1/0
no ip address
no ip directed-broadcast
shutdown
!
interface POS5/0
no ip address
no ip directed-broadcast
shutdown
crc 16
!
interface POS5/1
no ip address
no ip directed-broadcast
shutdown
crc 16
!
interface POS5/2
no ip address
no ip directed-broadcast
```

```
shutdown
crc 16
!
interface POS5/3
no ip address
no ip directed-broadcast
shutdown
crc 16
!
interface GigabitEthernet6/0
no ip address
no ip directed-broadcast
shutdown
!
interface Ethernet0
no ip address
no ip directed-broadcast
no ip route-cache cef
no ip route-cache
no ip mroute-cache
shutdown
!
router ospf 100
network 192.168.60.0 0.0.0.255 area 0.0.0.0
!
router bgp 60
no synchronization
network 192.168.50.0
network 192.168.60.0
network 192.168.253.6 mask 255.255.255.255
network 192.168.254.0
neighbor 192.168.253.1 remote-as 65535
neighbor 192.168.253.1 ebgp-multihop 255
neighbor 192.168.253.1 update-source Loopback1
neighbor 192.168.253.2 remote-as 70
```

```

neighbor 192.168.253.2 ebgp-multihop 255
neighbor 192.168.253.2 update-source Loopback1
neighbor 192.168.253.3 remote-as 70
neighbor 192.168.253.3 ebgp-multihop 255
neighbor 192.168.253.3 update-source Loopback1
maximum-paths 2
default-metric 1
no auto-summary
!
ip classless
ip route 192.168.253.1 255.255.255.255 POS0/1
ip route 192.168.253.2 255.255.255.255 GigabitEthernet1/0
ip route 192.168.253.3 255.255.255.255 GigabitEthernet6/0
!
!
cos-queue-group test
logging trap emergencies
snmp-server engineID local 00000009020000D0FF646420
snmp-server community public RO
snmp-server community private RW
-----

```

ASP-DFT-GSR-C1(config)#**logging trap emergencies**

This command enables the logging of SNMP emergencies.

```

-----
line con 0
  transport input none
line aux 0
line vty 0 4
  password 7 12090442471C03162E
  login
!
no exception linecard slot 0 sqe-registers
no exception linecard slot 1 sqe-registers

```

```

no exception linecard slot 2 sqe-registers
no exception linecard slot 3 sqe-registers
no exception linecard slot 4 sqe-registers
no exception linecard slot 5 sqe-registers
no exception linecard slot 6 sqe-registers
ntp server 192.168.50.2
ntp server 192.168.60.2
ntp server 192.168.254.1
end

```

```

-----
ASP-DFT-GSR-C1(config)#no exception linecard slot 0 sqe-registers

```

This command disables the storage of crash information for a line card.

## Configuration for a Cisco Systems MGX Router That Is Located within the Access Layer

The following is the configuration for a Cisco Systems MGX router that is located within the Access layer.

```

ASP1-DFT-RPM-B1#show running-configuration

```

```

Building configuration...

```

```

Current configuration:

```

```

!
! Last configuration change at 09:53:45 PST Tue Feb 6 2001
!
version 12.1
no service pad
service timestamps debug uptime
service timestamps log uptime
service password-encryption
!
hostname ASP1-DFT-RPM-B1
!
boot system c:rpm-js-mz.121-2.T.bin

```

```
enable secret 5 $1$ShLc$HBf2vRWSEkd/GqQCI2.Ni0
enable password 7 08004257061700573305150B242E
!
!
class-map OverHead
  match ip dscp 8
class-map HTTP_Cache
  match input-interface Ethernet1/2
class-map Small_PKT_SERIAL
  match ip dscp 26
class-map Large_PKT_SERIAL
  match ip dscp 38
class-map test
  match ip dscp 14
class-map Small_PKT_SERIAL1
  match ip dscp 26
!
!
policy-map HTTP_Cache
policy-map test
  class test
    bandwidth 2000
policy-map TEST
  class test
    bandwidth percent 10
    service-policy test
policy-map switch1
  class OverHead
    bandwidth percent 10
    random-detect
  class Large_PKT_SERIAL
    bandwidth percent 30
    random-detect
  class Small_PKT_SERIAL
    bandwidth percent 30
```

```
random-detect
class HTTP_Cache
  bandwidth percent 10
  random-detect
class class-default
  bandwidth percent 20
  random-detect
!
clock timezone PST -8
clock summer-time PDT recurring
clock calendar-valid
ip subnet-zero
ip tftp source-interface Loopback1
ip domain-name dft.exn.com
ip name-server 192.168.1.11
!
!
ip vrf ip-mpls1
  rd 10.10.254.13:5
  route-target export 10.10.254.13:5
  route-target import 10.10.254.13:5
!
ip vrf lab1-access1
  rd 70:11
  route-target export 70:11
  route-target import 70:11
!
ip vrf lab1-access2
  rd 70:12
  route-target export 70:12
  route-target import 70:12
!
ip vrf lab1-access3
  rd 70:13
  route-target export 70:13
```

```
route-target import 70:13
!
ip vrf lab2-access1
rd 70:21
route-target export 70:21
route-target import 70:21
!
ip vrf lab2-access2
rd 70:22
route-target export 70:22
route-target import 70:22
!
ip vrf lab2-access3
rd 70:23
route-target export 70:23
route-target import 70:23
!
ip vrf lab3-access1
rd 70:31
route-target export 70:31
route-target import 70:31
!
ip vrf lab3-access2
rd 70:32
route-target export 70:32
route-target import 70:32
!
ip vrf lab3-access3
rd 70:33
route-target export 70:33
route-target import 70:33
ip cef
lane client flush
clns routing
```

```
cns event-service server
```

```
-----  
ASP1-DFT-RPM-B1(config)#lane client flush
```

This command enables the flush mechanism of a LAN emulation client (LEC). The *flush* command helps to ensure that cell packets arrive in order.

```
-----  
  
interface Loopback0  
  no ip address  
!  
interface Loopback1  
  ip address 192.168.253.13 255.255.255.255  
!  
interface Ethernet1/1  
  description Cache Engine VPN Network  
  no ip address  
  ip wccp web-cache redirect out  
  no ip mroute-cache  
  shutdown  
!  
interface Ethernet1/2  
  description Cache Engine Legal Network  
  ip address 192.168.200.1 255.255.255.0  
  ip wccp web-cache redirect out  
  no ip mroute-cache  
  shutdown  
  tag-switching ip  
!  
interface Ethernet1/3  
  no ip address  
  no ip mroute-cache  
  shutdown  
!
```



```

interface Ethernet1/4
  no ip address
  no ip mroute-cache
  shutdown
!
interface Switch1
  no ip address
  no ip mroute-cache
  no atm ilmi-keepalive
!
interface Switch1.101 point-to-point
  description Lab1 64k Frame 32K Cir to 1.1.0.16
  ip wccp web-cache redirect out
  pvc lab1_access1 0/11
    vbr-nrt 64 32 256

```

```
-----
ASP1-DFT-RPM-B1(config-int)#pvc lab1_access1 0/11
```

A virtual connection is permanently established to lab\_access1. The PVC saves bandwidth that is associated with circuit establishment and tear down where virtual connections exist all the time.

```
ASP1-DFT-RPM-B1(config-int)#vbr-nrt 64 32 256
```

This command enables nonreal-time variable bit rate (VBR-nrt) that uses sustained cell rate (SCR), peak cell rate (PCR), and maximum burst size (MBS).

- SCR defines the sustained rate at which you can expect to transmit data traffic.
- PCR defines the maximum rate at which you expect to transmit data traffic.
- MBS defines the duration (in kbps) at which the router sends at the peak cell rate.

```
-----
tag-switching ip
```

```
!  
interface Switch1.102 point-to-point  
  description Lab1 128k Frame 64K Cir to 1.2.0.16  
  ip wccp web-cache redirect out  
  no ip mroute-cache  
  pvc lab1_access2 0/12  
    vbr-nrt 128 64 512  
  !  
  tag-switching ip  
!  
interface Switch1.103 point-to-point  
  description Lab1 256k Frame 128K Cir to 1.3.0.16  
  ip wccp web-cache redirect out  
  no ip mroute-cache  
  pvc lab1_access3 0/13  
    vbr-nrt 256 128 768  
  !  
  tag-switching ip  
!  
interface Switch1.104 point-to-point  
  description Lab1 512k Frame 256K Cir to 1.4.0.16  
  ip address 192.168.248.129 255.255.255.128  
  ip wccp web-cache redirect out  
  no ip mroute-cache  
  pvc lab_access4 0/14  
    service-policy output switch1  
    vbr-nrt 512 256 1024  
  !  
  tag-switching ip  
!  
interface Switch1.105 point-to-point  
  description LAB T1 Frame 768k Cir to 1.5.0.16  
  ip vrf forwarding ip-mpls1  
  ip address 10.10.254.13 255.255.255.252
```

```

ip wccp web-cache redirect out
pvc lab1_access5 0/15
    service-policy output switch1
    vbr-nrt 1536 768 1536
!
tag-switching ip
!
interface Switch1.107 point-to-point
ip address 192.168.244.2 255.255.255.0
ip wccp web-cache redirect out
pvc GSR-B2_5_0_105 0/19
    vbr-nrt 155000 155000 60000
!
tag-switching ip
!
interface Switch1.108 point-to-point
ip address 192.168.215.2 255.255.255.0
ip wccp web-cache redirect out
pvc GSR-B2_5_0_106 0/100
    vbr-nrt 150000 150000 60000
!
tag-switching ip
!
interface Switch1.201 point-to-point
description Lab2 64k Frame 32K Cir to 2.1.0.16
ip wccp web-cache redirect out
pvc lab3_access1 0/21
    vbr-nrt 64 32 256
!
tag-switching ip
!
interface Switch1.202 point-to-point
description Lab2 128k Frame 64K Cir to 2.2.0.16
ip address 192.168.249.129 255.255.255.128
ip wccp web-cache redirect out

```

```
no ip mroute-cache
pvc lab2_access2 0/22
    vbr-nrt 128 64 512
!
tag-switching ip
!
interface Switch1.203 point-to-point
description Lab2 256k Frame 128K Cir to 2.3.0.16
ip vrf forwarding lab2-access3
ip address 10.30.254.13 255.255.255.252
ip wccp web-cache redirect out
no ip mroute-cache
pvc lab2_access3 0/23
    vbr-nrt 256 128 768
!
tag-switching ip
!
interface Switch1.204 point-to-point
description Lab2 512k Frame 256K Cir to 2.4.0.16
ip vrf forwarding lab1-access1
ip wccp web-cache redirect out
no ip mroute-cache
pvc lab2_access4 0/24
    vbr-nrt 512 256 1024
!
tag-switching ip
!
interface Switch1.205 point-to-point
description LAB2 T1 Frame 768k Cir to 2.5.0.16
ip wccp web-cache redirect out
pvc lab2_access5 0/25
    vbr-nrt 1536 768 1536
!
tag-switching ip
!
```

```
interface Switch1.301 point-to-point
description Lab3 64k Frame 32K Cir to 2.1.0.16
ip wccp web-cache redirect out
tag-switching ip
!
interface Switch1.302 point-to-point
description Lab3 128k Frame 64K Cir to 2.2.0.16
ip address 192.168.228.129 255.255.255.128
ip wccp web-cache redirect out
no ip mroute-cache
pvc lab3_access2 0/32
    vbr-nrt 128 64 512
!
tag-switching ip
!
interface Switch1.303 point-to-point
description Lab3 256k Frame 128K Cir to 2.3.0.16
ip vrf forwarding lab3-access3
ip address 10.30.254.13 255.255.255.252
ip wccp web-cache redirect out
no ip mroute-cache
pvc lab3_access3 0/33
    vbr-nrt 256 128 768
!
tag-switching ip
!
interface Switch1.304 point-to-point
description Lab3 512k Frame 256K Cir to 2.4.0.16
ip vrf forwarding lab3-access1
ip address 10.10.254.13 255.255.255.252
ip wccp web-cache redirect out
no ip mroute-cache
pvc lab3_access4 0/34
    vbr-nrt 512 256 1024
!
```

```
tag-switching ip
!
interface Switch1.305 point-to-point
description LAB3 T1 Frame 768k Cir to 2.5.0.16
ip wccp web-cache redirect out
pvc lab3_access5 0/35
vbr-nrt 1536 768 1536
!
tag-switching ip
!
router ospf 99
redistribute static subnets
network 192.168.200.0 0.0.0.255 area 0.0.0.0
network 192.168.215.0 0.0.0.255 area 0.0.0.0
network 192.168.248.0 0.0.0.255 area 0.0.0.0
network 192.168.249.0 0.0.0.255 area 0.0.0.0
network 192.168.253.13 0.0.0.0 area 0.0.0.0
!
router rip
version 2
network 10.0.0.0
default-information originate
no auto-summary
!
address-family ipv4 vrf lab1-access3
version 2
network 10.0.0.0
default-information originate
no auto-summary
exit-address-family
!
address-family ipv4 vrf lab1-access2
version 2
network 10.0.0.0
default-information originate
```

```
no auto-summary
exit-address-family
!
address-family ipv4 vrf lab1-access1
version 2
network 10.0.0.0
default-information originate
no auto-summary
exit-address-family
!
address-family ipv4 vrf ip-mpls1
version 2
network 10.0.0.0
default-information originate
no auto-summary
exit-address-family
!
router bgp 70
no synchronization
no bgp default ipv4-unicast
network 10.10.254.12 mask 255.255.255.252
network 192.168.200.0
network 192.168.253.13 mask 255.255.255.255
redistribute ospf 99 metric 1
redistribute rip
neighbor 192.168.253.2 remote-as 70
neighbor 192.168.253.2 update-source Loopback1
neighbor 192.168.253.2 activate
neighbor 192.168.253.3 remote-as 70
neighbor 192.168.253.3 update-source Loopback1
neighbor 192.168.253.3 activate
neighbor 192.168.253.9 remote-as 70
neighbor 192.168.253.9 update-source Loopback1
neighbor 192.168.253.9 activate
default-information originate
```

```
!  
address-family ipv4 vrf lab3-access3  
no auto-summary  
no synchronization  
exit-address-family  
!  
address-family ipv4 vrf lab3-access2  
no auto-summary  
no synchronization  
exit-address-family  
!  
address-family ipv4 vrf lab3-access1  
no auto-summary  
no synchronization  
exit-address-family  
!  
address-family ipv4 vrf lab2-access3  
no auto-summary  
no synchronization  
exit-address-family  
!  
address-family ipv4 vrf lab2-access2  
no auto-summary  
no synchronization  
exit-address-family  
!  
address-family ipv4 vrf lab2-access1  
no auto-summary  
no synchronization  
exit-address-family  
!  
address-family ipv4 vrf lab1-access3  
no auto-summary  
no synchronization  
exit-address-family
```



```
!  
address-family ipv4 vrf lab1-access2  
no auto-summary  
no synchronization  
exit-address-family  
!  
address-family ipv4 vrf lab1-access1  
redistribute connected  
redistribute static  
neighbor 192.168.253.9 remote-as 70  
neighbor 192.168.253.9 update-source Loopback1  
neighbor 192.168.253.9 activate  
neighbor 192.168.253.9 send-community both  
no auto-summary  
no synchronization  
network 10.10.11.0 mask 255.255.255.0  
network 10.10.254.12 mask 255.255.255.252  
exit-address-family  
!  
address-family ipv4 vrf ip-mpls1  
redistribute connected  
redistribute rip  
no auto-summary  
no synchronization  
exit-address-family  
!  
address-family vpv4  
neighbor 192.168.253.9 activate  
neighbor 192.168.253.9 send-community both  
default-information originate  
exit-address-family  
!  
ip classless  
ip route 0.0.0.0 0.0.0.0 192.168.215.1  
ip route 192.168.248.0 255.255.255.128 Switch1.104
```

```

ip route 192.168.249.0 255.255.255.128 Switch1.105
ip route 192.168.253.3 255.255.255.255 Switch1.108
ip route vrf lab1-access1 10.10.11.0 255.255.255.0 Switch1.103
no ip http server
!
snmp-server engineID local 000000090200005054AD9480
snmp-server community public RO
snmp-server community private RW
snmp-server packetsize 2048
snmp-server host 192.168.1.11 public
-----

```

```

ASP1-DFT-RPM-B1(config)#snmp-server packetsize 2048

```

This command is used to create the maximum Simple Network Management Protocol (SNMP) packet size that is permitted when the SNMP server is receiving a request or generating a reply.

```

ASP1-DFT-RPM-B1(config)#snmp-server host 192.168.1.11 public

```

This command is used to specify the recipient of an SNMP notification.

```

-----
line con 0
  transport input none
line aux 0
line vty 0 4
  password 7 051B075A745B411B1D
  no login
!
ntp master
ntp update-calendar
ntp server 10.10.254.14
ntp server 192.168.216.1
ntp server 192.168.249.130
rpmrsctrn PAR 100 100 0 255 0 3840 4070
addcon auto_synch off
addcon vcc switch 1.101 11 rname MGX-B1 rslot 1 1 0 16 master local

```

```

addcon vcc switch 1.102 12 rname MGX-B1 rslot 1 2 0 16 master local
addcon vcc switch 1.103 13 rname MGX-B1 rslot 1 3 0 16 master local
addcon vcc switch 1.104 14 rname MGX-B1 rslot 1 4 0 16 master local
addcon vcc switch 1.105 15 rname MGX-B1 rslot 1 5 0 16 master local
addcon vcc switch 1.201 21 rname MGX-B1 rslot 2 1 0 16 master local
addcon vcc switch 1.202 22 rname MGX-B1 rslot 2 2 0 16 master local
addcon vcc switch 1.203 23 rname MGX-B1 rslot 2 3 0 16 master local
addcon vcc switch 1.204 24 rname MGX-B1 rslot 2 4 0 16 master local
addcon vcc switch 1.205 25 rname MGX-B1 rslot 2 5 0 16 master local
addcon vcc switch 1.108 100 rname MGX-B1 rslot 0 5 1 1
end

```

```

-----
ASP1-DFT-RPM-B1 (config) #ntp master

```

This command is used to configure the IOS software as a Network Time Protocol (NTP) master clock. This allows peers to synchronize themselves when an external NTP source is not available.

```

ASP1-DFT-RPM-B1 (config) #rpmrscprtn PAR 100 100 0 255 0 3840 4070

```

This command is used to set up resource partitioning. It uses the following switches in this configuration: Ingress Percent Bandwidth, Egress Percent Bandwidth, Minimum VPI Value, Maximum VPI Value, Minimum VCI Value, Maximum VCI Value, Number of Logical Connections (LCNs).

```

ASP1-DFT-RPM-B1 (config) #addcon auto synch off

```

This command disables automatic synchronization between the connections.

```

ASP1-DFT-RPM-B1 (config) #addcon vcc switch 1.101 11 rname MGX-b1 rslot 1
1 0 16 master local

```

This command is used to add a connection to the PVC, using VCC. This instance used the following switches: Add a connection {VCC}, Switch Virtual Interface, Chassis slot number, Switch interface number, local VCI value, remote node name {name}, Remote slot number, Remote interface, Remote VPI, Remote VCI, Remote VPI, Remote VCI, Master end of the ATM connection, and Local option.

## Summary

As you can see, these configurations are based mostly on the Distribution and Access layers of the figures presented in this Appendix. The reason that I have only included these is that your core will vary greatly depending on what you decide to implement. If you decide to do voice and video, or everything over IP (XoIP), then you will want more robust, leading-edge equipment. If you are looking to provision bandwidth for your application and customers, then you will probably use the gear that we listed.

Remember, there is no such thing as a perfect, permanent infrastructure. There will always need to be support and upgrades for your network. One of the major concerns that I had while writing this appendix is that most of the equipment that you see here will be obsolete within the next two years. That is the nature of the Information Technology and Internet beast.

I hope that this appendix has given you a basic understanding of the complexity that is involved, even in creating a “test” network. All I can offer is to say, “Don’t become overwhelmed.” Things are constantly changing, vendors always want to add more functionality, and users will always look for ease of installation and support.

EngineX Networks Inc. would like to say good luck in all of your current and future endeavors.



## ASP Configuration Handbook Fast Track

**This Appendix will provide you with a quick, yet comprehensive, review of the most important concepts covered in this book.**

## ❖ Chapter 1: An Introduction to ASPs for ISPs

### Why This Book Is for You

- ☑ According to the International Data Corporation (IDC), worldwide spending for outsourcing services should reach approximately \$142 billion by the year 2002.
- ☑ The ASP market began capturing the interest and commitment from a large number of venture capitalists and the telecommunications industry in the late 1990s.
- ☑ The ASP concept is the advent of a new computing era, with small to medium-sized companies searching for IT alternatives, and a gradual acceptance among larger enterprises.
- ☑ The IT infrastructure has evolved from a self-contained environment to a distributed computing model and now toward a net-centric infrastructure that links multiple areas of operation.

### Definitions of Common ASP Terms

- ☑ An Internet service provider (ISP) is an organization that provides access to the Internet. ISPs can provide service via modem, or dedicated or on-demand access.
- ☑ The ASP Industry Consortium, an alliance of companies formed to promote and educate the IT industry, offers the following definition: *“An ASP manages and delivers application capabilities to multiple entities from a data center across a wide area network.”*
- ☑ The definition of a *pure* ASP is an ASP that joins with a particular ISV, and performs the initial application implementation and integration.
- ☑ Information technology (IT) outsourcing is the transfer of an organization’s internal IT infrastructure, staff, processes, or applications to an external resource provider.
- ☑ Business process outsourcing (BPO) and information utilities providers are primarily concerned with economic and efficient outsourcing for the highly sophisticated but repetitive business processes.

## Chapter 1 Continued

- ☑ Platform IT outsourcing offers an array of data center services, such as facilities management, onsite and offsite support services, data storage and security, and disaster recovery.

## The Elements That Make an ASP Viable

- ☑ The initial purchase price of system software such as a Unix platform or a Microsoft Windows platform and their licensing are considered part of the initial system software purchase.
- ☑ Initial application software acquisition is any application that assists in the productivity of the organization.
- ☑ Hardware upgrade costs are associated with obligatory improvements to hardware when your company will need to support expanded applications databases, and a more robust operating system.
- ☑ Operating system software may need to be upgraded to support newer, vigorous applications.
- ☑ Applications are constantly being improved due to customer and client demands.

## Possible Business Models and Offerings

- ☑ ASPs host services work on an extensive array of hardware, so at any given time that hardware will have a substantial amount of its processing power idle. The ASP will find that this ability to provision and partition that extra horsepower can be the basis for a very valuable and profitable differentiation service offering.
- ☑ The ability to offer different types of service to different types of clients is an incredibly valuable way for ASPs and ISPs to provide granular and real-world service degrees of difference.

## Types of ASP Firms

- ☑ There are several types of ASP-enabled firms. These organizations can be separated into professional consulting, project-based service providers, outsourcing providers, staff augmentation providers, education and training providers, and value-added resellers.



## Chapter 1 Continued

- ☑ Professional consulting firms focus on corporate-level business and strategic engagements. This can be broken down into three subcategories: IT consulting, Strategic management consulting, and Business process consulting.
- ☑ Clients that select project-based service providers for projects are opting for well-defined tangible deliverables and scopes. Contract designs range from a billable-hours approach to fixed-price engagements for components and entire projects. These companies focus on industry expertise, either in specific technologies or industry applications.
- ☑ Outsourcing providers are organizations that provide process automation services, facilities management, and operations for clients who require an assortment of technical answers.
- ☑ Staff augmentation organizations specialize in providing IT professionals, on a temporary or long-term contract basis, to clients who need specific skill sets and support for internal systems and development projects.
- ☑ Education and training companies provide training and help desk consulting to firms that have implemented custom-designed or packaged software products.
- ☑ Value-added reseller (VAR) organizations are solution-oriented vendors who can provide integration for hardware and software systems.

## ISO-OSI Seven Layer Model

- ☑ The OSI reference model is a conceptual model composed of seven layers, each specifying particular network functions. The OSI model divides these communications involved with the moving of information between networked computers into seven smaller, more manageable layers.
- ☑ The *upper layers* of the OSI model handle application issues and are generally implemented in software.
- ☑ The *lower layers* of the OSI model are also known as the Data Transport layer.
- ☑ These *pseudo layers* are not actual OSI model layers, but they will directly influence the way in which you will implement your equipment and policies.

## Chapter 1 Continued

### Choosing the Best Platform for Your ASP

- ☑ ASPs take advantage of existing Internet connectivity to offer corporations the opportunity to outsource not only peripheral applications but also mission-critical applications.
- ☑ Traditional ISPs are experiencing an explosion of data traffic on their networks. The Internet and its dramatic growth have fueled the need to satisfy this increasing data demand by forcing the migration towards multiservice network platforms.
- ☑ Reliability is one of the most important considerations to make when choosing a server platform, but it is likely that ASPs should be most concerned about the operating system. This often boils down to a choice between some form of Unix, Linux, or Microsoft Windows platform.

### Business Drivers for the Conversion to ASP

- ☑ The new value proposition that is offered by various services using Internet and intranet technologies is the ability of the ASP to free its customers from having to develop, maintain, and provide services for themselves.
- ☑ There is also the added benefit of using an ASP's application management expertise. Over the lifetime of an application, such as Enterprise Resource Planning (ERP), an ASP can estimate that software licensing, hardware and basic infrastructure costs will account for less than one-fifth of the total cost of ownership (TCO) over a five-year period.
- ☑ By using outsourced resources, companies can become more efficient with their internal business processes, and that can make the difference between success and failure in this intensely competitive market.

### Performance Issues

- ☑ As ASPs become a more viable alternative for corporate IT, the demands that are placed on service providers to deliver high-level 24x7 service will continue to escalate. This places the burden of reliability and scalability on the ASPs' underlying system platforms.

## Chapter 1 Continued

- ☑ *Five nines* means that in a year's time, a system will be “down” or offline for no longer than five minutes.
- ☑ Clustering is the combination of multiple servers that will allow for failover and data reclamation from storage in case of a catastrophic occurrence.

## Problems That Could Arise from a Conversion

- ☑ ISPs that are converting to ASPs face an assortment of hurdles in trying to break into their chosen markets. Perhaps the greatest obstacle is the acquisition, training, and retention of intellectual property, all of which will allow an ASP to offer stellar implementation, service, and support.

## Major Issues in the Implementation of an ASP Model

- ☑ The contractual assurances that an ASP must make to its clients is usually some form of negotiated contract that specifies acceptable levels of service, availability, security, and performance collectively called a service level agreement (SLA).
- ☑ The software applications must conform to a company's business guidelines by being able to discriminate between customers, partners, and suppliers and provide the best business value, and return on a company's investments (ROI) in time and resources.

## What Is Needed to Sell Your Services

- ☑ An ASP must draw together resources that traditionally have operated independently of one another.
- ☑ To successfully deploy a dynamic and interactive application, you will need to integrate several components, while providing access to other network resources.

## ❖ Chapter 2: The Business Case

### ISP Market Conditions

- ☑ Internet access reached 50-percent market penetration in less than eight years of existence. The growth rate in the United States is projected to be anywhere from 40 to 110 percent for at least the next few years.
- ☑ According to *Boardwatch Magazine*, there are currently more than 7700 ISPs (early 2001) that are doing business in the United States alone.
- ☑ The reality of the DSL market is that providers must rely on the Incumbent Local Exchange Carrier (ILEC) for the all-important connection to the customer. That forces ISPs into the position of commodity resellers in direct competition with their suppliers.
- ☑ While broadband connections seem to be following the same economic pattern as their slower counterparts, their significance should not be overlooked. Increasing broadband access speeds will be the foundation for the value-added services that will allow ISPs to differentiate their offerings.

### Service Provider Business Requirements

- ☑ In order to break out of the current cycle, many service providers and ISPs in particular will have to address these factors: commoditized offering, significant pricing pressure, high customer churn, drastically reduced valuations, restricted access to capital.
- ☑ The current demands of the financial community once again include traditional terms such as *differentiation*, *barriers to entry*, and *profitability*. The easy money is gone.

### The Evolving ISP

- ☑ The evolving ISP must overcome the issues that are facing its core business, the demands of its customers, and the demands of the investor community.
- ☑ Among the first required steps to migrate to value-added offerings is to develop a highly reliable service model.

## Chapter 2 Continued

- ☑ Current implementations of hosted applications and Web sites are accessed across existing connections, sometimes with significant delay, but saturated links and latency will not be tolerated in the future.

## The Service Provider of the Future

- ☑ ISPs must ask themselves what type of services they will need to have available in two, three, and five years in order to remain competitive and profitable.
- ☑ Businesses and consumers will not purchase services from a provider that cannot include all required data, voice, and entertainment offerings. Over the next few years, providers who have not embraced new offerings and developed methods for continuously developing new offerings will not be facing commoditization, but extinction.

## The Case for Application Service Provider Conversion

- ☑ The ASP offering is a revolutionary response to the inefficiencies in our current distributed computing environment.
- ☑ Application hosting presents enormous potential for ISPs. It addresses many of the market realities that are currently plaguing the segment. Application hosting provides the opportunity to differentiate Internet connections and create additional high-margin revenue streams.
- ☑ International Data Corporation (IDC) placed worldwide ASP spending at \$300 million for 1999 and estimated spending of \$7.8 billion by 2003 based on 92-percent compound annual growth. Many other companies have projected much higher figures.

## Critical Success Factors

- ☑ *Application infrastructure provider (AIP)* is a term used to describe a provider that offers ASPs wholesale network and data center services.
- ☑ Leveraging channel partners with complementary offerings can be very effective, but these channels must be managed differently from direct sales methods.

## Chapter 2 Continued

- ☑ Current Analysis published the results of their survey of ASP customers that ranked the major decision criteria they used to choose an ASP provider. Major factors included support, expertise, price, and reputation.

## ❖ Chapter 3: Server Level Considerations

### Implementation, Where to Begin

- ☑ At the heart of an ISP/ASP are the server base and the application software packages. If they do not function efficiently, the ASP will not run effectively.
- ☑ Today, there are only two basic types of microprocessors available for computers: Complex Instruction Set Computers (CISC), and Reduced Instruction Set Computers (RISC).
- ☑ SMP is an architecture that provides better performance by using multiple processors in the same server.
- ☑ Fibre Channel has been introduced as a replacement for the SCSI architecture. Fibre Channel provides a method for transmitting data between computers at a rate of 100 Mbps, and scales up to 1 Gigabit per second (Gbps).
- ☑ Link aggregation allows a single server to use two or more installed network interface cards (NICs) to aggregate bandwidth across several links.

### Software Solutions for Your ASP

- ☑ *System software* describes software packages that provide the basis for all other applications that are run on a computer.
- ☑ Unix is not a proprietary operating system, and the source code has been available to the public since its inception. Currently, the leading Unix environment is Solaris from Sun Microsystems.
- ☑ Windows 2000 Advanced Server offers all of the features available in the standard version, but includes more reliability and scalability, as well as additional features for applications that require a higher level of scalability.
- ☑ Novell offers a powerful network operating system called NetWare. This operating system was originally designed for use in small to enterprise businesses and networks, and typically used a protocol stack called Internet Packet eXchange (IPX).

## Chapter 3 Continued

### Application Software Types

- ☑ *Applications* is the term used to describe a group of programs or code designed to perform a specific function directly for users or other application packages.
- ☑ *Internet Information Server (IIS)* is a scalable Web server offering from Microsoft Corporation that runs under the Windows family of operating systems.
- ☑ Apache HTTP Server is an open-source software package that is organized by the Apache Software Foundation.
- ☑ A database can be defined as a collection of data that is organized for management and access.
- ☑ *Middleware* can be considered the “glue” that holds applications together. It is a general term for any computer application whose purpose is to combine or mediate between two applications in order to allow them to share data between them.

### Network Service Considerations

- ☑ *Network storage* defines the ability to store information on a remote system connected over a network.
- ☑ NFS was first released in 1984 by Sun Microsystems Corporation.
- ☑ Today, many systems use NFS to connect servers to centralized storage. Since NFS was designed on the Unix platform, it has remained a Unix tool, for the most part. It is possible to find NFS servers and clients that run under other operating systems, such as Windows, but they are not very desirable since they are not native to the particular operating system.

### Data Backups and How They Can Affect You

- ☑ Although hardware platforms have become more reliable over the years, the fact still remains that your data is stored on what is essentially a mechanical device; a disk that rotates at very high speeds with another bit of metal called a head that floats left and right across the surface of the disk many times a second.

## Chapter 3 Continued

- ☑ You will most likely use a third-party backup program as opposed to the generic ones that sometimes come with your operating system, or storage devices. Some of the products that you will run across such as ARCserve, Veritas Backup Exec, UltraBac, or NovaStor, will allow advanced scheduling with various levels of flexibility.
- ☑ One of the defining factors between backup systems is how tapes are rotated and what files get backed up to which tape. Each rotation method has different advantages that can be applied to systems and provide for different results.

## Virus Scanning Suggestions

- ☑ A virus can halt your servers, and can even remove data from your hard disks. What's worse is that it can spread to incorporate the computers throughout your entire network and into your client's networks, infecting every server along the way and leaving mass data destruction in its wake.
- ☑ When using an Internet Gateway product, make sure that you have a system that will allow you to queue incoming e-mail messages. If mail is received faster than it can be processed by an Internet gateway, it could start dropping or bouncing messages unless you have software that allows incoming messages to be queued.

## Thin Client Solutions

- ☑ One of the primary focuses for an ASP is to ensure the delivery of its products or services to each client's desktop.
- ☑ Independent Computing Architecture (ICA) allows the delivery of an application from a centralized server to any end-user desktop, regardless of the operating system or platform.

## Maintenance and Support Issues

- ☑ Eventually, every piece of hardware and software operated by your company will need an upgrade of some sort.



## Chapter 3 Continued

- ☑ When you consider that you might be performing hardware upgrades as well as software upgrades, and that one upgrade might cause another, it just does not make sense to even attempt to upgrade the servers all at once.
- ☑ Whenever performing an upgrade, always incorporate a back-out plan. In some cases, it may even be necessary to provide several back-out plans at every stage of a complicated upgrade.
- ☑ In order to catch problems before they arise, you will need to perform some type of system monitoring.

## ❖ Chapter 4: Performance Enhancement Technologies

### Web Caching and How It Works

- ☑ The intent of caching is to move Web content as close to the end users or the edge of the network as possible for quick access to improve the customers' satisfaction levels, and gives your ASP the competitive advantage.
- ☑ Hardware devices will cache frequently used data and instructions in order to speed tasks.
- ☑ Caching as much Web content as possible within the boundaries of an ISP while using modest amounts of upstream bandwidth is a way to grant clients what they require without creating a “black hole” for bandwidth investment on the part of the service provider.

### Deployment Models for Caching

- ☑ In the forward proxy cache configuration, a client's requests go through the cache on the way to the destination Web server.
- ☑ A transparent cache resides in the flow of the network and is invisible to a client's browser. Clients realize the benefits of caching without reconfiguring the browsers.
- ☑ Reverse cache servers can be deployed throughout the network to create a distributed site of hosted content; this model is commonly referred to as *site replication*.

## Chapter 4 Continued

- ☑ A cache appliance (this can also be called a *thin server*) can be defined as a device that offers a limited number of dedicated functions, and is able to deliver those functions more effectively than a multipurpose device can.

## Load Balancing in Your Infrastructure

- ☑ Load balancing, also called Layer 4–7 switching, occurs when cluster of Web servers are created to handle massive amounts of requests.
- ☑ Localized load balancing occurs when the load balancer determines which server should receive new requests.
- ☑ Distributed load balancing sends packets across dispersed networks, which can be located in geographically separate areas from the local server.

## Load Balancing Solutions from F5

- ☑ As more servers are added to the DNS round-robin rotation, traffic will be unevenly distributed. The older servers will tend to receive more traffic than newer servers, as the IP addresses of older servers are usually cached by more users than the addresses of newer servers are.
- ☑ When you implement a network device that is capable of high availability, you want it to guarantee that it can deliver IP-based services, which are always available. To do this, you must remember that it is imperative that both “quality of service” based high availability and load balancing are addressed so that your client has a good usability experience.

## Cisco Systems’ LocalDirector

- ☑ There are generally two approaches for scaling a server farm-based system. The first approach is to continuously upgrade the size and processing power of individual servers in the farm. The second approach is to add more servers as you require more capacity.
- ☑ Load-balancing technology does not normally consider variables such as bandwidth, server performance, and job size for optimizing the traffic loads among your server farms. Load balancing can allow you to incrementally scale the capacity of servers in your server farms in a more efficient manner.

## Chapter 4 Continued

- ☑ LocalDirector is considered a transparent device, as it is able to work with any TCP-based service or application. There is no special software required on the server, as these are external devices.
- ☑ The LocalDirector is considered a stateful device, as it is able to monitor and can track all TCP connections that are occurring between clients and servers.

## Foundry Networks' ServerIron

- ☑ Foundry's ServerIron Web switches provide high-performance content and application-aware traffic and server load balancing. ServerIron has the functionality of a traditional Layer 2 and Layer 3 switch built in, and is able to examine the content at Layer 4 and above through the packet header.
- ☑ ServerIron load-balancing characteristic is based on Layer 4 traffic such as HTTP, FTP, SSL, and email. This creates the ability to transparently distribute data traffic among multiple servers.

## Content Delivery Networks

- ☑ The networking industry's focus from Layer 3 connectivity issues is shifting to the creation of intelligent, Layer 4–7 networks that can support the rigorous response-time requirements of these new types of content. The emphasis is now turning to content delivery networks (CDN).
- ☑ CDNs are able to provide QoS to the Internet's IP-based backbone, which helps to eliminate or minimize delay.
- ☑ Content provider organizations build content for the Web, and are faced with delivering content that has dynamic characteristics to customers who require high levels of service.

## CDN Solutions from Various Vendors

- ☑ Content Distributor uses the agent/manager design and a proprietary communications protocol that can replicate content updates to a community of servers over any TCP/IP-based network.

## Chapter 4 Continued

- ☑ Cisco Systems' Content Delivery Network (CDN) system was developed to help service providers to deploy content delivery services so that they could realize new profit opportunities.

## ❖ Chapter 5: Storage Solutions

### Upfront Concerns and Selection Criteria

- ☑ Currently, there are many differing manufacturers of storage-based equipment, and several methods of delivering storage solutions to your servers and clients.
- ☑ With mass-storage products, some of the major manufacturers may only offer proprietary equipment, while others may standardize their equipment, using a technology such as fiber channel to ensure that their product will work with a similar offering from another manufacturer.
- ☑ Security should always be a concern, but it is especially important given the high visibility of ISPs and ASPs.
- ☑ Outboard security is any type of security feature that is located on the host. It might be an external authentication scheme that is provided by a firewall.
- ☑ You may already own storage devices that use interfaces other than fiber channel, such as small system computer interface (SCSI) or enhanced integrated drive electronics (EIDE) for host connections. It can sometimes prove difficult to port older hardware to some newer storage solutions.

### Directly Attached Storage in Your Infrastructure

- ☑ Server-to-storage access, or directly attached storage, has been in use in much of the history of computing, and still exists in over 90 percent of implementations today.
- ☑ In directly attached implementations, storage devices are directly connected to a server using either interfaces and/or bus architecture such as EIDE or SCSI.

## Chapter 5 Continued

### Network Attached Storage Solutions

- ☑ A NAS is a device that provides server-to-server storage. A NAS is basically a massive array of disk storage connected to a server that has been attached to a local area network (LAN).
- ☑ QoS has the ability to delegate priority to the packets traversing your network, forcing data with a lower priority to be queued in times of heavy use, and allowing for data with a higher priority to still be transmitted.
- ☑ When designing NAS in your network, probably the most effective solution for latency and saturation issues is the location of your NAS servers in relation to the hosts and systems that access their data.

### Storage Area Networks

- ☑ A storage area network (SAN) is a networked storage infrastructure that interconnects storage devices with associated servers. It is currently the most cutting-edge storage technology available, and provides direct and indirect connections to multiple servers and multiple storage devices simultaneously.
- ☑ A SAN can be thought of as a simple network that builds off the familiar LAN design.
- ☑ Distributed computing, client/server applications, and open systems give today's enterprises the power to fully integrate hardware and software from different vendors to create systems tailored to their specific needs.
- ☑ SANs remove data traffic—backup processes, for example—from the production network, giving IT managers a strategic way to improve system performance and application availability.
- ☑ Multihost arrays are the most simplistic and most common form of SAN virtualization implementation.

### Scalability and How It Affects Your Business

- ☑ A SAN is designed to span great distances, which allow it even more flexibility, since there is not a requirement for the SAN devices to be in close proximity to the hosts that access them.

## Chapter 5 Continued

- ☑ Wire speed plays an important role in delivering data to host devices. Whether your environment consists of directly attached storage, NAS, SAN, or a combination there of, you will still have bandwidth concerns that will limit the amount of actual data that can be sent across the wire at any given moment.

## Fault Tolerance Features and Issues

- ☑ One of the largest advantages a SAN has to offer is the true ability to share resources between other server and host systems.
- ☑ Remote mirroring is an excellent form of disaster recovery offered by SAN technology. Today, it allows for a complete copy of your data to be contained at a remote location that might be located up to 40 kilometers away.
- ☑ Redundant Array of Inexpensive Disks (RAID) provides methodology for storing the same data in different places on multiple hard disks.

## SAN Solutions Offered by Various Vendors

- ☑ IBM's SAN strategy involves the migration to a SAN infrastructure over time. It tries to deliver its SAN strategy in phases, to leverage new technologies once they are proven, and to help seamlessly integrate SAN technology into a company's IT infrastructure; all this while protecting your investments in application resources, servers, and storage.
- ☑ IBM's SAN solution uses Fiber Channel architecture for connectivity and device-level management.

## ❖ Chapter 6: ASP Security System Provisioning

### Security Policy

- ☑ An ASP needs to develop a general security policy that addresses how it manages and maintains the internal security posture of its infrastructure.
- ☑ A security policy defines how an ASP manages, protects, and distributes sensitive information and resources. Any ASP, before connecting to the Internet, should develop a usage policy that clearly identifies the solutions they will be using and exactly how those solutions will be used.

## Chapter 6 Continued

- ☑ An extension of the security policy is the *privacy policy*. The privacy policy should state what data the ASP considers to be confidential, and how that data can and cannot be used.

## Security Components

- ☑ As an ASP, to validate both the security policy and the privacy policy, a review of the various security mechanisms and methods used to implement those policies is required.
- ☑ One of the most important methods to provide accurate security is the ability to authenticate users and systems.
- ☑ A PIN provides another mechanism that you can use to enhance the security of a standard username and password system.
- ☑ Confidentiality is usually associated with data encryption mechanisms such as Secure Socket Layer (SSL) or Data Encryption Standard (DES), and targeted at protecting data as it traverses across a network, such as the Internet.

## Security Technologies and Attacks

- ☑ ASPs must deploy the best security technologies. Strong encryption is important, whether in the context of an SSL browser connection or a VPN connection.
- ☑ The two basic methods of VPN access are LAN-to-LAN VPNs and remote access VPNs.
- ☑ A perimeter firewall is a device, or software application, that controls access in to and out of a given network.
- ☑ Stateful inspection provides for the most robust of all firewall features.
- ☑ Embedded firewalls are software applications that are installed and run on a computer to guard it against attacks.
- ☑ Distributed denial of service (DDoS) is one of the newest and most troubling types of attack an ASP must face. This type of attack is perpetrated to cause the same undesired effects offered by DoS attacks, but on an even larger scale.

## Chapter 6 Continued

### Prevention Techniques

- ☑ As IP networking and the Internet began to come into widespread use, it became obvious that some companies used IP addresses for systems that were never intended to connect to the Internet. This meant that many of the dwindling IP addresses were wasted on private companies that used the addresses only to route internal traffic.
- ☑ Ingress filtering is used when these packets are filtered as they enter a network interface, and egress filtering is used when we filter these packets as they leave the interface.
- ☑ Most routers can be configured to limit the amount of data that will be processed for a particular time interval. This is known as *rate limiting*.
- ☑ It is possible to prevent most SYN attacks on your system using CARs to limit the amount of TCP traffic bursting allowed on your system. To accomplish this, you will need to configure rate limiting to allow for the full bandwidth of your connection, but reduce your *normal* and *excess bursting* sizes.

### Capturing Evidence

- ☑ If your organization has been the victim of an attack, it will be very important to capture and preserve as much evidence as possible. Any evidence you may be able to gather might prove useful in locating an attacker, and preventing further attack.
- ☑ Syslog is a software daemon that runs on a server to allow for logging of messages and events.
- ☑ Linux and SUN operating systems include an application called tcpdump that can be used to capture packets in real time.

## ❖ Chapter 7: Management and Monitoring

### The Effect of Outsourcing

- ☑ The service level agreement (SLA) allows the customer to set minimum (and maximum) limits to be met. There are three main areas in almost every SLA: Planning, Verification, and Troubleshooting.



## Chapter 7 Continued

- ☑ Frame Relay involves a number of system parameters that go beyond the standard parameters that can be monitored by the Simple Network Management Protocol (SNMP). Some of these elements cover the entire network, segmented networks, or even single circuits. The level at which an SLA can be defined depends entirely on the business need of the circuit.

## What Service Levels Should the Service Provider Consider?

- ☑ Most clients will want you to commit to a monthly guarantee of at least 99.5 (more often, 99.999) percent uptime. This guarantee generally includes all of the devices that are within your infrastructure, that connect to the local loop, or connect to the CPE. An uptime of 99.5 percent equals 3.6 total hours of downtime per month per site.
- ☑ Many of the largest companies guarantee a delay (round-trip) no greater than 300 milliseconds. You may be able to provide guarantees based on access line speeds, which can offer much lower delays for T1 and 64 kbps.
- ☑ Some service providers base effective throughput on the percentage of delivered frames based on a Committed Interface Rate (CIR) or frames that are labeled discard eligible (DE). Other providers base this calculation on the committed burst size rather than the excess burst size. You may be able to exclude configurations where the destination port is not configured to handle the bandwidth of the CIR.
- ☑ Response time can be whatever number of hours that you and the client agree upon. There is a pretty standard method that says that you will respond within four hours of reported outage. This also depends on the location of the service provider from the maintenance center. Usually this maintenance only covers CPE, as your facility will be handled on an internal basis.

## The Realities of Customer Compensation

- ☑ Many of your customers will want to know if you can find and fix issues (and potential issues) before they are affected. They will also most likely want to know if you will proactively fix issues, or wait for them to call and inform you. They will also wonder if you have the resources to meet the

## Chapter 7 Continued

demand of the time to resolution or repair that is included within their SLA. In the customer's mind, compensation for downtime is not the correct answer, nor will it ever be. They just want you to take care of them, so that they in turn can take care of their clients.

- ☑ What will your clients look for in these reports on SLAs? Here are some things that your clients will ask you to do:
  - Continually check that the WAN is capable of handling the services that they are providing.
  - Verify that service levels are being maintained. This request may require your ability to show monitoring in real time.
  - If services are not being met, then there must be an immediate path to resolution. This may be entirely your responsibility.
- ☑ Many tools are available to monitor the systems in the data center environment. These tools are generally used to collect usage statistics and the percentage of uptime for devices. These packages will also inform a centralized management station of the number of outages, the length of these outages, the mean time between failures (MTBF), and the mean time to repair (MTTR).
- ☑ By making your model more customer oriented, you can offer SLAs for things such as: emergency response, response time guarantees, call center availability, and remote troubleshooting.
- ☑ As the corporate infrastructure has evolved, so have the dynamics of the corporate network. What you are more apt to find in these changing times is an internal staff that handles and maintains very little of the overall network, remaining entirely within their walls or boundaries. External staff is comprised of the outsourced applications and infrastructure support. When you combine these two teams, you can encompass the range of support, including intranet-based Enterprise Resource Planning (ERP), electronic mail (e-mail), messaging, scheduling, desktop support, operating systems, remote access, security, and other miscellaneous company needs.

## Chapter 7 Continued

### How Service Providers Have Responded

- ☑ With all of the mission-critical applications that are available, many service providers are now offering services that are more advanced than the typical “leased line” connectivity that had been their bread and butter for so long. Leased lines were the lifelines to companies that needed direct access to their sites, and to their applications.

### The Operation Support System Model

- ☑ The Operations Support System (OSS) model usually refers to a system (or systems) that can perform the management necessary to maintain and monitor your SLA requirements. This model takes the following items into account: performance management, inventory control, system engineering, design, and support.
- ☑ In order to truly understand OSSs, you must first become familiar with some of the fundamental systems that are involved. These systems handle the functions of ordering, service fulfillment (such as voice, data, and other IP-based services), inventory, circuit provisioning, and activation.
- ☑ Many of today’s OSS solutions are considered commercial off-the-shelf (COTS) packages. These applications are able to offer some out-of-the-box utilities and are intended to be modified to meet customer needs. This customization could allow your company to integrate management capabilities and enable your customers to take advantage of your services, thus adding efficiency.

### Broadband Access Changes the Market

- ☑ Broadband access has changed the way we do business, and how we live at home. At this moment in time, DSL and cable are surpassing every other method of access across the United States. This isn’t to say that Frame or other connections are going to disappear; it is really saying that, like everything else, things change.
- ☑ Many of today’s service providers are struggling with the deployment of these technologies. It’s not because they don’t have the bandwidth; it’s because it is difficult to maintain and upgrade your infrastructure if you are unable to see your current copper allocation (for the local loop) and resource availability.

## Chapter 7 Continued

One of the ways that a central office (CO) can handle these issues is to have an up-to-date, dynamic inventory of provisioning.

- ☑ In order for a service provider to incorporate DSL within its infrastructure, there is the need to integrate two components: a splitter and a DSL Access Multiplexer (DSLAM). A splitter distributes voice traffic to the Plain Old Telephone System (POTS) cloud, and data traffic to the DSLAM. A DSLAM is able to communicate with the DSL router that is located on the customer's premises.

## Quality of Service

- ☑ Quality of Service (QoS) is a measurement of the service value. Measurement of QoS is very subjective; it depends on the technology on which it is implemented to see if there are acceptable levels of performance.
- ☑ You will need to maintain a high level of QoS to maintain and attract new customers. Therefore, you should implement and manage your solution so that it is capable of meeting your customers' expectations. QoS will vary from customer to customer, so tailor your SLAs to reflect client needs; for example, a bank that may need to implement high-speed transport (ATM) and VPNs.

## Management Systems for Your ASP

- ☑ Many of today's service providers use (at least at some level) the Telecommunications Management Network (TMN) model. The TMN model provides the outline for attaining interconnectivity and communications across diverse platforms and environments.
- ☑ TMN was developed by the International Telecommunications Union (ITU) as a tool to help support, manage, and deploy services. TMN was originally based on the common management information service element (CMISE).
- ☑ The TMN model outlines what is necessary to make your network infrastructure flexible, scalable, manageable, and highly available. TMN defines standard ways of handling management tasks and communications across networks. TMN allows you to distribute the appropriate levels for growth, efficiency, and communication performance.

## Chapter 7 Continued

### What Tools Do You Need to Automate TMN?

- ☑ A multitude of tools are available to automate the task of building TMN agent or manager applications. You can deploy and tailor the TMN agent and manager toolkits to match your company's GDMO/ASN.1 MIB representations. These products should have the following features in order to take advantage of the TMN model and to most productively support a TMN infrastructure: automated prototyping, conformance to all TMN standards, dynamic information modeling, Management Information Base (MIB), platform-independent interfaces and tools, Q adaption capability or compatibility, and system management functions (SMFs).

### The ASP Transformation

- ☑ To transform from an ISP to an ASP, you will need a service management solution that is designed specifically to manage the unique functions and processes of ASPs with carrier-class reliability and scalability.
- ☑ The ultimate goal is to build a unified system that automatically and dynamically builds and provisions a packaged service in response to customer clicks on a service portal icon even across multiple data centers and service sources. Standards-based interfaces are beginning to make that possible by allowing communications between provisioning systems and the applications.

### Pricing Models and Billing

- ☑ Usage-based billing is receiving a lot of attention; however, companies are struggling over what to measure, and how to measure it. If a company decides to measure usage, it must measure packets and relate the number of packets to some level of utilization. There are many methods of measurement. The interesting thing is, the way that the user data is gathered does not necessarily equate to how you present it back to the customer. If I say you use 75 units, how do I measure the units; is it the amount of bandwidth you use? The number of computer cycles you use? A formula that summarizes of all those? How much disk space you use?
- ☑ Pricing by transaction is gaining momentum. Still, defining a transaction and being able to capture the transactions for the billing system is no small task. Some applications could be open to pricing by the amount of data stored

## Chapter 7 Continued

within; for example, the number of customers stored within an application for a dentist office.

- ☑ Threshold pricing is another possible variation; for example, users pay a flat fee for usage up to a certain threshold. Beyond that, they would pay a small fee per unit (CPU cycles) used.
- ☑ The most common pricing model today is to charge a flat fee per month, often on a per-license/per-user basis. For the larger applications such as ERP software, some pricing occurs per seat/per license within the software.
- ☑ As an ASP provider, you will face various billing issues that are likely to be among your greatest challenges. Regardless of what is offered, the billing systems must go through many changes before they can effectively meet your billing needs in this new ASP business model.
- ☑ Directory services are the way to manage an installation of numerous servers. Many applications, though, are not directory enabled. It took Bell Laboratories the better part of 100 years to get telephone systems into a format that was reliable to handle millions of customers uninterrupted. Software as we know it is going to have to go through a massive transformation before the same can be said about software applications, especially in the ASP model.

## ❖ Chapter 8: Designing the Infrastructure

### Design Considerations

- ☑ There are generally three components when designing a large internetwork: data center networks, wide area networks (WAN), and remote users (in this case, your external clients).
- ☑ The *data center* is a building or set of buildings that house the infrastructure of your network.

### Site Considerations

- ☑ When you are building a new physical plant for your ASP, make sure that there is adequate space available and sufficient resources (power and cabling, as well as security) to suit your needs.

## Chapter 8 Continued

- ☑ *Routers* are Layer 3 network devices that connect separate networks and pass traffic between subnets.

## Designing with the Hierarchy in Mind

- ☑ One of the most beneficial tasks that you can perform in the design of your network is to create a hierarchical internetwork design that will modularize the elements of a large internetwork into layers of internetworking.
- ☑ Hierarchical internetworks are more scalable, because they allow you to grow your internetwork in a gradual way with the implementation of modules.
- ☑ The effect of broadcast traffic in your internetworks requires that you implement smaller groups of routers and switches, which will make your network more efficient.

## Frame Relay Internetwork Design Considerations

- ☑ A major concern when designing a Frame Relay implementation is scalability. As the number of remote clients and their links grows, your network must be able to grow to accommodate these growth spurts.
- ☑ Implementing a hierarchical mesh for Frame Relay environments can assist you in avoiding implementing an excessively large number of DLCIs.
- ☑ The cost-effective and strategic significance of the core network often forces network designers to implement a hybrid-meshed network for their WAN internetworks.

## Capacity Planning for Your Infrastructure

- ☑ If you have a general idea of where you stand for number of servers and expected growth, you can use those as a baseline for the capacity of your network.
- ☑ One of the best practices for planning is to map out where the different customer areas are located, and what the server count is going to be. Once these figures are determined, decide if the servers need one data link or multiple connections.

## Chapter 8 Continued

### Protocol Planning Concerns

- ☑ By determining the physical layout of the network, you will be able to map the correct topology and form a logical addressing scheme that will grow as your network grows.
- ☑ If your network is fairly simple in terms of the topology and number of routers, a distance-vector protocol such as RIP or IGRP (discussed later in this chapter) could work fine. If you're running a multivendor network, RIP, RIPv2, IS-IS, and OSPF are common protocols across many vendors' router implementations.

### Addressing Considerations

- ☑ The topology of a network is defined by sets of routers and the networks to which they connect. Routing protocols can also establish a logical topology depending on implementation.
- ☑ Broadcast traffic sets a practical limit to the size of the broadcast domain. Managing and troubleshooting a bridged campus becomes harder as the number of users increases because it adds to the broadcast domain.

### Application and Network Services

- ☑ When designing the data center, you should build the network as a modular building block using multilayer switching.
- ☑ Note that when using the Hot Standby Router Protocol (HSRP) (Cisco specific) or Virtual Router Redundancy Protocol (VRRP), which can also add redundancy, you should consider implementing Fast EtherChannel so you can scale bandwidth from Fast Ethernet, and from Gigabit Ethernet to Gigabit EtherChannel.

### Application-Aware Networking

- ☑ ASPs who want to deploy their applications need to realize that their success of mission-critical applications over both the internal LAN and clientele WAN is achieved by defining network policies, which assist in the apportioning of network resources with business objectives.



## Chapter 8 Continued

- ☑ Admission control is provided by a mechanism that can reject or remove applications based on user-defined policies. For example, a client can define a policy to temporarily stop the transmission of email packets, so that the mission-critical applications can use the necessary resources.

## Scalability Considerations

- ☑ Fast EtherChannel provides more efficient utilization of bandwidth by multiplexing multiple VLANs over one trunk.
- ☑ When designing your network, avoid creating STP loops in the backbone. STP takes 40 to 50 seconds to converge and does not allow for load balancing across multiple paths. When using ATM for your backbone, use PNNI to handle load balancing.

## Multimedia Services

- ☑ According to a study by the Telecommunications Industry Association, the multimedia application market (such as video on demand, VoIP, etc.) is expected to reach \$16 billion in 2001.
- ☑ Many of the new multimedia applications that customers want, require IP multicast for proper operation. Any network communication that needs to transmit information to multiple clients can benefit from the efficiency of multicast technologies.

## Planning for the Future Growth of Your Company's Infrastructure

- ☑ Distance routing protocols such as RIP, IGRP, SAP, and RTMP broadcast their complete routing tables on a periodic schedule. These updates will occur whether or not there have been any changes to the network.
- ☑ Cisco has implemented Data-Link Switching Plus (DLSw+) in their systems, which is an updated version of standard DLSw. This allows SNA frames from native SNA clients, which are then encapsulated in TCP/IP by a router.

## Chapter 8 Continued

### High-Availability Design

- ☑ Availability is the measurement of the uptime of database servers, mainframe applications, email, World Wide Web, multimedia, VoIP, and ERP (Enterprise Resource Planning).
- ☑ The network should be designed so that it can notify network operations personnel if there are failures, and be able to provide enough detail of the events that led up to the failure so that you can isolate and fix the issues.



3Com. *See* Dynamic Access  
 3-DNS controllers, 208  
 95th percentile measurement, 404  
 100BaseFX, 119  
 100BaseT, 119  
 286Gigabit Ethernet, interface  
 404 Object Not Found error, 212–213  
 1000BaseCX, 120  
 1000BaseLX, 119  
 1000BaseSX, 119

**A**

AboveNet Communications, 234  
 ABR. *See* Area border router  
 Abstract Syntax Notation One (ASN.1), 389, 396  
 Acceptable performance. *See* Service provider  
 Access. *See* Broadband delivery, 248  
   devices. *See* Customers getting, 386–387  
   layer, 426, 470, 537  
   overview, 481–482  
   usage, 427  
 lists, 318, 325  
   addition, 338  
   entry, 337  
   usage, 324, 335  
 method, 50  
 network connectivity, 92  
 providing, 295  
 Access control list (ACL), 229, 468. *See also* Extended ACL  
 Accounting, 390. *See also* Fault Configuration Accounting  
   Performance Security ability, 238  
   system, 232  
 ACK. *See* Acknowledgment  
 Acknowledgment (ACK), 353  
 ACL. *See* Access control list

Activation, 379–380  
 Active intercept, 353  
 Active Server Pages, 139  
 Active/Active configurations, 287  
 Active/Active features, 228  
 Active/Standby features, 228  
 Acts of God, 368  
 Adaptec, 127  
 Adapter Fault Tolerance (AFT), 129–131  
 Adapters. *See* Network  
 Adaptive Load Balancing (ALB) (Intel), 125  
 Address Resolution Protocol (ARP), 126  
 Addresses. *See* Internet Protocol; Media Access Control  
   spaces  
     classes, 450  
     filtering. *See* Request For Comment  
 Addressing, concerns, 450–453, 475  
 Administration  
   improvement. *See* Storage area network (SAN)  
   resources. *See* Network  
 Administrator-level access, 468  
 Admission control, 455–456  
 Advanced Micro Devices (AMD), 114  
 Advanced storage solutions, 161  
 Advanced store, 32  
 AFS file management, 31  
 AFT. *See* Adapter Fault Tolerance  
 Agent roles, 393  
 Agent-based content management, 206  
 Aggregation. *See* Link aggregation  
 Aggressive mode, 354  
   thresholds, 353–355  
 Agility Edge, 244

AHP. *See* Alliance Hosting Partner  
 AI. *See* Artificial Intelligence  
 AIP. *See* Application infrastructure provider  
 AIX (IBM), 31, 134, 140  
 Akamai, 28, 237, 249–250  
 ALB. *See* Adaptive Load Balancing  
 Alcatel networks, 26  
 Algorithms. *See* Asymmetric algorithms; Symmetric algorithms  
   types, 315–316  
 All fiber solutions, mixed solutions (contrast), 277–280  
 Allaire products, 138, 140  
 Alliance Hosting Partner (AHP), 44  
 Allotment, 379  
 ALPHA, 113, 135  
 AMD. *See* Advanced Micro Devices  
 American National Standards Institute (ANSI), 142  
 AMO. *See* Application maintenance outsourcing  
 Analytical applications, 97  
 Anonymity, 202  
 ANSI. *See* American National Standards Institute  
 Anti-virus  
   application, 169  
   software, 170, 171  
 AO. *See* Application outsourcing  
 Apache  
   HTTP Server, 138, 139  
   Software Foundation, 139  
 API. *See* Application programming interface  
 AppleTalk, 419. *See also* Routing Table  
   Maintenance Protocol  
 Appliance-based load-balancing product, 208  
 Application, 50. *See also* Database applications; World Wide Web

- attack, 332–333
- automatic synchronization, 245
- availability, 217
- capabilities, 98
- deployment, acceleration, 37, 87
- development, 14
  - acceleration, 87
- hosing, 70
- infrastructure hosting, 92
- integration, 95
- layer (Layer 7), 17, 19–21, 204, 233
  - attacks, 323
- management, 92
  - services, 373
- monitoring, 95
- movement, 244
- real-time updating, 245
- services, 453–454, 475
- SLAs, 374–375
- software, 133
  - package acquisition, 9
  - types, 137–142, 181–182
  - upgrades, cost, 9, 12
- Application infrastructure provider (AIP), 94–95
- Application maintenance outsourcing (AMO), 3, 7, 15
- Application outsourcing (AO), 7, 15
- Application programming interface (API), 27. *See also* Dynamic APIs
  - functionality/gateways, 383–384
- Application Service Provider (ASP), 6, 91–94, 188. *See also* Pure ASP; Pure-play ASP
  - business model
    - barriers, 37–40
    - predictability, 43
    - strategies, 38–40
  - conversion
    - business drivers, 34–45, 55
    - case, 82–94, 105
  - customer value proposition, 86–88
  - firms, types, 13–16, 53–54
  - host services, 12
  - implementation, 10–11
    - time, reduction, 87
  - Industry Consortium, 34
  - infrastructure operations, 399–401
  - introduction, 2–3
  - management systems, 388–396, 412
  - management tools
    - deployment, industry examples, 398–399
  - model
    - business factors, 34–35
    - implementation, 49, 56
  - operating, 317
  - performance issues, 45–48, 55
  - platform, choosing, 22–33, 54–55
  - rollout, 100
  - security system provisioning, 303
    - introduction, 304–305
  - services, 91
    - improvement, 77
  - software solutions, 133–137, 181
  - strategy, 98
  - technical factors, 36–37
  - terms, definitions, 5–8, 52
  - transformation, 397–401, 412–413
  - viability, elements, 8–12, 52–53
- Application Service Provider (ASP) network
  - infrastructure, 416
  - introduction, 480–481
  - sample configuration, 479
- Application-aware networking, 455–458, 475
- Application-processing delays, 49
- Application-specific integrated circuit (ASIC), 466, 467
- ARCserver, 162
- Area border router (ABR), 459
- AristaSoft, 40
- ARP. *See* Address Resolution Protocol
- Arrays. *See* Multihost arrays
- Artificial Intelligence (AI), 324
- ASCII, 19
- ASIC. *See* Application-specific integrated circuit
- ASN.1. *See* Abstract Syntax Notation One
- ASO Industry Consortium, 375
- ASP. *See* Application Service Provider
  - services, 79
- Asymmetric algorithms, 316
- Asynchronous mirroring, 291
- Asynchronous Transfer Mode (ATM), 62, 385, 419, 456
  - backbones, 458
  - map statement, 521
  - port adapter, 517
  - usage, 459
- Asynchronous Transfer Mode Forum (ATMF), 390
- Athlon, 116
- ATM. *See* Asynchronous Transfer Mode
- ATME. *See* Asynchronous Transfer Mode Forum
- Attrieve, 30
- AT&T, 240, 480
- Attached storage. *See* Infrastructure; Network attached storage; Storage area network
- Attack signatures, 331
- Attacker, identification, 317
- Attacks. *See* Application; Buffer overflow attacks; Denial of Service; Distributed Denial of Service; Fraggle attack; Infrastructure; Internet Protocol; Physical

- attacks; Smurf attack; Synchronization types, 332–343
  - Audio broadcast, 460
  - Auditing. *See* Security
  - Authentication, 309–313. *See also* User mechanisms, 310
  - Automated prototyping, 396
  - Availability. *See* Application; High availability; LocalDirector; Server improvement. *See* Storage area network management, 281
  - Average network delay, 367, 369, 370
  - Average PVC delay, 367, 369, 370
  - Average rate, 349, 350
- B**
- Baan, 41
  - Baan Oracle PeopleSoft SAP (BOPS), 43
  - Backbone, 92. *See also* Asynchronous Transfer Mode; Internet Service Provider operations. *See* Enterprise backbone operations
  - Back-office ERP solutions, 40
  - Back-out plan, 176
  - Back-plane speed support, 207
  - Backup Exec (Veritas), 162
  - Backups. *See* Data; Differential backup; Full backup; Incremental backup; Remote backup effect. *See* Data frequency, 164 scheduling, 163–168
  - Backward Explicit Congestion Notification (BECN), 441
  - Bandwidth, 365, 443. *See also* Modal bandwidth; Scaling access, 70 allocation, 49 amount, 191. *See also* Outbound bandwidth availability. *See* Network cost, decrease, 36 expense, 216 limitations, 188 needs, reduction, 250 reduction. *See* Upstream bandwidth requirements, 172, 173 usage, 430, 448 reduction, 195
  - Bandwidth-intensive collaboration capabilities, 241
  - Barriers to entry. *See* Entry
  - Baseline metrics, 373
  - Basic store, 32
  - Bastion network, 321, 328–329 servers, 329
  - BBS. *See* Bulletin Board System
  - BEA Systems, 30
  - BECN. *See* Backward Explicit Congestion Notification
  - Bell Laboratories, 133
  - Bellcore, 390
  - Beowulf Cluster, 144
  - Berkeley Software Distribution (BSD), 135
  - BGP. *See* Border Gateway Protocol
  - BIG-IP (F5 Networks), 208
  - Billable-hours approach, 14
  - Billing, 95, 404–405. *See also* Pricing ability, 238 management, 405–406 system, 232
  - BIND, 332
  - bind (command), 221, 225, 227
  - Birds of a Feather (BOF), 242
  - Black hole creation, 195 policy, 202
  - Blacklist filtering. *See* Content; Sites
  - BLECs. *See* Building Local Exchange Carriers
  - Bluestone, 30
  - BMC Software, 375
  - BML. *See* Business Management Layer
  - BOF. *See* Birds of a Feather
  - BOPS. *See* Baan Oracle PeopleSoft SAP
  - Border Gateway Protocol (BGP), 448, 503, 520 connections, 505 enabling, 504 protocol, 490
  - Bottlenecks, 216, 427
  - BPO. *See* Business process outsourcing
  - Break/fix, 176–177
  - Break-ins, 332
  - Bridge, 425
  - Bridged protocol needs, 467–468
  - Bridging. *See* Multilayer model
  - Bright Tiger, 30
  - Broadband, 67 access, 386–387, 411 usage, 64–65
  - Broadcast domain. *See* Spanning Tree Protocol issues, 417, 431–432. *See also* Frame relay media, 439 queue, creation. *See* Interface traffic, 435 levels, 434 optimization, 427, 433
  - Broadcast and unknown server (BUS), 463
  - Broadvision, 43
  - Brocade, 275
  - Browser-based interface, 202
  - Browser-defined differences. *See* Content
  - Browser-enabled desktops, 35
  - Browsers. *See* World Wide Web
  - Brute force, 312
  - BSD. *See* Berkeley Software Distribution

BSS, 400  
 Buffer overflow attack, 334–335  
 Buffer size, 336  
 Building Local Exchange Carriers (BLECs), 63, 67  
 Built-in SSH, 229  
 Bulletin Board System (BBS), 169  
 Bundled services, 76  
 Burst size. *See* Excess burst size; Normal burst size  
 BUS. *See* Broadcast and unknown server  
 Business  
   drivers. *See* Application Service Provider  
   factors. *See* Application Service Provider  
   flexibility, increase. *See* Storage area network  
   models, 12–13, 53, 94–96. *See also* Application Service Provider; Long-term business model  
   objectives, resource focus. *See* Core business objectives  
   offerings, 12–13, 53  
   process  
     consulting, 14  
     design. *See* Implementation/business process design  
   requirements. *See* Service provider  
   scalability, impact, 282–288, 299  
 Business cases, 59  
   introduction, 60–61  
 Business Management Layer (BML), 394  
 Business process outsourcing (BPO), 7, 15  
 Business-to-business extranet, 401

## C

C++, 141  
 C (language), 133  
 Cable, 81, 385  
   modems, 61  
 Cable & Wireless, 240  
 Cache. *See* Proxy  
   appliance makers, 28–29  
   appliances  
     cost effectiveness, 201  
     definition, 201–204  
     installation/management, ease, 201–202  
     performance/speed, 201, 203–204  
     scalability/flexibility, 201, 203  
   economic potential, 199  
   hierarchies, 199–200. *See also* HyperText Transfer Protocol  
   locations/placement, 199  
   requesting cached objects, 196  
   server, 232  
 Caching. *See* Nontransparent caching; Transparent caching; World Wide Web  
   benefit, 191–192  
   definition, 190. *See also* World Wide Web  
   deployment models, 197–204, 252  
   solution, 192–194  
   key requirements, 195–196  
   systems, 207–208  
   usage, 239  
 CAD. *See* Computer aided design  
 Call center availability, 375  
 Candle, 375  
 Canonical names (C-names), 236  
 Capacity management, 281  
 Capacity planning. *See* Infrastructure  
 Capital  
   markets, 64  
   requirements, 95  
   restricted access, 68, 89  
 CAR. *See* Committed Access Rate  
 Carrier Sense Multiple Access Collision Detect (CSMA/CD), 118, 452  
 Catalyst series  
   router, 130  
   switch, 127  
 C-bit framing, 502  
 CDN. *See* Content Delivery Network  
 CDNP. *See* Content Delivery Network Peering  
 CDP. *See* Cisco Discover Protocol  
 CDS. *See* Content Delivery Suite  
 C.E. Unterberg, Towbin, 92, 94  
 CEF. *See* Cisco express forwarding  
 Central office (CO), 386  
 Central Processing Unit (CPU), 46, 112–114, 190  
   clock cycles, 264  
   cycles, 273, 402, 432  
   memory utilization, 206  
   overhead, 459  
   utilization, 204, 216, 465  
 CGMP. *See* Cisco Group Multicast Protocol  
 Chainlink, 399  
 CHAP secrets, 486  
 Child tape rotation scheme, 165–166  
 Choke point, 199  
 Churn. *See* Customers  
 CIBER Enterprise Outsourcing, 40, 41  
 CIFS. *See* Common Internet File System  
 CIGP. *See* Common Interconnection Gateway Platform  
 CIR. *See* Committed Interface Rate

- CISC. *See* Complex Instruction Set Computers
- Cisco. *See* Content Delivery Network; Fast EtherChannel; Gigabit EtherChannel; LocalDirector router, 345, 352
- Cisco Discover Protocol (CDP), 513
- Cisco express forwarding (CEF), 495
- Cisco Group Multicast Protocol (CGMP), 461–462
- Cisco Resource Manager (CRM), 471
- Cisco Systems, 26, 374, 400
  - 7200 router, configuration, 486–509
  - configuration, commands/references, 485–553
  - Gigabit Switch Router (GSR), configuration, 509–537
  - MGX router, configuration, 537–552
- Citrix Systems, 172, 400
- Citrix-based environment, 399
- Class of Service (CoS), 456, 527
- class-map match-all (command), 488–492
- class-map match-any (command), 488–492
- Clearinghouse services, providing ability. *See* Third-party clearinghouse services
- Clear-text passwords, 311
- CLEC. *See* Competitive Local Exchange Carrier
- CLI. *See* Command-line interface
- Clients
  - availability, guarantee, 214–215
  - reference, 100
  - requests, responsiveness, 196
- Client/server
  - applications, 268
  - environment, 36
  - overhead, 173
- CLNS. *See* Connectionless Network Services
- Cluster (Digital), 144
- Cluster Enterprise (Legato), 145
- Cluster implementations, 47
- Cluster Server
  - Microsoft, 145
  - Veritas, 144
- Clustering, 47–48, 202, 206, 208
  - solutions. *See* Lower-end clustering solutions
  - systems, 208
  - technology, 218
- Clusters (CustomSystems), 144
- CMIP. *See* Common Management Information Protocol
- CMIS. *See* Common Management Information Services
- CMISE. *See* Common management information service element
- C-names. *See* Canonical names
- CO. *See* Central office
- Cogent Communications, 66
- ColdFusion, 140
- Collaboration
  - services, 97
  - applications, 97
- Collision domains, 424
- Collocation hosting, 92
- Command-line interface (CLI), 202, 228, 247, 480
- Commercial off-the-shelf (COTS), 383
- Committed Access Rate (CAR), 337, 338, 342, 512
  - usage, 349
- Committed Interface Rate (CIR), 370, 432, 434, 440–442
- Commoditization, onset, 63–65
- Commoditized offering, 67, 89
  - pricing, improvement, 70
- Common Interconnection Gateway Platform (CIGP), 384
- Common Internet File System (CIFS), 148, 159
- Common Management Information Protocol (CMIP), 380, 389, 393
- Common management information service element (CMISE), 388
- Common Management Information Services (CMIS), 390
- Common Object Request Broker Architecture (CORBA), 382
- Communication links, 4
- Compaq Computer Corporation, 24, 25, 48, 117, 127, 275, 374, 400. *See also* ProLiant
- Competitive Local Exchange Carrier (CLEC), 381–383
  - interfaces, 384
- Complex Instruction Set Computers (CISC), 112
  - microprocessors, 113
  - processors, 114
- Computer
  - breaches, 305
  - viruses, 305
- Computer aided design (CAD), 278
- Compuware, 375
- Concentrator, 424
- Confidentiality, 309
  - protection, 313–317
- Configuration, 390. *See also*
  - Fault Configuration Accounting
  - Performance Security management, 281
- Configuration Terminal (command), 220



- Congestion avoidance, 457–458
- Connection speeds, 188, 434
- Connectionless Network Services (CLNS), 511
- Connectivity. *See* Access; Internet
- Consulting. *See* Business; Information Technology; Strategic management consulting
- Consumers, 241
- Content. *See* Mission-critical Internet-based content
  - automatic synchronization, 245
  - blacklist filtering, 202
  - browser-defined differences, 202
  - delivery, 231–233
  - distribution, 247, 248
    - integration, 245
    - internetworking, 240
  - failure, 213, 214
  - management, 247, 248. *See also* Agent-based content management tools, 231
    - visibility, 236–237
  - monitoring tools, 231
  - movement, 244
  - providers, 230, 234–235
  - publishers requirement. *See* Content Delivery Network
  - real-time updating, 245
  - routing, 247, 248
  - signaling technologies, 238
  - switching, 247
  - usage visibility, 236–237
  - user-defined differences, 202
- Content Alliance, 238, 241–243
- Content Bridge Alliance, 238, 241–244
- Content Delivery Network (CDN) (Cisco), 189, 230–244, 247–248, 254. *See also* Distributed CDN; Facilities-based CDN; Hybrid CDN; Multinetwork CDN
- component product makers, interaction. *See* Network
- content publishers, requirement, 235–238
- deployment basics/considerations, 239
- function, explanation, 232–233
- functional components, 232
- Group, 248
- industry standardization efforts, 242–244
- migrations, 237
- need/benefit, 233–235
- network infrastructure, 236
- product manufacturers, 240–241
- server provider specialist, 233
- service providers, 235–238
  - interaction. *See* Network requirements, 238–241
- services landscape, 241
- solutions, vendor impact, 244–250, 254
- Content Delivery Network Peering (CDNP), 242
- Content Delivery Suite (CDS) (Inktomi), 244–250
- Content Distributor (Inktomi), 245–247
- Content Manager (Inktomi), 245–247
- Content routing, 202
  - definition, 189
- Content-aware applications, 248
- Conversion. *See* Application Service Provider
  - problems, 48, 56
- Cookie ID-based switching, 228
- Coordination efforts, improvement, 35
- Copy, marketing, 50
- CORBA. *See* Common Object Request Broker Architecture
- Core business objectives, resource focus, 86–87
- Core competencies, 75, 78–80, 95
- Core layer, 426, 470, 484–485
  - usage, 427
- Corio, 39, 398, 399
- Corporate IT, 45
- Corporate LAN, 193
- CoS. *See* Class of Service
- Cost of Ownership, life cycle, 8–12
- Cost structure, 100
- COTS. *See* Commercial off-the-shelf
- Covad Communications, 65
- CPE. *See* Customer premise equipment
- CPU. *See* Central Processing Unit
- CRC. *See* Cyclic redundancy check
- CRM. *See* Cisco Resource Manager; Customer Relationship Management
- Cryptographic considerations, 316–317
- Cryptography. *See* Public-key cryptography
- CSMA/CD. *See* Carrier Sense Multiple Access Collision Detect
- CustomAuctions, 64
- Customer premise equipment (CPE), 365, 371
- Customer Relationship Management (CRM), 37, 41, 83, 96, 460
  - applications, 90, 99
  - implementation time, reduction, 87
  - solutions, 42
  - suite, 10
- Customer-induced downtime, 368
- Customer-oriented companies, 375–376
- Customers
  - access devices, 368

- base, reaching ability, 231
  - churn, 67, 89
  - compensation, realities, 371–376, 409–410
  - confidence, building. *See* Security
  - data, loss, 160
  - DSU/CSU, 368
  - examination. *See* Implemented SLA
  - issues, 99–102
  - relationship management, 97
  - router, 368
  - satisfaction. *See* Long-term customer satisfaction
  - usage pattern, 175
  - value proposition. *See* Application Service Provider
  - CustomSystems. *See* Clusters
  - CyberCop Scanner, 319
  - Cyclic redundancy check (CRC), 513
- D**
- Daemons, 158
    - running, 356
  - Data. *See* Mission-critical data
    - analysis, 206
    - archiving, 278
    - backup, 289–290
      - services, 110
      - strategy, 160
    - backups, effect, 159–168, 182
    - integrity, enhancement, 295
    - link layer (Layer 2), 17, 18
      - switching, 462, 466
    - loss, 160. *See also* Customers
    - lost. *See* Scheduled maintenance
    - network. *See* Packet-switching data network
    - recovery, 293
    - sabotage, 305
    - services
      - activation, 386
      - provisioning, 385
      - support, 384–386
    - sharing, 265
    - storage, 68
    - switching, 2
    - synchronization, 206
    - throughput rates, 69
    - traffic explosion, 25–26
    - transmission, 370
  - Data center, 92. *See also* Terminal data centers
    - definition, 419
    - environment, end-to-end network services, 416
    - networks, 418
      - design, 454
      - services, 95
  - Data Communication Network (DCN), 392, 393
  - Data Encryption Standard (DES), 313
    - encryption, 315
  - Data link connection identifier (DLCI), 433–435, 437–441
  - Database
    - applications, 141–142
    - replication, 460
  - Data-based infrastructure, 25
  - Data-Link Switching Plus (DLSw+), 467
  - DCE. *See* Distributed Computing Environment
  - DCN. *See* Data Communication Network
  - D-COM. *See* Distributed Component Object Model
  - DDoS. *See* Distributed Denial of Service
  - DE. *See* Discard eligibility
  - DEC. *See* Digital Equipment Corporation
  - Decision-making capabilities, 444
  - Decision-making process, 163, 267
  - Delay, 444
  - Dell (computers), 117
  - Deloitte Consulting, 95
  - Demilitarized zone (DMZ), 321
  - Denial of Service (DoS). *See* Distributed Denial of Service
    - attacks, 305, 333–334, 339
    - avoidance, 211
  - Dense mode, 461
  - Departmental-level servers, 25
  - DES. *See* Data Encryption Standard
  - Design and Assign system, 379
  - Dial-up
    - access, 65, 69
    - services, 63
  - Dictionary attack, 487
  - Differential backup, 164, 165
  - Differentiation, 71, 78, 104
  - Digex, 240
  - Digital. *See* Cluster
  - Digital certificates, 310–311
  - Digital Equipment Corporation (DEC), 113
  - Digital Subscriber Line Access Multiplexer (DSLAM), 386–387
  - Digital Subscriber Line (DSL), 61, 62, 74–75, 304, 385
    - concerns, 387
    - offerings, 67
    - providers, 65
    - router, 386
  - Digital Versatile Disk (DVD), 61
  - Directed broadcast functionality, 337
  - Directly attached storage. *See* Infrastructure
  - DirectPC satellite platforms, 240
  - Dirty network, 321
  - Disaster recovery purposes, 278
  - Discard eligibility (DE), 370, 441
  - Disk Operating System (DOS), 136

- Distance Vector Multicast Routing Protocol (DVMRP), 462
  - Distance vector protocol, 466
  - Distributed CDN, 239
  - Distributed Component Object Model (D-COM), 382
  - Distributed computing, 268
  - Distributed Computing Environment (DCE), 31
  - Distributed Denial of Service (DDoS), 304–305
    - attack, 339–343
    - daemon, 340
    - software, 343
    - tactics, 349
  - Distributed load balancing, 204–205
  - Distributed Lock Manager (DLM), 144
  - Distribution layer, 426, 470, 482–484
    - usage, 427
  - Distribution services, 232
  - DLCI. *See* Data link connection identifier
  - DLM. *See* Distributed Lock Manager
  - DLSw+. *See* Data-Link Switching Plus
  - DMZ. *See* Demilitarized zone
  - DNS. *See* Domain Name System
  - Document management, 68, 70, 77, 97–99
  - Domain
    - expansion, 202
    - servers. *See* Storage
  - Domain Name System (DNS), 2, 6, 93, 328, 332. *See also* Round-robin DNS entries, 223
    - round-robin rotation, 212
    - specification, 494
    - tables, 220
  - Domino Advanced Enterprise Server, 140
  - Domino products, 139
  - DOS. *See* Disk Operating System
  - DoS. *See* Denial of Service
  - Double-Take (Network Specialists), 145
  - Downtime. *See* Customer-induced downtime
  - Drop mode, 353, 354
  - DSL. *See* Digital Subscriber Line
  - DSLAM. *See* Digital Subscriber Line Access Multiplexer
  - DSU/CSU. *See* Customers
  - Dual servers, multihomed servers (contrast), 131–133
  - Dual-homed servers, 132
  - Dual-loop configuration, 48
  - Duron, 116
  - DVD. *See* Digital Versatile Disk
  - DVMRP. *See* Distance Vector Multicast Routing Protocol
  - Dynamic access, 126–127
  - Dynamic Access (3Com), 125
  - Dynamic APIs, 382
  - Dynamic information modeling, 396
  - Dynamic Web pages, 189
- E**
- Earthlink, 240
  - EBay, 64
  - EBCDIC, 19
  - E-business. *See* Electronic business
  - ECC. *See* Error checking and correcting
  - E-commerce. *See* Electronic commerce
  - EDI. *See* Electronic data interchange
  - EDS. *See* Electronic Data Systems
  - Education/training providers, 13, 16
  - Effective throughput, 367, 369, 370
  - Egress filtering, 346–348
  - EIDE. *See* Enhanced Integrated Drive Electronics
  - EIGRP. *See* Enhanced Interior Gateway Routing Protocol
  - ELANs. *See* Emulated LANs
  - E-learning. *See* Electronic learning
  - Electronic business (E-business), 267
    - growth, 268
    - solutions, 36
  - Electronic commerce (E-commerce), 96, 97
    - applications, 46
    - capabilities, 238
    - servers, 211
    - sites, owners, 234
    - solutions, 36
    - transactions, 230
  - Electronic data interchange (EDI), 382
  - Electronic Data Systems (EDS), 23, 44
  - Electronic learning (E-learning) developers, 234
  - Electronic mail (E-mail), 110, 160, 375
    - loss, 160
    - messages, sending, 334
    - usage, 229
  - Electronic Privacy Information Center (EPIC), 308
  - Electronic training (E-training) content, 241
  - Element management layer (EML), 390, 394
  - Element Management System (EMS), 395
  - E-mail. *See* Electronic mail
  - Embedded firewalls, 328–329
    - usage, 327
  - EMC, 117, 275
  - Emergency response, 375
  - Emerging technologies, acceptance, 37

- EML. *See* Element management layer
  - EMS. *See* Element Management System
  - Emulated LANs (ELANs), 459, 462–464
  - enable (command), 220
  - enable password (command), 486
  - enable secret (command), 487
  - Encryption. *See* Secure Sockets Layer
    - key, 314
  - End-to-end
    - application, 455
      - testing tools, 33
    - connectivity, 280
    - customer care, 364
    - delivery solution, 91
    - model, 366
    - network
      - services. *See* Data center solution, 468
    - offering, 100
    - routing/switching solutions, 27
    - services, 26
    - solution, 26, 28, 89, 94, 279
    - topology, 365
  - End-user satisfaction, 173
  - Enemies, identification, 327
  - Engineering/provisioning, 379
  - Enhanced Integrated Drive Electronics (EIDE), 116, 261–263
    - usage, 285
  - Enhanced Interior Gateway Routing Protocol (EIGRP), 431, 447–449, 464
  - Enterprise Monitoring Package, 375
  - Enterprise resource management, 97
  - Enterprise Resource Planning (ERP), 8–10, 34, 97, 460, 468. *See also* Back-office ERP solutions; Intranet-based ERP
    - application, 90, 99
    - implementation time,
      - reduction, 87
    - outsourcing divisions, 40
    - servicing process, 215
    - software, 402
      - applications, 38
  - Enterprise Storage Resources Management (ESRM), 280, 281
  - Enterprise Systems Connection (ESCON), 116–118, 275–276
    - interfaces, 276
  - Enterprises, 241
  - Entertainment companies, 234
  - Entry, barriers, 71, 104
  - EPIC. *See* Electronic Privacy Information Center
  - ERP. *See* Enterprise Resource Planning
  - Error checking and correcting (ECC), 291, 293
  - Error messages, 210
  - ESCON. *See* Enterprise Systems Connection
  - ESRM. *See* Enterprise Storage Resources Management
  - EtherChannel, 127. *See also* Fast EtherChannel
    - configuration, 286
  - Ethernet, 81, 431. *See also* Fast Ethernet; Gigabit ethernet
    - technologies, 118
    - usage, 452
  - Ethernet WAN/MAN
    - services, 67
  - E\*Trade, 30
  - E-training. *See* Electronic training
  - ETSI. *See* European Telecommunications Standards Institute
  - European Telecommunications Standards Institute (ETSI), 390
  - Evidence, capturing, 355–357, 360–361
  - Excess burst size, 349
  - Exodus Communications, 42, 234
  - Extended ACL, 229
  - External protocols, 448
  - Externally provided local loop, 368
  - Extranet, 138. *See also* Business-to-business extranet
  - Extreme Networks, 26, 480
- ## F
- F5 Networks, 26, 27, 249–250. *See also* BIG-IP
  - F5 products, 250
  - F5 usage. *See* Load balancing
  - Facilities-based CDN, 239
  - Facilities-based providers, 239
  - Facility experience, 100
  - Failover, 47
    - capabilities, 468
    - mechanism, 218
    - protection, 209
  - Farm, scaling. *See* Server
  - Fast EtherChannel (Cisco), 125, 458
  - Fast Ethernet, 123, 125, 127, 419
    - adapter, 285
  - Fast SCSI, 277
  - Fast serial interface processor (FSIP), 513
  - Fault, 390. *See also* Fault Configuration Accounting Performance Security
    - Fault Configuration Accounting Performance Security (FCAPS), 390
  - Fault tolerance, 124, 201, 202–203, 258. *See also* Adapters
    - features/issues, 288–295, 300
    - improvement, 124
  - Fault-tolerant system, 288, 289
  - FC-AL. *See* Fiber Channel Arbitrated Loop

- FCAPS. *See* Fault Configuration Accounting Performance Security
- FCC. *See* Federal Communications Commission
- FDDI, 467
- FECN. *See* Forward Explicit Congestion Notification
- Federal Communications Commission (FCC), 381
- Federal Trade Commission (FTC), 309
- FEP. *See* Front-end processor
- Fiber Channel, 267
  - adapter, 285
  - benefits, 276–277
  - connector, 279
  - implementations, 280
  - interfaces, 276
  - networks, 279
  - SCSI comparison, 275–280
  - technology, 268
- Fiber Channel Arbitrated Loop (FC-AL), 275, 276
- Fiber Channel-to-SCSI bridge products, 275
- Fiber Communications Channel, 276
- Fiber Connectivity (FICON), 116–118
- Fiber-optic cable, 120, 278, 338
- Fibre Channel, 116–117
- FICON. *See* Fiber Connectivity
- File Transfer Protocol (FTP), 138, 159, 197, 203, 332
  - file transfers, 194
  - protocol, 489
  - server, 137, 194
  - services, 139
  - usage, 229
- Filtering. *See* Egress filtering; Ingress filtering; Packet; Request For Comment
- Finance services, 96
- Financial
  - fraud, 305
  - information, loss, 160
  - layer (Layer 10), 22
  - viability, 98
- Fireproofing, 421
- Firewalls, 313. *See also* Perimeter firewalls
  - auditing, 319
  - capability, 211
  - implementations, 322
- First Sense Software, 375
- Five nines, 38, 46–47, 69
- Fixed-price engagements, 14
- Flash favoring groups, 244
- Flexibility. *See* Cache
- Flood attacks, 350–351
- Flow control, 18
- flush (command), 541
- Forrester Research, 3
- Forward Explicit Congestion Notification (FECN), 441
- Forward proxy, 198, 203
- Foundry Networks, 26, 27, 480. *See also* ServerIron
- Fraggle attack, 334, 338
- Fragmentation attack. *See* Internet Protocol
- Frame relay, 384–385
  - benefits, 376
  - circuits, SLA components, 366–368
  - interface, 438
  - internetworks. *See* Hierarchical meshed frame relay internetworks; Hybrid-meshed frame relay internetworks
  - hierarchical design, 433–434
  - internetworks, design considerations, 432–442, 474
  - map class name, 508
  - networks, 376
    - broadcast issues, 439–440
    - regional topologies, 437–439
    - switches, 441
- Free Software Foundation, 135
- FreeFlow, 250
- Front-end processor (FEP), 468
- FSIP. *See* Fast serial interface processor
- FTC. *See* Federal Trade Commission
- FTP. *See* File Transfer Protocol
- Full backup, 164
- Full duplex
  - configuration, 285
  - contrast. *See* Half duplex
  - conversation, 122
  - operation, 122
- Fully meshed topologies, 428–430, 436–438
- Functionality, 290
- FutureLink, 398
- FutureLink Distribution, 40
- ## G
- Gateways, 276. *See also* Application programming interface
- GDMO. *See* Guideline for Definition of Managed Objects
- Genuity, 240
- GIF. *See* Graphics Interchange Format
- Gigabit EtherChannel (Cisco), 125
- Gigabit Ethernet, 74–75, 119, 125–127, 275, 419
  - adapter, 285, 287
  - port, 457
- Gigabit Switch Router (GSR), 509, 522, 532
  - configuration. *See* Cisco Systems
- GigaMAN, 66
- Global Recruiting Solutions, 40
- Globalization, 267
- GNU, 135

- Grandparent tape rotation scheme, 165–166
  - Granularity. *See* Quality of Service
  - Graphical user interface (GUI), 36, 247, 379. *See also* Management GUI; World Wide Web
    - management interfaces, 197
  - Graphics Interchange Format (GIF), 19
  - Great Plains Software, 42, 44, 400
  - Greenwich Mean Time (GMT), 510
  - GSR. *See* Gigabit Switch Router
  - GUI. *See* Graphical user interface
  - Guideline for Definition of Managed Objects (GDMO), 389, 396
- H**
- HA. *See* High availability
  - Half duplex, full duplex (contrast), 120–123
  - Hardware, 23–25, 112. *See also* Server
    - acquisition, initial cost, 9
    - interfaces, 426
    - maintenance/costs, 9, 10
    - management, 427
    - packages, 260
    - short-circuiting, 338
    - upgrades
      - cost, 11
      - performing, 174
  - Hardware-based routing, 467
  - Hardware-software network appliance, 208
  - Head-end layer, 484
  - Heat ventilation and air conditioning (HVAC), 422
  - Hello-timer interval. *See* Internet Group Management Protocol
  - Help Desk, 12
  - Heterogeneous operating systems, 172–173
  - Hewlett-Packard (HP), 24, 46–47, 127, 140, 374. *See also* Openview; WebQoS
    - HP-UX, 134, 144
  - Hierarchical design. *See* Frame relay
  - Hierarchical internetworks
    - design flexibility, 427
    - manageability, 427
    - scalability, 426–427
  - Hierarchical meshed frame relay internetworks, 434–436
  - Hierarchical service policy, 518
  - Hierarchy, impact. *See* Infrastructure design
  - High availability (HA), 24, 217–218, 262, 269, 365
    - design, 469–471, 476–477
    - implementation, considerations, 469–471
  - High-end system, 160
  - High-margin revenue streams, 78
  - High-speed packet-switched infrastructures, 65
  - High-speed switched links, 454
  - HIP. *See* HSSI interface processor
  - Host
    - independence, 259–260
    - services. *See* Application service provider
  - Host-based
    - security, 261
    - solution, 271
  - Hosting. *See* Collocation
    - hosting; Infrastructure experience, 100
    - service providers, 247
    - services, 76, 79
  - Host-name expansion, 202
  - Hot Standby Routing Protocol (HSRP), 217, 454, 471, 499
  - Hot-spare load balancing, 208
  - Hot-swappable
    - component, 175
    - modules, 283
  - HP. *See* Hewlett-Packard
  - HP-UX. *See* Hewlett-Packard
  - HSRP. *See* Hot Standby Routing Protocol
  - HSSI interface processor (HIP), 513
  - HTML. *See* HyperText Markup Language
  - HTRC Group, 234
  - HTTP. *See* HyperText Transfer Protocol
  - Hub and spoke design, 428
  - Hubs, 267, 276, 424
    - router, 430
  - Hughes Network Systems, 240
  - Human resources, 96
  - Hunter Group, 39
  - HVAC. *See* Heat ventilation and air conditioning
  - Hybrid CDN, 239
  - Hybrid-meshed frame relay internetworks, 436–437
  - HyperText Markup Language (HTML). *See* Static HTML
    - document, 139
    - format, 139
    - pages, 32
  - HyperText Transfer Protocol (HTTP), 137, 159, 195–197, 332
    - attacks, 351
    - cache hierarchy, 202
    - header information, 210
    - host field, 210
    - objects, 236
    - requests, 195
    - server, 193
    - sessions, 204
    - traffic, 197, 215
    - usage, 229, 326
    - version 1.1, 140, 203
    - versions 0.9, 203

- I**
- IANA. *See* Internet Assigned Numbers Authority
  - IBM, 24, 31, 46, 117. *See also* AIX; Storage area network
    - S/390, 135
  - ICA. *See* Independent Computing Architecture
  - ICMP. *See* Internet Control Message Protocol
  - Icon CMT, 41
  - ID-based switching. *See* Cookie ID-based switching; Secure Sockets Layer
  - IDC. *See* International Data Corporation
  - IDE. *See* Integrated Drive Electronics
  - Identification mechanisms, 310
  - IDS. *See* Intrusion Detection System
  - IEEE, 134
  - IETF. *See* Internet Engineering Task Force
  - IGMP. *See* Internet Group Management Protocol
  - IGRP. *See* Interior Gateway Routing Protocol
  - iHost (Oracle), 44
  - IIS. *See* Internet Information Server
  - ILEC. *See* Incumbent Local Exchange Carrier
  - ILMI. *See* Integrated Local Management Interface
  - Image-editing programs, 137
  - Implementation. *See* Infrastructure; Server level considerations
  - Implementation/business process design, 92
  - Implemented SLA, customer examination, 372
  - In-band virtualization, 270, 272–273
  - Incident response, 309, 317–319
  - Incremental backup, 164, 165
  - Incremental tape rotation method, 167
  - Incumbent Local Exchange Carrier (ILEC), 65, 104, 381
  - Independent Computing Architecture (ICA) protocol, 172–173
  - Independent provider, 94
  - Independent software vendor (ISV), 3, 6, 42–45, 96, 245. *See also* Third-party ISVs
    - companies, 43–45
    - information, 98
    - relationships, 87
    - sourcing, 364
  - Industry standardization efforts. *See* Content Delivery Network
  - Information
    - modeling. *See* Dynamic information modeling
    - securing, 37–38
  - Information Technology (IT), 2, 49–50. *See also* Corporate IT
    - budgets, 3
    - consulting, 4, 14, 23
    - managers, 268
    - outsourcing, 6–8. *See also* Platform IT outsourcing
    - professionals, 70
    - requirements, reduction. *See* Internal IT requirements
    - responsibilities, 86
    - skilled labor, shortage, 37
    - staff, 35
      - augmentation. *See* Pure IT staff augmentation
    - staffing, difficulties, 87
    - systems, 267
  - Information utilities, 15
  - Infrastructure, 50. *See also* Data-based infrastructure; High-speed packet-switched infrastructures; Internal infrastructure; Network; Nonhierarchical infrastructures; Shared infrastructure
    - attack, 334
    - capacity planning, 442–443, 474
    - connection, 442–443
    - deployment, 22
    - designing, 376, 415
    - directly attached storage, 263–264, 298
    - expansion, 442–443
    - growth, planning, 465–468, 476
    - hosting. *See* Application implementation, 22
    - load balancing, 204–205, 252–253
    - operations. *See* Application Service Provider
    - planning, best practices, 443
    - provider, 373
    - storage, 282–284
  - Infrastructure design
    - considerations, 417–421, 473
    - hierarchy, impact, 425–432, 473–474
    - introduction, 417–417
    - process, 418–420
  - Ingress filtering, 346–348
  - In-house
    - deployment, 88
    - system, 91
  - Inktomi, 28–29. *See also* Content Delivery Suite; Content Distributor; Content Manager; Object Store; Traffic Server
  - Innovation, 98
  - Input/Output (I/O), 46, 47
    - channel, 275
  - Installation delays, 70
  - Integrated Drive Electronics (IDE), 116. *See also* Enhanced Integrated Drive Electronics
  - Integrated Local Management Interface (ILMI), 516

- Integrated long-distance services, 70
- Integrators/implementers. *See* System
- Intel, 24, 28, 114, 127. *See also* Adaptive Load Balancing; NetStructure Cache Appliance
  - caching appliances, 190
- Intelligent network services, 247
- Intelligent storage, 79
- INTER\_AS metric attribute, 530
- Interconnection. *See* Operations Support System
  - challenges, 382
- Interconnects, 276
- Interexchange Carrier (IXC), 382
- Interface. *See* Hardware; Platform-independent interfaces/tools; Standard interfaces
  - broadcast queue, creation, 440–442
  - points, 207
- interface ethernet (command), 221
- Interior Gateway Routing Protocol (IGRP), 445, 446, 506. *See also* Enhanced Interior Gateway Routing Protocol
- Interior protocols, 444–448
  - choice, 448–449
- Interliant, 42
- Intermediate System-to-Intermediate System (IS-IS), 447, 451, 459, 466
  - usage, 497, 503, 519–520
- Internal infrastructure, 205
- Internal IT requirements, reduction, 100
- Internal servers, 132
- InterNAP, 237
  - data centers, 240
- InterNAP-owned data center, 237
- International Data Corporation (IDC), 3, 84, 97, 258
- International Organization for Standardization (ISO), 17, 142. *See also* Open System
  - Interconnection-International Organization for Standardization
- International Telecommunications Union (ITU), 388
  - M.3000 recommendation series, 389
- Internet
  - abuse, 305
  - access, 36
  - banking sites, 30
  - browsers, 6, 234
  - connection, 231
    - mission-critical nature, 69
  - connectivity, 22
  - content. *See* Mission-critical Internet-based content
  - slowdowns, 214
  - usage, 36
- Internet Assigned Numbers Authority (IANA), 344
- Internet Control Message Protocol (ICMP)
  - echo, 336
  - flooding, 350
  - packets, 334, 337, 341–342, 350
  - requests, 336, 337
- Internet Data Center Solutions (Nortel), 28
- Internet Engineering Task Force (IETF), 232, 444, 446
  - standards, 240
  - Working Groups, 242
- Internet Gateway products, 171
- Internet Group Management Protocol (IGMP), 461
  - hello-timer interval, 431
  - snooping, 461
- Internet Information Server (IIS), 138, 139
- Internet Packet eXchange (IPX), 137, 149, 264
- Internet Protocol (IP), 264, 276. *See also* Voice-over Internet Protocol
  - addresses, 129, 309, 312–313. *See also* Unregistered IP addresses; Virtual IP address
  - spoofing, 329
  - filtering, 211
  - fragmentation attack, 334, 336
  - IP-based traffic, 507
  - IP-Video Conferencing, 60
  - level, 198
  - multicast, 461–462
  - networks, 25, 446
  - packets, spoofing, 344
  - precedence, 531
  - protocol, 135
  - spoofing, prevention, 211
  - support, 208
  - technologies, 241
- Internet Security Systems (ISS), 331
- Internet Service Provider (ISP), 13, 41–42, 91–94, 188, 258
  - backbones, 239
  - definition, 5–6
  - evolution, 72–80, 104–105
  - industry, 60
  - investment, 190
  - market conditions, 61–66, 104
  - networks, 444
  - usage, 135
  - value proposition, 88–91
- Internet-based appliances, 135
- Internetwork
  - design considerations. *See* Frame relay



ease, 427  
 Internetwork Operating System (IOS)  
   software, 552  
   version 11.1, 337  
 Internetwork Packet eXchange (IPX), 433, 497, 498  
   networks, 446  
 Interoperability Forum (SIF), 390  
 Interoperability protocols, 203  
 Inter-Switch Link, 456, 458, 463, 499  
 Intranet. *See* Private intranets  
 Intranet-based ERP, 375  
 Intrusion Detection System (IDS), 318, 330–332  
   auditing, 319  
 Inventory, 379  
   control, 377  
 Investor demands, 71–72  
 I/O. *See* Input/Output  
 IOS. *See* Internetwork Operating System  
 IOS-based routers, 467  
 IP. *See* Internet Protocol  
 ip address (command), 220  
 ip audit notify log (command), 496  
 ip audit po local (command), 496  
 ip inspect name (global configuration command), 495  
 ip subnet-zero (command), 493  
 ip tcp intercept (command), 352  
 ip tcp intercept timeout (command), 354  
 IP-over-Photon architecture, 74  
 IPX. *See* Internet Packet eXchange; Internetwork Packet eXchange  
 IPX/SPX, 419  
 is (command), 221  
 is real (command), 225  
 is virtual (command), 225, 227

IS-IS. *See* Intermediate System-to-Intermediate System  
 ISL. *See* Inter-Switch Link  
 ISO. *See* International Organization for Standardization  
 ISP. *See* Internet Service Provider  
 ISS. *See* Internet Security Systems  
 ISV. *See* Independent software vendor  
 IT. *See* Information Technology  
 ITU. *See* International Telecommunications Union  
 IXC. *See* IntereXchange Carrier

## J

Java, 140, 141  
 JBOD. *See* Just-a-bunch-of-disks  
 J.D. Edwards, 41, 42  
 Joint Photographic Experts Group (JPEG), 19  
 JPEG. *See* Joint Photographic Experts Group  
 Juniper Networks, 26, 27, 480  
 JUNOS Internet software, 27  
 Just-a-bunch-of-disks (JBOD), 273

## K

K6, 114  
 keepalive (command), 501  
 Kerberos access control, 31  
 Key length, 314–315  
 KPMG, 41, 42, 95

## L

LAN. *See* Local Area Network  
 LAN emulation client (LEC), 541  
 Land attacks, avoidance, 211

LANE configuration server (LECS), 463  
 Large-scale  
   implementations, 83  
   internetworks, 431  
 Latency, 247, 440. *See also* Zero latency  
   assurances, 365  
   level, 266  
 Lawson Software, 39, 41  
   application, 45  
 Layer 1. *See* Physical layer  
 Layer 2. *See* Data  
 Layer 3. *See* Network  
 Layer 4. *See* Transport layer  
 Layer 5. *See* Session  
 Layer 6. *See* Presentation layer  
 Layer 7. *See* Application  
 Layer 8. *See* Political layer  
 Layer 9. *See* Religion layer  
 Layer 10. *See* Financial layer  
 Layers. *See* Lower layers; Pseudo layers; Upper layers  
 LCN. *See* Logical Connection  
 LDAP. *See* Lightweight Directory Access Protocol  
 Learning curve, 43  
 Least connections, 229  
 LEC. *See* LAN emulation client  
 LECS. *See* LANE configuration server  
 Legacy NEs, integration, 397  
 Legacy support, 259, 261–262  
 Legacy-switched voice technology, 25  
 Legal addresses, 347  
 Legator. *See* Cluster Server  
 Life cycle. *See* Cost of Ownership  
 Life-cycle cost, 9  
 LifeKeeper (NCR), 144  
 Lightweight Directory Access Protocol (LDAP), 140  
 Link aggregation, 123–129  
   standardization, 125  
 Links. *See* Weakest link

- Linux, 135, 136, 140, 144
    - operating system, 172
  - Load, 444
  - Load balancers, 215
  - Load balancing, 189, 198. *See also* Distributed load balancing; Hot-spare load balancing; Infra-structure; Localized load balancing; Parallel load balancing; Per-destination load balancing; Per-packet load balancing
    - appliance, 214
    - capabilities, 239, 468
    - product, 208, 213. *See also* Appliance-based load-balancing product
    - software, 29–30
    - solution
      - availability, 210–211
      - criteria, 209–211
      - dependability, 209
      - F5, usage, 212–215, 253
    - systems. *See* Software-only load-balancing systems
      - comparison, 205–212
    - usage. *See* Network
  - Local Area Network (LAN), 4, 275, 452. *See also* Corporate LAN; Emulated LANs; Virtual LANs
    - attachment, 118, 264
    - connection, 132
    - facilities, 15
    - group, 204
    - implementations, 101
    - interface, 464
    - LAN-based services, 73
    - LAN-to-LAN connections, 321
    - LAN-to-LAN VPNs, 320
    - linking, 418
    - resources, 277
    - segments, 327
    - specifications, 17
    - storage performance, 277
    - technology, 74
    - traffic, filtering, 337
  - Local loop. *See* Externally provided local loop
  - LocalDirector (Cisco), 215–228, 253
    - availability, 217
    - configuration samples, 220–223
    - security, 219
  - Localized load balancing, 204
  - Log file formats, 202
  - Logic unit number (LUN)
    - masking, 270–272
  - Logical Connection (LCN), 552
  - Logical network, overview, 481
  - Logical partitions, 289
  - Logical TMN model, 394–396
  - Logistics, 96
  - Long-distance services. *See* Integrated long-distance services
  - Long-term business model, 260
  - Long-term customer satisfaction, 100
  - Lost data. *See* Scheduled maintenance
  - Lotus Notes, 41, 171
  - Low voltage differential (LVD), 277
  - Lower layers, 21
  - Lower-end clustering solutions, 208
  - Low-security systems, 307
  - Lucent Technologies, 26, 27
  - LUN. *See* Logic unit number
  - LVD. *See* Low voltage differential
- ## M
- M.3000 recommendation series. *See* International Telecommunications Union
  - MAC. *See* Media Access Control
  - MacO/S, 172
  - Macromedia, 30
  - Maintenance. *See* Scheduled maintenance; Service level agreement
    - issues, 174–178, 183
  - Malicious programs, 169
  - MAN. *See* Metropolitan Area Network
  - Managed security, 70
  - Managed service provider (MSP), 91
  - Managed VPNs, 76
  - Management. *See* Billing; Configuration; Network; Performance; Service management
    - GUI, 202
    - systems. *See* Application Service Provider
    - tools deployment, industry examples. *See* Application Service Provider
  - Management Information Base (MIB), 202, 390, 396
    - objects, 509
    - representations, 396
  - Management/monitoring, 363
    - introduction, 364
  - Manager roles, 393
  - Margin, increase, 71
  - Marimba Incorporated, 400
  - Market
    - change, 386–387, 411
    - conditions. *See* Internet Service Provider
    - factors, 84–86
    - opportunities, 42
    - penetration, 68, 98
  - Marketing, 49. *See also* Copy
  - Market-leading expertise, 99
  - Market-ready databases, 110
  - Masking. *See* Logic unit number
  - Mass storage, 115–118
  - Mass-storage products, 260
  - Master/slave architecture, 341
  - match (command), 508
  - Match-all, Match-any (contrast), 490

- Maximum burst size (MBS), 542
  - Maximum transmission unit (MTU), 501
  - Max-incomplete high, 355
  - Max-incomplete low, 355
  - MBONE. *See* Multicast backbone
  - MBS. *See* Maximum burst size
  - McAfee, 170
  - McData, 275
  - MCF. *See* Message communication function
  - M-Class systems, 24
  - MC/ServiceGuard. *See* Multi-Computer/ServiceGuard
  - MD. *See* Mediation device
  - MD5. *See* Message Digest 5
  - MDS. *See* Multicast distributed switching
  - Mean time between failures (MTBF), 374
  - Mean time to repair (MTTR), 374
  - Media Access Control (MAC) addresses, 126, 127, 462, 464
  - Mediation device (MD), 391
  - Mercury Interactive, 101
  - Mergers, reason, 45
  - Message communication function (MCF), 392–393
  - Message Digest 5 (MD5), 487
  - Messaging, 375
  - Metamor Worldwide, 41
  - Metric support, 444
  - Metrics, 402, 449. *See also* Baseline metrics testing, measurement, 370
  - Metropolitan Area Network (MAN), 66, 74
  - MGX router. *See* Cisco Systems
  - MIB. *See* Management Information Database
  - Middleware, 133, 138, 142–143
  - Midstream, 243
  - Mirroring. *See* Asynchronous mirroring; Remote mirroring; Synchronous mirroring
  - Mission-critical applications, 47, 87, 136, 230, 376
    - data, 289
    - information, 102
    - integrity, 38
    - Internet-based content, 234
    - services, 70
  - Mixed vendor support, 259, 260
  - MLS. *See* Multilayer switching
  - Modal bandwidth, 120
  - Modems. *See* Cable
  - Monitoring, 79. *See also* Management/monitoring; System implementation guidelines, 372–373 solution, 373
  - Moving Pictures Expert Group (MPEG), 19, 243
  - MP3, 73
  - MPEG. *See* Moving Pictures Expert Group
  - mpls traffic-eng tunnels (command), 496
  - MSFC. *See* Multi-layer Switch Feature Card
  - MSP. *See* Managed service provider
  - MTBF. *See* Mean time between failures
  - MTTR. *See* Mean time to repair
  - MTU. *See* Maximum transmission unit
  - Multicast. *See* Internet Protocol control traffic, optimization, 427, 433 policy, 461
  - Multicast backbone (MBONE), 462
  - Multicast distributed switching (MDS), 512
  - Multicasting capabilities, 70
  - Multi-Computer/ServiceGuard (MC/ServiceGuard), 47
  - Multihomed servers, contrast. *See* Dual servers
  - Multihost arrays, 270–271
  - Multilayer model
    - bridging, 468
    - security, 468
  - Multi-layer Switch Feature Card (MSFC), 467
  - Multilayer switching (MLS), 496, 498
  - Multilevel hierarchies, 199–200
  - Multimedia
    - content, 110, 231
    - server, 137
    - services, 79, 460–464, 476
  - Multimode fiber, 120
  - Multinetwork CDN, 239
  - Multiple real servers, 219–220, 226–228
  - Multiple virtual servers, 223–228
  - Multiplexing schemes, 387
  - Multiprocessing. *See* Symmetric multiprocessing
  - Murphy's Law, 143
- ## N
- NAI. *See* Network Associates name (command), 221, 224
  - NAP. *See* Network Access Point
  - NAS. *See* Network Attached Storage
  - NASA, 313
  - NAT. *See* Network Address Translation
  - NaviSite, 42, 244
  - N-Class systems, 24
  - NCR. *See* LifeKeeper
  - NE. *See* Network element neighbor remote-as (command), 505
  - NerveCenter (Veritas), 178
  - NET. *See* Network entity title NetBEUI protocols, 149

- NetBIOS, 467, 468
  - protocol, 149
- NetBIOS Message Block Daemon (NMBD), 158
- Net-centric
  - environment, 4
  - software, 4
- Netegrity, 399
- Net-hosted application, 92
- Netscape Enterprise Server, 140
- Netscape Fast Track Server, 138
- NetStructure Cache Appliance (Intel), 28
- NetWare (Novell), 137, 162
- Network. *See* Alcatel networks;
  - Bastion network; Dirty network; Extreme networks; F5 networks; Foundry networks; Internet Protocol; Juniper networks; Nortel networks; Storage area network
- adapters, 118–133
- administration resources, 9, 12
- appliances, 208–209, 213
- availability, 367, 369
  - contrast. *See* Network-based availability
- bandwidth availability, 247
- block, 346
- closeness, implication, 205
- connectivity, 283. *See also* Access
- delays, 49. *See also* Average network delay
- development, 15
- enhancement/extension,
  - load balancing usage, 211
- equipment, 25–28
  - auditing, 319
  - basics, 424–425
- infrastructure, 197, 295
- intelligence, 238
- investment, 212
- layer (Layer 3), 17, 18
  - switching, 466–467
- layers, development/deployment speed, 230
- management, 380–381
- modification, 420
- monitoring, 76
- NAS location, 266–267
- outage, 131, 371
- performance, 267
- provider, 373
- reach, scaling, 295
- regional topologies. *See* Frame relay
- sabotage, 305
- scalability, 465–467
- service considerations, 147–159, 182
- service providers, 239. *See also* Satellite-based network service providers
  - requirements, CDN service providers/CDN component product makers interaction, 240
- services, 453–454, 475. *See also* Data center; Intelligent network services; Smart network services
- SLAs, 374
- sniffing, 311–312
- storage, 147–159
- streamlining, 420
- support, 380–381
- testing, 481–485
- transport. *See* Transmission Control Protocol/Internet Protocol
  - unavailability, 213, 214
  - usage, rules, 424
- Network Access Point (NAP), 193, 199, 202, 248
- Network Address Translation (NAT), 211, 228, 450
  - usage, 218
- Network Appliance, 127
- Network Associates (NAI), 170, 319
- Network Attached Storage (NAS), 2, 148, 258, 264–267
  - devices, 159, 265, 266
  - location. *See* Network SAN, contrast, 274–275
  - servers, 265
  - solutions, 298–299
  - usage, 460
- Network element (NE), 389, 394–395, 434
  - integration. *See* Legacy NEs
- Network Engines, 243
- Network entity title (NET), 503
- Network File System (NFS) protocol, 148–149
- Network Interface Card (NIC), 118, 123–125. *See also* Non-operational NIC
- Network Management Forum (NMF), 390
- Network Management Layer (NML), 394
- Network News Transfer Protocol (NNTP), 195, 203
  - news cache, 203
  - server, 137
- Network Operating System (NOS), 133, 260, 272, 400–401. *See also* Unix-based NOS
- Network Operations Center (NOC), 2, 240, 381
- Network Specialists. *See* Double-Take
- Network Time Protocol (NTP), 471, 552
- Network Working Group, 344
- Network-based availability, site-based availability (contrast), 369
- Network-based solutions, 111
- Networking. *See* Application-aware networking
- Never-cache, 202
- New-model service provider, 247

- News feeds, dissemination, 460
- News organizations, 234
- Next-generation content-based services, 247–248
- Next-hop device, 335
- NFS. *See* Network File System
- NIC. *See* Network Interface Card
- NMBD. *See* NetBIOS Message Block Daemon
- NMF. *See* Network Management Forum
- NML. *See* Network Management Layer
- NNTP. *See* Network News Transfer Protocol
- no bgp default ipv4-unicast (command), 505
- no cdp enable (command), 513
- no ip directed broadcast (command), 511
- no shutdown (command), 513
- NOC. *See* Network Operations Center
- Nodes, 269, 311. *See also* Routing; Routing Information Protocol
- Nonbroadcast environments, 429
- Nonfunctional system, 176
- Nonhierarchical infrastructures, 426
- Nonoperational NIC, 124
- Nontransparent caching, 203
- Normal burst size, 349
- Nortel Networks, 28, 374, 480. *See also* Internet Data Center Solutions
- NorthPoint Communications, 65
- NOS. *See* Network Operating System
- NovaStor, 162
- Novell, 208. *See also* NetWare; Web and Application Services
- NTP. *See* Network Time Protocol
- O**
- Object Management Group (OMG), 382
- Object Store (Inktomi), 28
- Objects
  - delivering. *See* Requested object
  - queue processing, 196
  - requesting. *See* Cached objects
- OC-3, 458
- OC-12 trunks, 458
- OC-48 connection, 66
- OC-192, 458
  - transport link, 66
- ODBC. *See* Open Database Connectivity
- OEM. *See* Original Equipment Manufacturer
- Off-campus programs, 16
- Offerings, determination, 96–99
- OFO. *See* Open File Options
- OMG. *See* Object Management Group
- One-minute high, 354–355
- One-minute low, 354
- One-minute threshold, 354
- Onyx Software Corporation, 400
- Open Database Connectivity (ODBC), 140
- Open File Options (OFO), 163
- Open Proxy Extension Services (OPES), 242
- Open Shortest Path First (OSPF), 18, 445–447, 451, 464, 519
  - routing traffic, 459
- Open System Interconnection (OSI), 19
  - model, 20, 21
- Open System Interconnection-International Organization for Standardization (OSI-ISO), seven layer model, 16–22, 54
- Open systems, 268
- Openview (HP), 178
- OpenVMS, 144
- Operating system (O/S // OS), 133, 145. *See also* Heterogeneous operating systems; Network operating system; Unix dependencies, 207 management, 95
- Operation Support System (OSS) model, 377–386, 395, 410
  - basics, 378–381
  - design/support, 377
  - efficiencies, 383
  - flexibility, 383
  - functionality, 380
  - interconnection, 381–383
  - upgrading, 382–383
- Operational capacity, 195–196
- Operations systems (OS), 391
- OPES. *See* Open Proxy Extension Services
- OR (operator), 293
- Oracle, 42, 142. *See also* Baan Oracle PeopleSoft SAP; iHost
  - application, 45
  - Business OnLine, 41, 43
  - platforms, 44
- Ordering, 379
  - functions, 378
- Original Equipment Manufacturer (OEM), 46, 96
- OS. *See* Operating system; Operations systems
- O/S. *See* Operating system
- OSI. *See* Open System Interconnection
- OSPF. *See* Open Shortest Path First
- OSS. *See* Operation Support System
- Outboard security, 261
- Out-of-order packets, 450
- Out-of-the-box utilities, 383

- Outsourcing. *See* Application outsourcing; Business process outsourcing; Information Technology; Platform IT outsourcing
  - effect, 364–368, 408
  - providers, 13, 15
- Oversubscription. *See* Server
- P**
- Packet
  - balancing, 207
  - capturing, 357
  - direction, 349
  - filtering, 322, 324–327
  - forwarding, 499
  - replication, 435
- packet inter-network groper (ping), 336, 370
  - attacks, avoidance, 210
- Packet-based technologies, 463
- Packet-filtering rules, applying, 326
- Packet-service infrastructure, 419
- Packet-switching data network (PSDN), 418, 427, 434, 436
  - devices, 418
  - environment, 440
  - infrastructure, 417
  - networks, 433
- Packrats, 282
- Parallel load balancing, 208
- Parent tape rotation scheme, 165–166
- Parity information, 294
- Partially meshed topologies, 428, 430–431, 438–439
- Partnered deployment, 94
- Pascal, Blaise, 111
- Password-based security mechanisms, 311
- Passwords, 310
  - exchange, 311
  - sniffing, 311
- PAT. *See* Port Address Translation
- Pattern matching, 228
- PBX. *See* Private Branch eXchange
- PeopleSoft, 44, 399. *See also* Baan Oracle PeopleSoft SAP
  - application, 45
- PeopleSoft Financial Management, 39, 41, 42
- Per-destination load balancing, 449–450
- Performance, 100, 101, 390. *See also* Cache; Fault Configuration Accounting Performance Security; Service provider
  - contrast. *See* Price decay, 207
  - enhancement technologies, 187
    - introduction, 188–189
  - increase, 69
  - issues, 417. *See also* Application Service Provider; Wide Area Network
  - management, 281, 377
  - monitoring/reporting, 245
  - unpredictability, 188
- Performance-reporting capabilities, 374
- Perimeter firewalls, 321–327
- PERL, 140
- Permanent virtual connection (PVC), 433, 442, 517, 520–521, 542
  - addition/reconfiguration, 370
  - availability, 367, 369–370
  - delay. *See* Average PVC delay
- Perot, Ross, 23
- Per-packet load balancing, 449–450
- Personal applications, 97
- Personal Identification Number (PIN), 310
- Per-use agreement, 419
- Per-user charge, 101
- Physical attacks, 334, 338
- Physical equipment
  - environment, 422
  - power, 422
  - space, 421–424
  - weight, 422
- Physical layer (Layer 1), 17
- Physical media, 440
- PIM. *See* Protocol independent multicast
- PIN. *See* Personal Identification Number
- ping. *See* packet inter-network groper
- Pin-in-cache, 202
- Plain Old Telephone System (POTS), 386
- Planned upgrades, 174–176
- Planning, 365
- Platform IT outsourcing, 7–8, 15
- Platform support, 272
- Platform-independent interfaces/tools, 397
- PNNI, 459
- Point-of-failure issues, 438
- Point-of-origin server, 232, 236, 244
- Points of presence (POPs), 33, 193, 199, 201, 248, 482
- Point-to-point
  - interconnections, 428
  - networks, 464
- POISSON model, 377
- police (command), 512
- policy-map (global configuration command), 492
- PolicyMaps, 490–493
- Political layer (Layer 8), 21
- POPs. *See* Points of presence
- Port Address Translation (PAT), 229, 450
- Portable Operating System Interface (POSIX), 134
  - standard, 135
- Portal Software Incorporated, 400

- Ports. *See* Real server; Virtual server
- POS1/0, 507
- POSIX. *See* Portable Operating System Interface
- POTS. *See* Plain Old Telephone System
- PowerPC, 135
- predictor (command), 226
- Presentation layer (Layer 6), 17, 19
- Prevention
  - tactics, 360
  - techniques, 343–355
- Price, performance contrast, 259, 262–263
- Priceline.com, 64
- Pricing, 101–102. *See also* Transaction-based pricing; Usage-based pricing
  - billing, 401–406, 413–414
  - models, 401–406, 413–414
  - pressure, 67, 89
  - structure, 100
- Priority-level maps, 456
- Privacy policy, 308–309
- Private Branch eXchange (PBX), 2. *See also* Virtual Private Branch eXchange
- Private intranets, 138
- Private investors, 64
- Procurement, 96
- Productivity, improvement, 88
- Productivity-enhancing packages, 85
  - leveraging, 88
- Professional consulting, 13, 14
- Profitability, 71, 104
- Progress Software Corporation, 400
- Project-based service providers, 13–15
  - integration/implementation, 14–15
- ProLiant (Compaq), 48
- prompt (command), 487
- Proprietary information, 305
- Protocol. *See* External protocols; Interior protocols; Routing
  - dependency, 425
  - needs. *See* Bridged protocol needs
  - planning, concerns, 444–450, 474–475
  - routing, 433–434
  - usage, 276
- Protocol independent multi-cast (PIM), 461–462
- Prototyping. *See* Automated prototyping
- Providers. *See* Education/training providers; Internet Service Provider; Outsourcing; Project-based service providers; Service provider; Staff augmentation providers
  - capabilities, 235
- Provisioning, 95. *See also* Data; Engineering/provisioning
- Proxy. *See* Forward proxy; Reverse proxy
  - caches, 242
- PSDN. *See* Packet-switching data network
- Pseudo layers, 21
- PSINet, 240
- PSTN. *See* Public Switched Telephone Network
- Public Switched Telephone Network (PSTN), 61–62
- Public-key algorithms, 316
- Public-key cryptography, 316
- Pure ASP, 6, 39–40
- Pure IT staff augmentation, 16
- Pure-play ASP, 5
- PVC. *See* Permanent virtual connection
- Q**
- Q adapter (QA), 391
- Q adoption capability/compatibility, 397
- QA. *See* Q adapter
- QAF, 393
- Qlogic, 275
- QoS. *See* Quality of Service
- Quality of Service (QoS), 49, 173, 188, 210, 215, 387–388, 411, 841
  - abilities, enhancement, 233
  - benefits, 199
  - implementation, 195
  - improvement, 189, 191, 230
  - issues, 188
  - levels, 241
  - needs, 385
  - options, 13
  - QoS-based availability, 206
  - QoS-type granularity, 248
  - usage, 248, 266
- Queues, processing. *See* Objects
- Queuing techniques, 464
- QuickTime, 19
- Qwest Communications International, 41
- Qwest Cyber.Solutions, 41, 95
- R**
- Radio stations, 234
- RAID. *See* Redundant Array of Inexpensive Disks
- RAM. *See* Random Access Memory
- Random Access Memory (RAM), 29, 46, 114–115, 189–190
- Random early detection (RED), 457, 531. *See also* Weighted Random Early Detection
- Rate limiting, 348–352
  - command, 349
- rate-limit (command), 492
- RateXchange Trading System, 63, 64
- RBOCs. *See* Regional Bell Operating Companies
- RCDD. *See* Registered communication distribution designer

- RD. *See* Route distinguisher
- RDIST. *See* Remote File Distribution
- real (command), 221, 224, 227
- Real server, 223–226. *See also* Multiple real servers ports, 223–224
- Real-time site replication, 161
- Real-to-Virtual-to-Real (RVR), 218
- RED. *See* Random early detection
- Redirection service, 232
- Reduced Instruction Set Computers (RISC), 112, 467  
microprocessors, 113  
processors, 24
- Redundancy, 262. *See also* Server  
amount, 143  
internal requirements, 216  
level, 429, 470
- Redundant Array of Inexpensive Disks (RAID), 258, 276, 291–295  
adapter, 292  
architecture, 280  
controller, 270, 280  
mirroring solution, 294  
RAID-0, 292, 293  
RAID-1, 292, 293  
RAID-2, 292, 293  
RAID-3, 292, 293  
RAID-4, 292, 294  
RAID-5, 292, 294  
RAID-6, 292, 294  
RAID-7, 292  
RAID-10, 292, 294–295  
RAID-53, 292, 295  
usage, 273
- Regional Bell Operating Companies (RBOCs), 67, 381–383  
customers, 384  
interfaces, 384
- Regional topologies, 433. *See also* Frame relay
- Registered communication distribution designer (RCDD), 423
- Reliability, increase, 69
- Religion layer (Layer 9), 22
- Remote backup, 278
- Remote File Distribution (RDIST), 245
- Remote links, definition, 419
- Remote mirroring, 290–291
- Remote Operations Center (ROC), 2
- Remote Procedure Call (RPC), 323
- Remote troubleshooting, 375
- Remote users, 418
- Replication. *See* Database; Real-time site replication; Traffic; World Wide Web
- Reply traffic. *See* Server-to-client reply traffic
- Reputation, 100
- Request For Comment (RFC), address spaces filtering, 344–346
- Requested object, delivering, 196
- Requests, responsiveness. *See* Clients
- Residential connections, 69
- Resolution, contextualization, 242
- Resonate, 30, 32–33  
Commander, 33  
Global Dispatch, 32
- Respondents, 305
- Response time, 367, 369, 371  
guarantees, 375
- Response-time delays, 232
- Retailers, 234
- Retransmissions, reduction, 424
- Return on investment (ROI), 8, 23, 49, 101, 111, 379
- Revenue  
growth, 71  
streams. *See* High-margin revenue streams
- Revenue-generating services, 76
- Revenue-sharing model, 98
- Reverse proxy, 198–199, 203
- RFC. *See* Request For Comment
- Rhythms NetConnections, 65
- RIP. *See* Routing Information Protocol
- RISC. *See* Reduced Instruction Set Computers
- Risk  
assessment, 309, 319–320  
transfer, 37
- Ritchie, Dennis, 133
- ROC. *See* Remote Operations Center
- ROI. *See* Return on investment
- Rollback, 245
- Root-level access, 468
- Round robin, 229. *See also* Weighted round robin rotation. *See* Domain Name System
- Round-robin DNS, 212
- Route distinguisher (RD), 494
- Route Switch Module (RSM), 467
- Routers, 206, 208, 267, 276, 426. *See also* Customers; Digital Subscriber Line connection, 429  
interfaces, increased costs, 436  
multicast performance, 461  
Route selection, 449–450
- Routing, 464  
algorithm, 452  
nodes, 432  
protocols, 444–448  
tables, 466
- Routing Information Protocol (RIP), 18, 439, 445–446, 449, 459  
advertisements, 506  
RIP-based nodes, 431  
usage, 503



version 2 (RIPv2), 447, 503  
 Routing Table Maintenance Protocol (RTMP) (AppleTalk), 459, 466  
 RPC. *See* Remote Procedure Call  
 RSM. *See* Route Switch Module  
 RTMP. *See* Routing Table Maintenance Protocol  
 RVR. *See* Real-to-Virtual-to-Real

## S

Sales, 49  
 Sales training companies, 234  
 Samba, 150, 158  
 SAN. *See* Storage Area Network; Storage area networking  
 SAP, 41, 42. *See also* Baan Oracle PeopleSoft SAP  
 Satellite-based network service providers, 240  
 Scalability, 100, 209, 258, 464.  
*See also* Cache; Hierarchical internetworks; Network agility, 88 considerations, 458–459, 476 impact. *See* Business  
 Scaling bandwidth, 458 considerations 844–845  
 Scanning. *See* Viruses  
 Scheduled maintenance, 368 lost data, 370  
 Scheduling, 457, 458  
 SCM. *See* Supply Chain Management  
 SCO Unix operating system, 172  
 SCR. *See* Sustained cell rate  
 SCSI. *See* Small Computer System Interface  
 Secure socket shell (SSH), 211.  
*See also* Built-in SSH  
 Secure Sockets Layer (SSL), 229, 313 browser, 320

encryption, 203, 211  
 session ID tracking, 207, 210  
 session ID-based switching, 228  
 transaction, 210  
 Security, 69, 102, 259–261, 390, 421. *See also* Fault Configuration Accounting Performance Security; Host-based security; Managed security; Multilayer model; Outboard security  
 auditing, 309, 319–310, 330  
 components, 309–320, 359  
 enhancement, 295  
 mechanisms, 335  
 offerings, 69  
 policy, 306–309, 359 development, 306–308  
 services, 79  
 system  
 customer confidence, building, 306  
 provisioning. *See* Application Service Provider  
 technology, 320–332, 359–350  
 Segue, 33  
 Server, 24. *See also* Domain; Dual-homed servers; Internal servers; Multiple real servers; Multiple virtual servers; Real server; Virtual server  
 availability, 217  
 communities, 245  
 connections, management, 218–219  
 contrast. *See* Dual server  
 failure, 213, 214  
 farm, 111, 470  
 scaling, 215–228  
 hardware, 111–133  
 management, 95  
 manufacturers, 115  
 oversubscription, 145–147  
 ports. *See* Real server; Virtual server

redundancy, 143–145  
 selection, ability, 228  
 software, 139  
 traffic, growth, 215  
 undersubscription, 145–147  
 uptime, 231  
 Server Advertisement Protocol (SAP), 459  
 Server level considerations, 109  
 implementation, 111–133, 180–181  
 introduction, 110–111  
 Server Message Block (SMB), 148  
 protocol, 149–159  
 Server Provider Edition (SPE), 32  
 Server-based services, 2  
 Server-free backup, 273  
 ServerIron (Foundry Network), 228–230, 254  
 Server-less backup solution, 161  
 Server-to-client reply traffic, 228  
 Server-to-server storage, 264  
 Server-to-storage access, 263  
 Server-to-storage device connection, 277  
 Service level agreement (SLA), 12, 38, 48–49, 100, 307, 365–366. *See also* Application; Network; System level SLAs  
 advantage, 67  
 components. *See* Frame relay customer examination. *See* Implemented SLA  
 excluded items, 268  
 guarantees, 365–366  
 maintenance, 372, 373  
 needs, 385  
 Service levels  
 consideration. *See* Service provider verification, 372  
 Service licensing agreement (SLA), 23  
 Service management, 379–380

- Service Management Layer (SML), 394
- service password-encryption (command), 486
- Service provider. *See* Application Service Provider; Content Delivery Network; Internet Service Provider; Network; Project-based service providers; Satellite-based network service providers
- acceptable performance, 376–377
- bonus, 377
- business requirements, 67–72, 104
- finances, 82
- future, 80–82, 105
- model, 68–72
- response, 376–377, 410
- service levels, consideration, 368–371, 408–409
- ServiceNet, 40
- Services. *See* Bundled services; Hosting; Multimedia services; Revenue-generating services; Security; Value-added services
- activation. *See* Data addition, 70
- breadth/depth, 38
- components, 50
- considerations. *See* Network deployment, 75–80
- fulfillment, 378
- quality, 38
- requirements, 49, 56
- support. *See* Data upgrades, 70
- Session
  - layer (Layer 5), 17, 19
  - persistence, 228
- SFA, 85
- SGML, 140
- Shared applications, use increase, 36
- Shared infrastructure, 92
- Shared network, 35
- Shared resources, 289
- Shared services, resolution, 238
- Shared-device model, 144
- Shared-nothing model, 144–145
- Short-wave lasers, 120
- show bind (command), 225, 227–228
- show configuration (command), 222–223
- show tcp intercept connections (command), 355
- show tcp intercept statistics (command), 355
- shutdown interface (command), 221
- Siebel, 41, 43
- Siebel Systems, 399
- SIF. *See* Interoperability Forum
- Signaling technologies. *See* Content
- Silicon Graphics, 127
- Simple Mail Transfer Protocol (SMTP), 328, 332
- Simple Network Management Protocol (SNMP), 280, 366, 380, 509
  - emergencies, 536
  - management, 202
  - packet size, 551
  - SNMP-based management, 228
- Single-mode fiber, 120
- Single-source solution provider, 39
- Site-based availability, contrast. *See* Network-based availability
- SiteMinder, 399
- Sites
  - availability, 369
  - blacklist filtering, 202
  - considerations, 421–425, 473
  - downtime, reasons, 213–214
  - mirroring, 278
  - replication. *See* Real-time site replication
- SLA. *See* Service level agreement; Service licensing agreement
- Small Computer System Interface (SCSI), 116, 261, 263, 276. *See also* Fast SCSI; Wide SCSI
  - comparison. *See* Fiber channel
  - drives. *See* Ultra-2 SCSI drives
  - implementations, 279
  - interfaces, 276
  - limitations, 277
  - SCSI over IP, 275
  - SCSI-3, 277
  - topologies, 280
  - usage, 285
- Small- to medium-sized enterprises (SMEs), 3, 42
  - market, 44
- Smart network services, 248
- SmartNet Package, 217
- SMB. *See* Server Message Block
- SMBD, 158
- SMDS. *See* Switched Multi-megabit Data Service
- SMEs. *See* Small to medium-sized enterprises
- SME. *See* System management function
- SML. *See* Service Management Layer
- SMP. *See* Symmetric multiprocessing; Symmetrical multiprocessing
- SMP-based systems, 47
- SMS. *See* Systems Management Server
- SMTP. *See* Simple Mail Transfer Protocol
- Smurf attack, 334, 336–338
- SNA. *See* Systems Network Architecture
- Sniffing. *See* Network
- SNMP. *See* Simple Network Management Protocol
- Social engineering, 327

- Software, 29–33. *See also* Load balancing  
 adaptability, 38  
 applications, 101, 142  
 downloads, 460  
 failure, 213, 214–215  
 packages, 260  
   acquisition. *See* Application; System products, 208  
 selection, 162–168  
 solutions. *See* Application Service Provider  
 upgrades, cost. *See* Application; System vendors, 95. *See also* Independent software vendor
- Software-based connection, 494
- Software-only  
 load-balancing systems, 205–206  
 solutions, 205–206
- Solaris, 140, 144  
 operating system, 172, 357  
 servers, 31  
 versions, 134
- Solomon–Wolff Associates, 63
- SONET. *See* Synchronous Optical Network
- Spanning Tree Protocol (STP), 128–130, 463  
 broadcast domain, 452  
 loops, 459
- Spanning trees, issues, 131
- SPARC, 113, 135
- Sparse mode, 461
- SPE. *See* Server Provider Edition
- Spectra, 141
- Speed. *See* Cache
- Splitter, 386
- Spoofing, 312–313
- Sprint, 240
- SQL. *See* Structured Query Language
- SSH. *See* Secure socket shell
- SSL. *See* Secure Sockets Layer
- Stacheldraht, 340, 342–343
- Staff augmentation providers, 13, 15–16
- Staffing, 16
- Standard interfaces, 393–394
- Star topologies, 428–429, 437
- Stateful inspection, 322–324
- Static HTML, 193
- Static routes, 434
- Static routing, 503
- Stock quotes, dissemination, 460
- Storage. *See* Data; Infrastructure; Intelligent storage; Mass storage; Network devices, 276  
 concerns, 259  
 direct attachment. *See* Infrastructure; Network attached storage  
 domain servers, 270, 273–274  
 performance. *See* Local Area Network  
 services, 97  
 systems, 268
- Storage Area Network (SAN), 2, 148, 258, 267–281, 299  
 administration, improvement, 269  
 attached storage, 296  
 availability, improvement, 269  
 benefits, 268–269  
 business flexibility, increase, 269  
 configurations, 281  
 contrast. *See* Network attached storage  
 devices, 266  
 environment, 262  
 fabric, 270  
 management, 280–281  
 need, 267–268  
 optimized storage, 296  
 optimized systems, 296  
 solution (IBM), 295–296  
 solutions. *See* Vendors  
 storage devices, 284  
 usage, 460  
 virtualization, 270–274, 288
- Storage solutions, 257. *See also* Advanced storage solutions  
 introduction, 258  
 selection criteria, 259–263, 298  
 upfront concerns, 259–263, 298
- StorageTek, 117, 275
- Store. *See* Advanced store; Basic store
- StorEdge products, 25
- STP. *See* Spanning Tree Protocol
- Strategic management  
 consulting, 14
- Streaming broadcasts, 232
- Streaming media, 243–244  
 appliances, 243
- Striping, 291
- Structured Query Language (SQL), 142
- Success factors, 94–102, 105–106
- Sun Microsystems, 24–25, 47, 117, 127, 134, 148, 374  
 operating systems, 357
- Supply Chain Management (SCM), 37, 41, 43, 96, 377
- Support, 100–101. *See also* Network capabilities, 100  
 improvement, 70  
 issues, 174–178, 183  
 quality, 38
- Sustained cell rate (SCR), 543
- SVC. *See* Switched virtual circuit
- SwitchBack technology, 228
- Switched links. *See* High-speed switched links
- Switched Multi-megabit Data Service (SMDS), 376
- Switched virtual circuit (SVC), 442, 518, 521
- Switches, 206, 207, 267, 276, 425. *See also* Frame relay

- performance, 461
  - Switching. *See* Data
  - Switch-to-switch deployment, 365
  - Sybase application, 45
  - Symantec, 170
  - Symmetric algorithms, 316
  - Symmetric keys, implementations, 316
  - Symmetric multiprocessing (SMP), 114, 202–203
  - Symmetrical multiprocessing (SMP), 24
    - configuration, 25
  - SYN. *See* Synchronization
  - Synchronization (SYN). *See* Data
    - attacks, 334–336, 351–352
    - floods, 335
    - prevention, 211
    - requests, 335
  - Synchronizing methods, 147
  - Synchronous mirroring, 290–291
  - Synchronous Optical Network (SONET), 390
  - Syslog, 356
  - System
    - availability, 259, 262
    - development, 14
    - engineering, 377
    - general-management
      - functionality, 390
    - integrators/implementers, 40–41
    - management. *See* Operating system
    - monitoring, 95, 177–178
    - penetration, 305
    - software, 133–137
      - package acquisition, 9, 10
      - upgrades, cost, 9, 11
    - uptime, 365
      - amount, 46–47
  - System level SLAs, 374
  - System management function (SMF), 397
  - Systems Management Server (SMS), 374
  - Systems Network Architecture (SNA), 467
    - bridging, 419
- ## T
- T1 line, 265, 370
  - T3 line, 350
  - Tagged Image File Format (TIFF), 19
  - Tag-switching packets, 501
  - Tape library, 275, 276
  - Tape rotation
    - method. *See* Incremental tape rotation method
    - scheduling, 163–168
    - scheme. *See* Child tape rotation scheme; Grandparent tape rotation scheme; Parent tape rotation scheme
  - Tcl. *See* Tool Command Language
  - TCO. *See* Total cost of ownership
  - TCP. *See* Transmission Control Protocol
  - TCP/IP. *See* Transmission Control Protocol/Internet Protocol
  - TDM. *See* Time Division Multiplexing; Time division multiplexing
  - Teardrop attacks, avoidance, 211
  - Technical expertise, obtaining, 37
  - Technologies. *See* Security
    - access, 87
    - adjusting, 37
    - enabling, 35–36, 64–65
  - Telecommunication
    - companies, 41–42
    - providers, 25
  - Telecommunications
    - Management Network (TMN)
      - automation tools, 396–397
      - model, 384. *See also* Logical TMN model
        - building blocks, 391–392
      - model, OSI interaction, 392–394
      - outline, 388–391
      - standards, 389–391
        - conformance, 396
    - Telecomputing ASA, 40
    - TeleManagement forum, 384
    - Television stations, 234
    - Terminal data centers, 454
    - TFN. *See* Tribal Flood Network
    - TFN2K. *See* Tribal Flood Networks 2000
    - TFTP connections, 494
    - Thin client, 171–173
      - solutions, 183
      - computing, 36
      - mechanism, 402
      - technology, 454
    - Third-party application, 328
    - Third-party clearinghouse services, providing ability, 238
    - Third-party ISVs, 39
    - Third-party organization, 313
    - Third-party partnerships, 39
    - Third-party providers, 44
    - Third-party vendor products, 374
    - Thompson, Ken, 133
    - Three-bit field, 456
    - Throughput. *See* Effective throughput; Wire-speed throughput
      - handling, 197
      - rates. *See* Data
    - TIFF. *See* Tagged Image File Format
    - Time Division Multiplexing (TDM), 74
    - Time division multiplexing (TDM), 442
    - Time to resolution or repair, 368, 369, 371
    - Time-to-live (TTL), 466
    - Time-to-market

- improvement, 87–88
  - issues, 95
  - TL1. *See* Transaction Language 1
  - TLV. *See* Type length value
  - TMN. *See*
    - Telecommunications Management Network
  - Tool Command Language (Tcl), 140
  - Topologies, 451–453. *See also*
    - End-to-end topology;
    - Frame relay; Fully meshed topologies;
    - Partially meshed topologies;
    - Regional topologies;
    - Star topologies
  - design, 417
  - types, 428–432
  - Torvald, Linus, 135
  - ToS. *See* Type of Service
  - Total cost of ownership (TCO), 34, 35
  - Total service availability (TSA), 373
  - Towbin. *See* C.E. Unterberg, Towbin
  - TP. *See* Transaction process
  - Traffic. *See* Server-to-client reply traffic
    - classification, 456–457
    - detection, 455
    - engineering router identifier, 529
    - explicitly allowing, 326
    - explicitly denying, 327
    - explosion. *See* Data flow, 326
    - optimization. *See* Broadcast; Multicast
    - patterns, 372
    - prioritization, 210–211
    - requests, 29
    - requirements, 211
    - saturation, 213, 214
  - Traffic Server (Inktomi), 240
  - Training, 12
  - Transaction Language 1 (TL1), 380, 397
  - Transaction process (TP), 382
  - Transaction times, 234
  - Transaction-based pricing, 406
  - Transformation, 268
  - Transitioning firms, 16
  - Transmission Control Protocol (TCP), 441, 455, 457, 507
    - communication ports, 341
    - connections, 216–218
    - Intercept, 352–355
      - mode, 353
      - timers, 353, 354
    - packets, 325, 351, 508
    - ports, 326, 343
    - protocol, 335
    - traffic, 226
      - connections, 219–220
      - flow, 467
  - Transmission Control Protocol/Internet Protocol (TCP/IP), 18, 149, 467
    - network transport, 193
  - TCP/IP-based network, 246
  - Transparency, 217
  - Transparent caching, 198, 203
  - Transport layer (Layer 4), 17–19, 204, 207, 229
    - switching, 467
  - Tribal Flood Network (TFN), 340–342
  - Tribal Flood Networks 2000 (TFN2K), 340, 342
  - Trinoo, 340, 341
  - Trojan horses, 169
  - Troubleshooting, 365. *See also* Remote troubleshooting
  - Tru64 Unix, 144
  - Trunking, 463
  - TSA. *See* Total service availability
  - TTL. *See* Time-to-live
  - Type length value (TLV), 519
  - Type of Service (ToS), 456
- U**
- UDP. *See* User Datagram Protocol
  - Ultra-2 SCSI drives, 279
  - UltraBac, 162
  - Undersubscription. *See* Server
  - Unicast environment, 460
  - Uniform Resource Locator (URL), 193
    - hostname, 205
    - parsing, 206
    - switching policies, 228
  - Uninterruptible power supply (UPS), 422
  - Universal Service Order Code (USOC), 384
  - Unix, 29, 133–136, 273. *See also* Tru64 Unix
    - operating system, 134
    - solutions, 24
    - servers, 150, 357
    - systems, 357
    - Unix-based platforms, 139
    - versions, 135
  - Unix-based NOS, 135
  - Unterberg. *See* C.E. Unterberg, Towbin
  - Upgrades. *See* Planned upgrades
    - budget constraints, 216
  - Upper layers, 21
  - UPS. *See* Uninterruptible power supply
  - Upstream bandwidth, reduction, 195
  - Uptime, 262
    - amount. *See* System
  - URL. *See* Uniform Resource Locator
  - Usage, monitoring/reporting, 245
  - Usage-based pricing, 406
  - User
    - authentication, 206, 309–312
    - base, 72
    - environment, 211
    - traffic, 440
  - User Datagram Protocol (UDP), 323, 441, 455, 507
    - communication ports, 341
    - echo requests, 338
    - packets, 342, 351
    - protocol, 338

- traffic flow, 467
  - User-definable metrics, 248
  - User-definable password, 316
  - User-defined differences. *See* Content
  - User-defined policies, 455, 456
  - Usenames, 310, 311
  - USinternetworking (USi), 39, 398
  - USOC. *See* Universal Service Order Code
  - Utilization spikes, 4
  - UUNET, 42
  - UUNet, 240
- V**
- Valuations, reduction, 67–68, 89
  - Value
    - chain, migration, 91
    - offering steps, 73–75
    - proposition. *See* Internet Service Provider
    - providing, 78
  - Value-added offerings, 73
  - Value-Added Resellers (VAR), 9, 13, 16, 96, 406
  - Value-added services, 65, 72, 78–80, 258
    - deployment, 75
    - providing, 191
  - VAR. *See* Value-Added Resellers
  - Variable bit rate (VBR), 542
  - Variable-length subnet masks (VLSM), 446, 504
  - VBR. *See* Variable bit rate
  - VC. *See* Virtual channel
  - VCC, 552
  - VCI value, 552
  - Vendor-neutral representation, 16
  - Vendors. *See* Independent software vendor
    - credibility/support infrastructure, 211–212
    - equipment, 446
    - impact. *See* Content Delivery Network
    - partnerships, 400
    - products. *See* Third-party vendor products
    - SAN solutions, 295–296, 300
  - Venture firms, 64
  - Verification, 365
  - Verio, 240
  - Veritas. *See* Backup Exec; Cluster Server; NerveCenter
  - Versatile interface processor (VIP), 468
  - Vertical applications, 97
  - Video broadcast, 460
  - Video streams, 460
    - transmission, 70
  - Video traffic, 25
  - Videoconferencing, 460
  - Vingage, 243
  - VIP. *See* Versatile interface processor
  - Virtual channel (VC), 517, 518, 521
  - Virtual circuits, 19, 426
  - virtual (command), 221, 225, 227
  - Virtual connection, 542
  - Virtual IP address, 218
  - Virtual LANs (VLANs), 451, 458, 462–463
    - encapsulation, 499
    - number, 459
  - Virtual local area network Trunking Protocol (VTP), 497
  - Virtual Private Branch eXchange (VPBX), 2
  - Virtual Private Network forwarding routing (VRF), 494
  - Virtual Private Network (VPN), 23, 69, 248, 320–321, 377. *See also* Local Area Network; Managed VPNs
    - concentrator, 321
    - connection, 320
    - implementation, 388
    - services, 70
    - tunnel, 314, 483
    - usage, 419
  - Virtual Router Redundancy Protocol (VRRP), 454, 471
  - Virtual server. *See* Multiple virtual servers
    - ports, 224
  - Virtual universities, 234
  - Virtualization. *See* In-band virtualization; Storage area network
  - Virtual server, 219–220
  - Viruses. *See* Computer application. *See* Anti-virus application
    - scanning, 168–171
    - possibilities, 170
    - suggestions, 182–183
    - software. *See* Anti-virus software
  - Vividon, 243
  - Vixel, 275
  - VLANs. *See* Virtual LANs
  - VLSM. *See* Variable-length subnet masks
  - Voice
    - services, 68
    - transmission, 70
  - Voice-over Internet Protocol (VoIP), 60, 321, 448, 468
  - VoIP. *See* Voice-over Internet Protocol
  - VPBX. *See* Virtual Private Branch eXchange
  - VPI value, 552
  - VPI/VCI pair, 517, 521
  - VPN. *See* Virtual Private Network
  - VRF. *See* Virtual Private Network forwarding routing
  - VRRP. *See* Virtual Router Redundancy Protocol
  - VTP. *See* Virtual local area network Trunking Protocol

**W**

WAN. *See* Wide Area Network  
 Warp, 140  
 Watch mode, 353  
 WBEM. *See* Web-Based Enterprise Management  
 WCCP. *See* Web Cache Communication Protocol  
 Weakest link, 373–375  
 Web and Application Services (Novell), 139  
 Web-Based Enterprise Management (WBEM), 280  
 Web Cache Communication Protocol (WCCP), 493  
 WebManage, 30  
 WebQoS (HP), 24, 30–31  
 WebSphere Commerce Suite, 32  
 Weighted percentage, 230  
 Weighted Random Early Detection (WRED), 457  
 Weighted round robin (WRR), 458  
 What-if monster, 177  
 What's Up Gold, 178  
 White House, 313  
 Wide Area Network (WAN), 4, 32, 66, 320, 418  
   availability, 207  
   capability, checking, 372  
   circuit services, 427  
   connection, 265  
   connectivity, 377  
   core, 436  
   definition, 419  
   facilities, 15  
   internetworks, 436

  link, 101, 454  
     considerations, 464  
   linking, 418  
   model, 420  
   performance issues, 432  
   service levels, 365  
     data, 366  
     specifications, 17  
 Wide SCSI, 277  
 Windows 2000, 136, 144  
 Windows (Microsoft), 136–137  
 Windows NT, 144  
 WinFrame for Windows Terminals, 171  
 WinNuke, 334  
 Wire speed, impact, 284–287  
 Wire-speed throughput, 228  
 Workflow engine, 378  
 Workstation (WS), 391  
 World Wide Web (WWW/Web), 469  
   applications, 137–141  
   performance, 32  
   browsers, 172  
   caching, 242  
   definition, 189  
   function, 189–197, 252  
   designers, 32  
   development, 15  
   hosting, 91, 199  
   links, 33  
     editing/redirecting ability, 235–236  
   object database, 28  
   page, 169. *See also* Dynamic Web pages  
     load testing, 33  
   replication, 242  
   request redirection systems, 231

  servers, 139, 193, 198, 204, 244  
   sites, overwhelmed, 188  
 Web Collaboration, 60  
 Web-based GUI, 228  
   interfaces, 70  
   applications, 96  
   technologies, 15  
   caching, 460  
 Worm, 169. *See also* Xerox Corporation  
 WRED. *See* Weighted Random Early Detection  
 write memory (command), 221  
 write terminal (command), 221  
 WRR. *See* Weighted round robin  
 WS. *See* Workstation  
 WSE, 393

**X**

X.25 networks, 432  
 x86, 114  
 Xerox Corporation, 118  
   worm, 169  
 Xevo, 399, 401  
 XoIP, 230

**Y**

Yankee Group, 84

**Z**

Zero latency, 268



## **The Global Knowledge Advantage**

Global Knowledge has a global delivery system for its products and services. The company has 28 subsidiaries, and offers its programs through a total of 60+ locations. No other vendor can provide consistent services across a geographic area this large. Global Knowledge is the largest independent information technology education provider, offering programs on a variety of platforms. This enables our multi-platform and multi-national customers to obtain all of their programs from a single vendor. The company has developed the unique Competus™ Framework software tool and methodology which can quickly reconfigure courseware to the proficiency level of a student on an interactive basis. Combined with self-paced and on-line programs, this technology can reduce the time required for training by prescribing content in only the deficient skills areas. The company has fully automated every aspect of the education process, from registration and follow-up, to “just-in-time” production of courseware. Global Knowledge through its Enterprise Services Consultancy, can customize programs and products to suit the needs of an individual customer.

## **Global Knowledge Classroom Education Programs**

The backbone of our delivery options is classroom-based education. Our modern, well-equipped facilities staffed with the finest instructors offer programs in a wide variety of information technology topics, many of which lead to professional certifications.

## **Custom Learning Solutions**

This delivery option has been created for companies and governments that value customized learning solutions. For them, our consultancy-based approach of developing targeted education solutions is most effective at helping them meet specific objectives.

## **Self-Paced and Multimedia Products**

This delivery option offers self-paced program titles in interactive CD-ROM, videotape and audio tape programs. In addition, we offer custom development of interactive multimedia courseware to customers and partners. Call us at 1-888-427-4228.

## **Electronic Delivery of Training**

Our network-based training service delivers efficient competency-based, interactive training via the World Wide Web and organizational intranets. This leading-edge delivery option provides a custom learning path and “just-in-time” training for maximum convenience to students.



# Global Knowledge Courses Available

## Microsoft

- Windows 2000 Deployment Strategies
- Introduction to Directory Services
- Windows 2000 Client Administration
- Windows 2000 Server
- Windows 2000 Update
- MCSE Bootcamp
- Microsoft Networking Essentials
- Windows NT 4.0 Workstation
- Windows NT 4.0 Server
- Windows NT Troubleshooting
- Windows NT 4.0 Security
- Windows 2000 Security
- Introduction to Microsoft Web Tools

## Management Skills

- Project Management for IT Professionals
- Microsoft Project Workshop
- Management Skills for IT Professionals

## Network Fundamentals

- Understanding Computer Networks
- Telecommunications Fundamentals I
- Telecommunications Fundamentals II
- Understanding Networking Fundamentals
- Upgrading and Repairing PCs
- DOS/Windows A+ Preparation
- Network Cabling Systems

## WAN Networking and Telephony

- Building Broadband Networks
- Frame Relay Internetworking
- Converging Voice and Data Networks
- Introduction to Voice Over IP
- Understanding Digital Subscriber Line (xDSL)

## Internetworking

- ATM Essentials
- ATM Internetworking
- ATM Troubleshooting
- Understanding Networking Protocols
- Internetworking Routers and Switches
- Network Troubleshooting
- Internetworking with TCP/IP
- Troubleshooting TCP/IP Networks
- Network Management
- Network Security Administration
- Virtual Private Networks
- Storage Area Networks
- Cisco OSPF Design and Configuration
- Cisco Border Gateway Protocol (BGP) Configuration

## Web Site Management and Development

- Advanced Web Site Design
- Introduction to XML
- Building a Web Site
- Introduction to JavaScript
- Web Development Fundamentals
- Introduction to Web Databases

## PERL, UNIX, and Linux

- PERL Scripting
- PERL with CGI for the Web
- UNIX Level I
- UNIX Level II
- Introduction to Linux for New Users
- Linux Installation, Configuration, and Maintenance

## Authorized Vendor Training

### Red Hat

- Introduction to Red Hat Linux
- Red Hat Linux Systems Administration
- Red Hat Linux Network and Security Administration
- RHCE Rapid Track Certification

### Cisco Systems

- Interconnecting Cisco Network Devices
- Advanced Cisco Router Configuration
- Installation and Maintenance of Cisco Routers
- Cisco Internetwork Troubleshooting
- Designing Cisco Networks
- Cisco Internetwork Design
- Configuring Cisco Catalyst Switches
- Cisco Campus ATM Solutions
- Cisco Voice Over Frame Relay, ATM, and IP
- Configuring for Selsius IP Phones
- Building Cisco Remote Access Networks
- Managing Cisco Network Security
- Cisco Enterprise Management Solutions

### Nortel Networks

- Nortel Networks Accelerated Router Configuration
- Nortel Networks Advanced IP Routing
- Nortel Networks WAN Protocols
- Nortel Networks Frame Switching
- Nortel Networks Accelar 1000
- Comprehensive Configuration
- Nortel Networks Centillion Switching
- Network Management with Optivity for Windows

### Oracle Training

- Introduction to Oracle8 and PL/SQL
- Oracle8 Database Administration



# Custom Corporate Network Training

## Train on Cutting Edge Technology

We can bring the best in skill-based training to your facility to create a real-world hands-on training experience. Global Knowledge has invested millions of dollars in network hardware and software to train our students on the same equipment they will work with on the job. Our relationships with vendors allow us to incorporate the latest equipment and platforms into your on-site labs.

## Maximize Your Training Budget

Global Knowledge provides experienced instructors, comprehensive course materials, and all the networking equipment needed to deliver high quality training. You provide the students; we provide the knowledge.

## Avoid Travel Expenses

On-site courses allow you to schedule technical training at your convenience, saving time, expense, and the opportunity cost of travel away from the workplace.

## Discuss Confidential Topics

Private on-site training permits the open discussion of sensitive issues such as security, access, and network design. We can work with your existing network's proprietary files while demonstrating the latest technologies.

## Customize Course Content

Global Knowledge can tailor your courses to include the technologies and the topics which have the greatest impact on your business. We can complement your internal training efforts or provide a total solution to your training needs.

## Corporate Pass

The Corporate Pass Discount Program rewards our best network training customers with preferred pricing on public courses, discounts on multimedia training packages, and an array of career planning services.

## Global Knowledge Training Lifecycle

Supporting the Dynamic and Specialized Training Requirements of Information Technology Professionals

- Define Profile
- Assess Skills
- Design Training
- Deliver Training
- Test Knowledge
- Update Profile
- Use New Skills

## **Global Knowledge**

Global Knowledge programs are developed and presented by industry professionals with “real-world” experience. Designed to help professionals meet today’s interconnectivity and interoperability challenges, most of our programs feature hands-on labs that incorporate state-of-the-art communication components and equipment.

### **ON-SITE TEAM TRAINING**

Bring Global Knowledge’s powerful training programs to your company. At Global Knowledge, we will custom design courses to meet your specific network requirements. Call (919)-461-8686 for more information.

### **YOUR GUARANTEE**

Global Knowledge believes its courses offer the best possible training in this field. If during the first day you are not satisfied and wish to withdraw from the course, simply notify the instructor, return all course materials and receive a 100% refund.

### **REGISTRATION INFORMATION**

In the US:

call: (888) 762-4442

fax: (919) 469-7070

visit our website:

[www.globalknowledge.com](http://www.globalknowledge.com)

Get More at [access.globalknowledge](http://access.globalknowledge.com)

## The premier online information source for IT professionals

You've gained access to a Global Knowledge information portal designed to inform, educate and update visitors on issues regarding IT and IT education.

Get what you want when you want it at the [access.globalknowledge](http://access.globalknowledge.com) site:

**Choose personalized technology articles** related to *your* interests. Access a new article, review, or tutorial regularly throughout the week customized to what you want to see.

**Keep learning in between Global courses** by taking advantage of chat sessions with other users or instructors. Get the tips, tricks and advice that you need today!

**Make your point** in the Access.Globalknowledge community with threaded discussion groups related to technologies and certification.

**Get instant course information** at your fingertips. Customized course calendars showing you the courses you want when and where you want them.

**Get the resources you need** with online tools, trivia, skills assessment and more!

All this and more is available now on the web at [access.globalknowledge](http://access.globalknowledge.com). VISIT TODAY!



<http://access.globalknowledge.com>



# Syngress Publishing's Sweepstake Terms

## OFFICIAL RULES - NO PURCHASE NECESSARY

### 1) TIMING

The contest (the "Contest") begins March 1, 2001 at 9:00 a.m. EST and ends November 30, 2001 at 11:59 p.m. EST (the "Entry Period"). You must enter the contest during the Entry Period.

### 2) THE PRIZES

Three (3) prizes will be awarded: (a) a Sony DVD Player ("1<sup>st</sup> Prize"); (b) a Palm Pilot V ("2<sup>nd</sup> Prize"); and (c) a Rio MP3 Player ("3<sup>rd</sup> Prize"). One of each prize will be awarded. The approximate retail value of the three prizes is as follows: (a) the Sony DVD Player is approximately \$595; (b) the Palm Pilot V is approximately \$399; and (c) the Rio MP3 Player is approximately \$299.

Sponsors make no warranty, guaranty or representation of any kind concerning any prize. Prize values are subject to change.

### 3) ELIGIBILITY REQUIREMENTS

No purchase is necessary. Contest is void in Puerto Rico, and where prohibited by law. Employees of Syngress Publishing, Inc. (the "Sponsor") and their affiliates, subsidiaries, officers, agents or any other person or entity directly associated with the contest (the "Contest Entities") and the immediate family members and/or persons living in the same household as such persons are not eligible to enter the Contest.

This contest is open only to people that meet the following requirements:

- legal residents of the United States
- Must be at least 21 years of age or older at the time of winning
- Must own a major credit card

**4) HOW TO ENTER:** No purchase is necessary to enter. Contestants can enter by mail (see below) or may enter on the Syngress website located at: [www.syngress.com/sweepstake.html](http://www.syngress.com/sweepstake.html). **ONLY ONE ENTRY PER PERSON OR E-MAIL ADDRESS PER HOUSEHOLD WILL BE ACCEPTED.**

No purchase is necessary to enter. To enter by mail, print your name, address, daytime telephone number, email address and age. Mail this in a hand-addressed envelope to: **Syngress Publishing Contest, Syngress Publishing, Inc., 800 Hingham Street, Rockland, MA 02370**. All mail entries must be postmarked before November 15, 2001.

Sponsor assumes no responsibility for lost, late, or misdirected entries or for any computer, online, telephone, or human error or technical malfunctions that may occur. Incomplete mail entries are void. All entries become the property of Sponsor and will not be returned.

If a prize notification or prize is returned to Sponsor or its fulfillment companies as undeliverable for any reason, it will be awarded to an alternate. If necessary, due to unavailability, a prize of equal or great value will be awarded at the discretion of the Sponsor. Prizes are not transferable, assignable or redeemable for cash.

By entering the Contest on the Sponsor Internet site, you may occasionally receive promotion announcements from Sponsor through e-mail. If you no longer wish to receive these e-mails, you may cease your participation in such promotions by sending an e-mail to [promotions@syngress.com](mailto:promotions@syngress.com) with your First Name, Last Name, and your e-mail address.

**5) WINNER SELECTION/DEADLINE DATES:** Random drawings will be conducted by the Sponsor from among all eligible entries. Odds of winning the prize depend on the number of eligible entries received. The first drawing will be for the winner of the 1<sup>st</sup> Prize, then a drawing will be held from all remaining eligible entries for the winner of the 2<sup>nd</sup> Prize and finally a drawing will be held from all remaining eligible entries for the winner of the 3<sup>rd</sup> Prize. These drawings will occur on December 1, 2001, at the offices of Syngress Publishing, Inc., 800 Hingham Street, Rockland, MA 02370. The decisions by the Sponsor shall be final and binding in all respects.

**6) GENERAL CONDITIONS:** Contest entrants agree to be bound by the terms of these official rules. The laws of the Commonwealth of Massachusetts and the United States govern this Contest, and the state and federal courts located in Suffolk and Middlesex Counties in the Commonwealth of Massachusetts shall be the sole jurisdiction for any disputes related to the Contest. All federal, state, and local laws and regulations apply. Winners will be notified via e-mail and/or U.S. Mail within two (2) weeks of prize drawing. Winners will be required to execute and return an Affidavit of Eligibility and Release of Liability and where legal, Publicity Release within 14 days following the date of issuance of notification. Non-compliance within this time period or return of any prize/prize notification as undeliverable may result in disqualification and selection of an alternate winner. Acceptance of prize constitutes permission for Sponsor to use winner's name and likeness for advertising and promotional purposes without additional compensation unless prohibited by law. BY ENTERING, PARTICIPANTS RELEASE AND HOLD HARMLESS SYNGRESS PUBLISHING, INC., AND ITS RESPECTIVE PARENT CORPORATIONS, SUBSIDIARIES, AFFILIATES, DIRECTORS, OFFICERS, PRIZE SUPPLIERS, EMPLOYEES AND AGENTS FROM ANY AND ALL LIABILITY OR ANY INJURIES, LOSS OR DAMAGE OF ANY KIND ARISING FROM OR IN CONNECTION WITH THE CONTEST OR ACCEPTANCE OR USE OF THE PRIZES WON.

**7) INTERNET:** If for any reason this contest is not capable of running as planned due to infection by computer virus, bugs, tampering, unauthorized intervention, fraud, technical failures, or any other causes beyond the control of the Sponsor which

corrupt or affect the administration, security, fairness, integrity, or proper conduct of this contest, the Sponsor reserves the right, at its sole discretion, to disqualify any individual who tampers with the entry process, and to cancel, terminate, modify, or suspend the online portion of the contest. The Sponsor assumes no responsibility for any error, omission, interruption, deletion, defect, delay in operation or transmission, communications line failure, theft or destruction or unauthorized access to, or alteration of, entries. Sponsor is not responsible for any problems or technical malfunction of any telephone network or telephone lines, computer on-line systems, servers, or providers, computer equipment, software, failure of any e-mail or entry to be received by Sponsor on account of technical problems, human error or traffic congestion on the Internet or at any Web site, or any combination thereof, including any injury or damage to participant's or any other person's computer relating to or resulting from participation in the Contest or downloading any materials in the Contest. CAUTION: ANY ATTEMPT TO DELIBERATELY DAMAGE ANY WEB SITE OR UNDERMINE THE LEGITIMATE OPERATION OF THE CONTEST IS A VIOLATION OF CRIMINAL AND CIVIL LAWS AND SHOULD SUCH AN ATTEMPT BE MADE, SPONSOR RESERVES THE RIGHT TO SEEK DAMAGES OR OTHER REMEDIES FROM ANY SUCH PERSON (S) RESPONSIBLE FOR THE ATTEMPT TO THE FULLEST EXTENT PERMITTED BY LAW. In the event of a dispute as to the identity of a winner based on an e-mail address, the winning entry will be declared made by the authorized account holder of the e-mail address submitted at time of entry. "Authorized account holder" is defined as the natural person who is assigned to an e-mail address by an Internet access provider, on-line service provider, or other organization (e.g., business, educational, institution, etc.) that is responsible for assigning e-mail addresses for the domain associated with the submitted e-mail address.

**8) WHO WON:** Winners who enter on the web site will be notified by e-mail and winners who had entered via mail will be notified by mail. The winners will also be posted on our web site. Alternatively, to receive the names of the winners please send a self addressed stamped envelope to: Syngress Publishing Contest, care of Syngress Publishing, Inc., 800 Hingham Street, Rockland, MA 02370.

The Sponsor of this sweepstakes is Syngress Publishing, Inc., 800 Hingham Street, Rockland, MA 02370.



# SYNGRESS SOLUTIONS...



AVAILABLE NOW  
ORDER at  
[www.syngress.com](http://www.syngress.com)

## Configuring Citrix MetaFrame for Windows 2000 Terminal Services

Citrix MetaFrame can deliver Windows-based applications to any user, anywhere regardless of network connection, LAN protocol, or client operating system. *Configuring Citrix MetaFrame for Windows 2000 Terminal Services* is written for system administrators who are deploying Citrix MetaFrame in a Windows 2000 environment. It examines MetaFrame's newest features and enhancements, as well as Citrix's Independent Computing Architecture (ICA). It explores how ICA, in conjunction with MetaFrame, will transform the way in which software is developed, deployed, and maintained in server-based computing environments. Finally, the book shows how to integrate Windows- and UNIX-based networks over the Web using MetaFrame.

ISBN: 1-928994-18-0

Price: \$49.95 US, \$77.95 CAN

AVAILABLE NOW  
ORDER at  
[www.syngress.com](http://www.syngress.com)

## Hack Proofing Your E-commerce Site

*From the authors of the bestselling Hack Proofing Your Network.* E-Commerce giants, previously thought to be impenetrable are now being exposed as incredibly vulnerable. This book gives e-commerce architects and engineers insight into the tools and techniques used by hackers to compromise sites. The security of e-commerce sites is even more imperative than non-commerce sites, because of the added responsibility of maintaining customers' personal and financial information. The book will provide web architects and engineers all the information they need to design and implement security measures.

ISBN: 1-928994-27-X

Price: \$49.95 US, \$77.95 CAN



AVAILABLE NOW  
ORDER at  
[www.syngress.com](http://www.syngress.com)

## Designing SQL Server 2000 Databases for .NET Enterprise Servers

Microsoft recently announced its .NET Enterprise Server line and its commitment to .NET as Microsoft's application architecture model. SQL Server 2000 is the first .NET Enterprise Server available offering the data storage and management component of the .NET line. *Designing SQL Server 2000 Databases* details the SQL Server 2000 product and its role in the .NET product line. Building on the successful delivery of SQL Server 7.0, SQL Server 2000 presents integration and maturity of many features that were announced with and after SQL Server 7.0.

ISBN: 1-928994-19-9

Price: \$49.95 US, \$77.95 CAN



solutions@syngress.com

SYNGRESS®