



Data Center Fundamentals

Understand Data Center network design
and infrastructure architecture, including
load balancing, SSL, and security

Data Center Fundamentals

Mauricio Arregoces

Maurizio Portolani

Copyright © 2004 Cisco Systems, Inc.

Published by:

Cisco Press

800 East 96th Street

Indianapolis, IN 46240 USA

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

ISBN: 1-58705-023-4

Library of Congress Cataloging-in-Publication Number: 2001086631

Printed in the United States of America 6 7 8 9 0

Sixth Printing August 2009

Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems, Inc., cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Warning and Disclaimer

This book is designed to provide information about Data Center technologies. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an “as is” basis. The authors, Cisco Press, and Cisco Systems, Inc., shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the author and are not necessarily those of Cisco Systems, Inc.

Feedback Information

At Cisco Press, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers’ feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this book or otherwise alter it to better suit your needs, you can contact us through e-mail at feedback@ciscopress.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

Corporate and Government Sales

Cisco Press offers excellent discounts on this book when ordered in quantity for bulk purchases or special sales. For more information, please contact:

U.S. Corporate and Government Sales 1-800-382-3419 corpsales@pearsontechgroup.com

For sales outside of the U.S. please contact: **International Sales** 1-317-581-3793 international@pearsontechgroup.com

Publisher	John Wait
Editor-In-Chief	John Kane
Cisco Representative	Anthony Wolfenden
Cisco Press Program Manager	Nannette M. Noble
Production Manager	Patrick Kanouse
Development Editors	Christopher Cleveland Betsey Henkels
Senior Project Editor	Sheri Cain
Copy Editors	Krista Hansing, Kris Simmons
Technical Editors	Mario Baldi, Robert Batz, Mark Gallo, Ron Hromoko, Fabio Maino, Scott Van de Houten, Stefano Testa, Brian Walck
Team Coordinator	Tammi Barnett
Cover Designer	Louisa Adair
Composition	Octal Publishing, Inc.
Indexers	Tim Wright, Eric Schroeder
Proofreader	Angela Rosio



Corporate Headquarters
 Cisco Systems, Inc.
 170 West Tasman Drive
 San Jose, CA 95134-1706
 USA
www.cisco.com
 Tel: 408 526-4000
 800 553-NETS (6387)
 Fax: 408 526-4100

European Headquarters
 Cisco Systems International BV
 Haarlerbergpark
 Haarlerbergweg 13-19
 1101 CH Amsterdam
 The Netherlands
www-europe.cisco.com
 Tel: 31 0 20 357 1000
 Fax: 31 0 20 357 1100

Americas Headquarters
 Cisco Systems, Inc.
 170 West Tasman Drive
 San Jose, CA 95134-1706
 USA
www.cisco.com
 Tel: 408 526-7660
 Fax: 408 527-0883

Asia Pacific Headquarters
 Cisco Systems, Inc.
 Capital Tower
 168 Robinson Road
 #22-01 to #29#01
 Singapore 068912
www.cisco.com
 Tel: +65 6317 7777
 Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the Cisco.com Web site at www.cisco.com/go/offices.

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Czech Republic
 Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel • Italy
 Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal
 Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
 Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2003 Cisco Systems, Inc. All rights reserved. CCIP, CCSP, the Cisco Arrow logo, the Cisco Powered Network mark, the Cisco Systems Verified logo, Cisco Unity, Follow Me Browsing, FormShare, iQ Net Readiness Scorecard, Networking Academy, and ScriptShare are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, The Fastest Way to Increase Your Internet Quotient, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, AIRST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, Internet Quotient, IOS, IPTV, iQ Expertise, the iQ logo, LightStream, MGX, MICa, the Networkers logo, Network Registrar, Packet, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, SlideCast, SMARTnet, StrataView Plus, Stratm, SwitchProbe, TeleRouter, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0303R)

Printed in the USA

Introduction

Data Centers are complex systems encompassing a wide variety of technologies that are constantly evolving. Designing and maintaining a Data Center network requires skills and knowledge that range from routing and switching to load balancing and security, including the essential knowledge of servers and applications.

This book addresses both fundamental information such as the protocols used by switches and routers; the protocols used in application environments; the network technology used to build the Data Center infrastructure and secure, scale, and manage the application environments; and design best practices. We hope this book becomes your Data Center reference on protocols, technology, and design.

Motivation for Writing This Book

While speaking to networkers abroad on the topic of server load balancing, we realized that we could only convey the benefits of the technology by explaining application layer information and describing the larger design issues common in application environments.

Often through discussions with customers, the subjects related to load balancing take a back seat as issues of integration with the entire Data Center take the forefront. This book attempts to cover the breadth and depth of the Data Center IP network. The storage network and distributed Data Center topics will be the subjects of other books.

Having designed campus and Data Center networks, and having developed and supported technologies that are often referred to as *content networking* (load balancing, Secure Socket Layer [SSL] offloading, and DNS routing), we felt the need for a book that described these topics in a single place and focused on what is relevant to the Data Center. This area is what this book is about: it is an all-encompassing view of Data Centers from routing and switching technologies to application-aware technologies.

Who Should Read This Book

This book is intended for any person or organization seeking to understand Data Center networks: the fundamental protocols used by the applications and the network, the typical network technologies, and their design aspects. The book is meant to be both a reference on protocols and technology and a design and implementation guide for personnel responsible for planning, designing, implementing, and operating Data Center networks.

Chapter Organization

This book has six parts. This book is designed to be read in order from the overview of the Data Center environment, through the server farms and infrastructure protocols, to security and load-balancing concepts, before you reach the Data Center design chapters. This organization also allows you to go directly to the desired chapter if you already know the information in the previous chapters.

Part I, “An Introduction to Server Farms,” includes chapters that contain an overview of the architecture of Data Centers, servers, and applications. This part also introduces the security and load-balancing technology:

- Chapter 1, “Overview of Data Centers,” presents Data Center environments, the Data Center architecture, and services that are used as a guide to the rest of the book.
- Chapter 2, “Server Architecture Overview,” explores the architecture of servers. This chapter covers topics such as how servers process TCP and User Datagram Protocol (UDP) traffic, how processes and threads are used, and server health.

- Chapter 3, “Application Architectures Overview,” explores the application environments and how applications are architected. This chapter includes discussions on the relation between the application architectures and the design of the Data Center, the n-tier model, HTML and XML, user-agent technologies, web server technologies, and clustering technologies. This chapter introduces application concepts that are developed in Chapter 18 and Chapter 19.
- Chapter 4, “Data Center Design Overview,” discusses the types of server farms on Data Centers, generic and alternative Layer 2 and Layer 3 designs, multitier designs, high availability, Data Center services, and trends that might affect Data Center designs.
- Chapter 5, “Data Center Security Overview,” discusses threats, vulnerabilities and common attacks, network security devices such as firewalls and intrusion detection systems (IDSs), and other fundamental security concepts such as cryptography; VPNs; and authentication, authorization and accounting (AAA).
- Chapter 6, “Server Load-Balancing Overview,” discusses reasons for load balancing, fundamental load-balancing concepts, high-availability considerations, and generic load-balancing architectures. The fundamental load-balancing concepts include Layer 4 and Layer 5 load balancing, session tracking, session persistence, and server health.

Part II, “Server Farm Protocols,” explores the fundamental protocols used in server farms:

- Chapter 7, “IP, TCP, and UDP,” explores the protocol headers details and their relevance to network design issues.
- Chapter 8, “HTTP and Related Concepts,” discusses key concepts such as Uniform Resource Identifiers (URIs) and URLs, Multipurpose Internet Mail Extension (MIME) and its relation to HTTP entities, and HTTP header details. Chapter 8 provides additional information on the operation of HTTP, the different versions and their performance characteristics.
- Chapter 9, “SSL and TLS,” discusses SSL operations with specific focus on SSL session establishment, cipher-suites, and SSL performance considerations. Chapter 15 provides additional information on the public-key infrastructure (PKI), certificates, and more security-related aspects of SSL.
- Chapter 10, “DNS Essentials and Site-Selection Considerations,” explores how the DNS namespace is organized, the DNS components in the Internet, how the DNS resolution process works, DNS configuration options, DNS server placement in the network, and how to use DNS to distribute application requests to multiple Data Centers.
- Chapter 11, “Streaming Protocols Overview,” discusses HTTP and real streaming, the use of TCP and UDP in streaming, analog and digital video, coders-decoders (codecs), packetization, the streaming transport formats, unicast, multicast and stream splitting, and encoding mechanisms.

Part III, “Infrastructure Protocols,” explores the fundamental Layer 2 and Layer 3 protocols as well as IBM Data Center technologies:

- Chapter 12, “Layer 2 Protocol Essentials,” discusses Ethernet frame types; the difference between unicast, multicast, and broadcast frames; physical layer characteristics of Ethernet technologies; jumbo frames; trunks and channels; and a variety of spanning-tree concepts. Chapter 20 provides the design best practices applied to the concepts described in this chapter.
- Chapter 13, “Layer 3 Protocol Essentials,” discusses the Address Resolution Protocol (ARP); gateway redundancy protocols such as Hot Standby Router Protocol (HSRP), VRRP and GLBP; and routing-protocol essentials for Open Shortest Path First (OSPF) and Enhanced Interior Gateway Routing Protocol (EIGRP). Chapter 20 provides the design best practices applied to the concepts described in this chapter.

- Chapter 14, “IBM Data Center Technology,” discusses mainframe attachment options, IBM networking, Systems Network Architecture (SNA) switching, Sysplex, TN3270, and current IBM Data Center designs.

Part IV, “Security and Server Load Balancing,” explores the security protocols and technology, load-balancing operations, server health management, session tracking and cookies, and persistence mechanisms on load balancers:

- Chapter 15, “Security Protocols and Technologies,” discusses cryptography, U.S. government-related topics about cryptography, PKI, transport security protocols (SSL and IP Security [IPSec]), authentication protocols and technologies, and network management security. This chapter also complements Chapter 9 with regards to the security design aspects of SSL and introduces the concept of SSL VPNs.
- Chapter 16, “Load-Balancing Modes and Predictors,” discusses the load-balancing modes of operation, server load-balancing algorithms, and cache farm load-balancing algorithms.
- Chapter 17, “Server Health Management,” discusses server health management through load balancers, SNMP, server failure detection and checking, in-band and out-of-band probes, and case studies on server checking for web hosting and e-commerce applications.
- Chapter 18, “Session Tracking and Cookies,” explores the concept of user sessions from an application point of view. This chapter explains nonpersistent cookies, cookies in general, how servers track user sessions, session persistence on clusters of servers, and the challenges of dealing with HTTP and HTTPS. Chapter 19 further expands the topic of session persistence in load-balancing deployments.
- Chapter 19, “Persistence Mechanisms on Load Balancers,” explains session persistence in relation to load balancing; discusses key persistence mechanisms, including source-IP sticky, cookie-URL sticky, HTTP redirection sticky, and SSL sticky; and presents a case study using an e-commerce application. Chapter 19 is based on the applications introduced in Chapter 3 and Chapter 18.

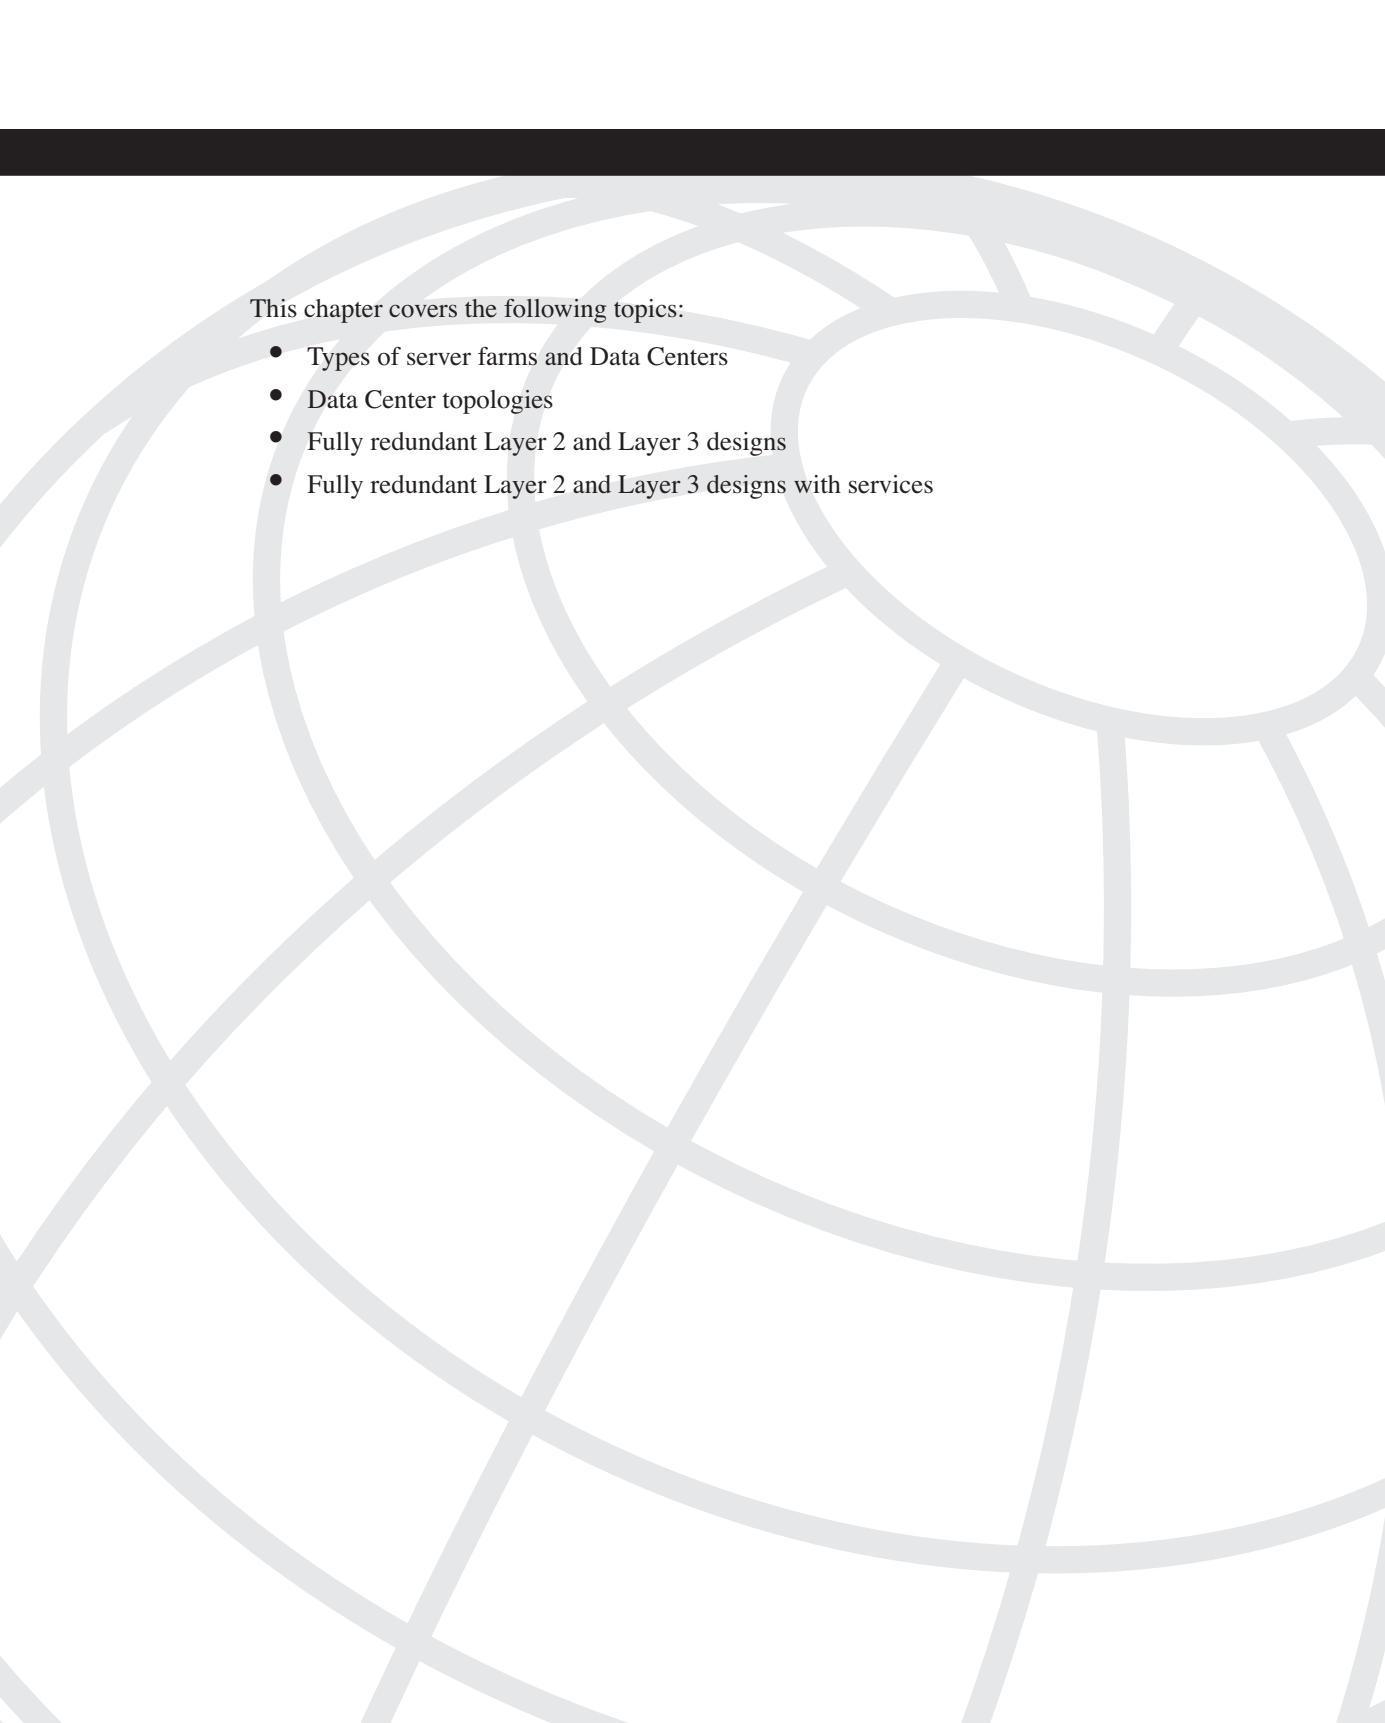
Part V, “Data Center Design,” explores the details behind designing the Data Center infrastructure, the integration of security into the infrastructure design, and the performance of Data Center devices:

- Chapter 20, “Designing the Data Center Infrastructure,” discusses router switching paths, essential Data Center design concepts, and the design best practices of the infrastructure by explaining the configuration of Layer 2 and Layer 3 features and protocols that are described in Chapter 12 and 13.
- Chapter 21, “Integrating Security into the Infrastructure,” discusses the concept of security zones and how to design application security at the Internet Edge and at intranet server farms. This chapter explains alternative designs and how to implement secure management.
- Chapter 22, “Performance Metrics of Data Center Devices,” discusses the Data Center traffic patterns and performance metrics of various Data Center devices, including proposed metrics for devices for which there are none and no standard methodology exists (such as load balancers and SSL offloaders).

Part VI, “Appendices,” is the final part of this book:

- Appendix A, “Character Sets,” covers multiple character sets, including ASCII, the extended ASCII sets, and the ISO-8859-1 set.
- Appendix B, “HTTP Header Fields,” explains the details of HTTP header fields that were not described in Chapter 8.
- Appendix C, “Video Encoding Mechanisms,” explains the removal of special and temporal redundancy in codecs with special focus on MPEG.
- Appendix D, “Loopback Interface Configuration Procedures,” provides an explanation about configuring a machine with multiple IP addresses used as loopbacks for certain load-balancing modes of operation.

- Appendix E, “Configuring Servers to Insert Cookies,” examines several alternatives for configuring cookie insertion on web servers.
- Appendix F, “Client-Side and Server-Side Programming,” provides excerpts of client-side programs to help you understand the differences and similarities between JavaScripts, Java applets, and ActiveX controls. The section on server-side programming explains the differences between CGI, servlets, and Active Server Pages (ASP) in terms of operating-system implications (threads versus processes). This appendix explains the adoption of certain technologies in today’s enterprise applications and the performance and availability implications.



This chapter covers the following topics:

- Types of server farms and Data Centers
- Data Center topologies
- Fully redundant Layer 2 and Layer 3 designs
- Fully redundant Layer 2 and Layer 3 designs with services

CHAPTER 4

Data Center Design Overview

This chapter focuses on three main properties of Data Center architectures: scalability, flexibility, and high availability. Data Centers are rapidly evolving to accommodate higher expectations for growth, consolidation, and security. Although the traditional Layer 2 and Layer 3 designs have not changed drastically over the last few years, stringent demands for uptime and service availability, coupled with new technology and protocols, make the design efforts more challenging and demanding.

Demands for scalability, flexibility, and high availability can be summarized as follows:

- **Scalability**—The Data Center must support fast and seamless growth without major disruptions.
- **Flexibility**—The Data Center must support new services without a major overhaul of its infrastructure.
- **High availability**—The Data Center must have no single point of failure and should offer predictable uptime (related to hard failures).

NOTE

A hard failure is a failure in which the component must be replaced to return to an operational steady state.

Scalability translates into the capability to sustain rapid growth in performance, the number of devices hosted in the Data Center, and the amount and quality of the services offered. Higher performance implies tolerance to very short-term changes in traffic patterns without packet loss and longer-term plans mapping growth trends to the capacity of the Data Center.

Scalability on the number of hosted devices refers to being capable of seamlessly adding more ports for servers, routers, switches, and any other service devices, such as server load balancers, firewalls, IDSs, and SSL offloaders. Higher density also includes slot density because the number of slots ultimately determines the potential growth of the system.

Flexibility translates into designs that accommodate new service offerings without requiring the complete redesign of the architecture or drastic changes outside the normal periods scheduled for maintenance. The approach to flexibility is a modular design in which the characteristics of the modules are known, and the steps to add more modules are simple.

High availability translates into a fully redundant architecture in which all possible hard failures are predictable and deterministic. This implies that each possible component's failure has a predetermined failover and fallback time, and that the worst-case scenario for a failure condition is still within the acceptable failover limits and is within the requirements as measured from an application availability viewpoint. This means that although the time of failure and recovery of a network component should be predictable and known, the more important time involves the user's perception of the time to recover application service.

NOTE

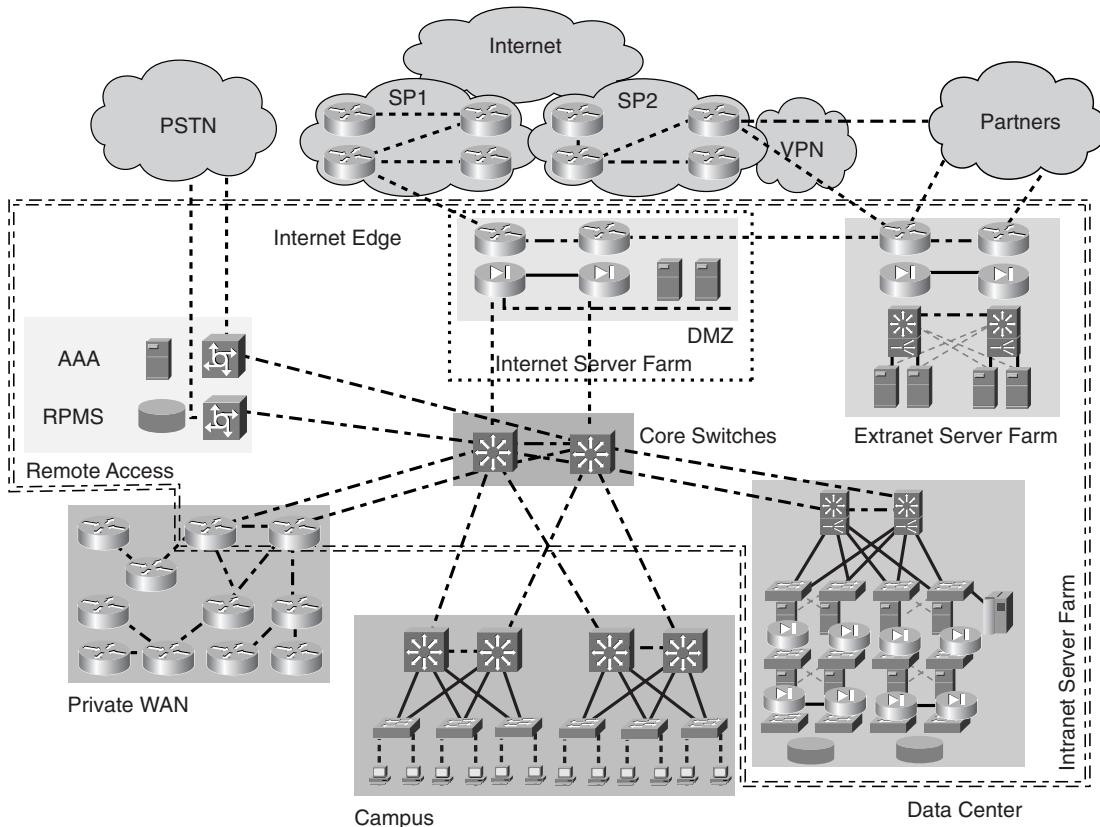
After a failure, the recovery time could be measured from the perspective of the Layer 2 environment (the spanning tree) or from a Layer 3 perspective (the routed network), yet the application availability ultimately matters to the user. If the failure is such that the user connection times out, then, regardless of the convergence time, the network convergence does not satisfy the application requirements. In a Data Center design, it is important to measure recovery time from the perspectives of both the network and the application to ensure a predictable network recovery time for the user (application service).

Figure 4-1 presents an overview of the Data Center, which, as a facility, includes a number of the building blocks and components of the larger enterprise network architecture.

This book deals primarily with the engineering of application environments and their integration to the remaining enterprise network. Different types of server farms support the application environments, yet this book focuses on understanding, designing, deploying, and maintaining the server farms supporting intranet application environments. The actual engineering of the different server farm types—Internet, extranet, and intranet server farms—does not vary much from type to type; however, their integration with the rest of the architecture is different. The design choices that differ for each type of server farm are the result of their main functional purpose. This leads to a specific location for their placement, security considerations, redundancy, scalability, and performance. In addition to the server farm concepts, a brief discussion on the types of server farms further clarifies these points.

NOTE

The figures in this chapter contain a wide variety of Cisco icons. Refer to the section, “Icons Used in This Book” (just before the “Introduction”) for a list of icons and their descriptions.

Figure 4-1 Overview of Data Center Topology

Types of Server Farms and Data Centers

As depicted in Figure 4-1, three distinct types of server farms exist:

- Internet
- Extranet
- Intranet

All three types reside in a Data Center and often in the same Data Center facility, which generally is referred to as the *corporate Data Center* or *enterprise Data Center*. If the sole purpose of the Data Center is to support Internet-facing applications and server farms, the Data Center is referred to as an *Internet Data Center*.

Server farms are at the heart of the Data Center. In fact, Data Centers are built to support at least one type of server farm. Although different types of server farms share many architectural requirements, their objectives differ. Thus, the particular set of Data Center requirements depends on which type of server farm must be supported. Each type of server farm has a distinct set of infrastructure, security, and management requirements that must be addressed in the design of the server farm. Although each server farm design and its specific topology might be different, the design guidelines apply equally to them all. The following sections introduce server farms.

Internet Server Farms

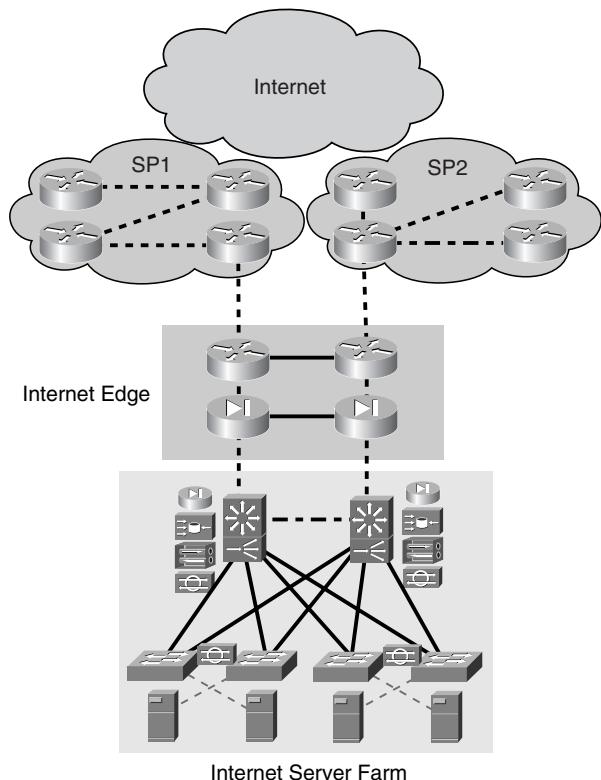
As their name indicates, Internet server farms face the Internet. This implies that users accessing the server farms primarily are located somewhere on the Internet and use the Internet to reach the server farm. Internet server farms are then available to the Internet community at large and support business-to-consumer services. Typically, internal users also have access to the Internet server farms. The server farm services and their users rely on the use of web interfaces and web browsers, which makes them pervasive on Internet environments.

Two distinct types of Internet server farms exist. The dedicated Internet server farm, shown in Figure 4-2, is built to support large-scale Internet-facing applications that support the core business function. Typically, the core business function is based on an Internet presence or Internet commerce.

In general, dedicated Internet server farms exist to sustain the enterprise's e-business goals. Architecturally, these server farms follow the Data Center architecture introduced in Chapter 1, "Overview of Data Centers," yet the details of each layer and the necessary layers are determined by the application environment requirements. Security and scalability are a major concern in this type of server farm. On one hand, most users accessing the server farm are located on the Internet, thereby introducing higher security risks; on the other hand, the number of likely users is very high, which could easily cause scalability problems.

The Data Center that supports this type of server farm is often referred to as an Internet Data Center (IDC). IDCs are built both by enterprises to support their own e-business infrastructure and by service providers selling hosting services, thus allowing enterprises to collocate the e-business infrastructure in the provider's network.

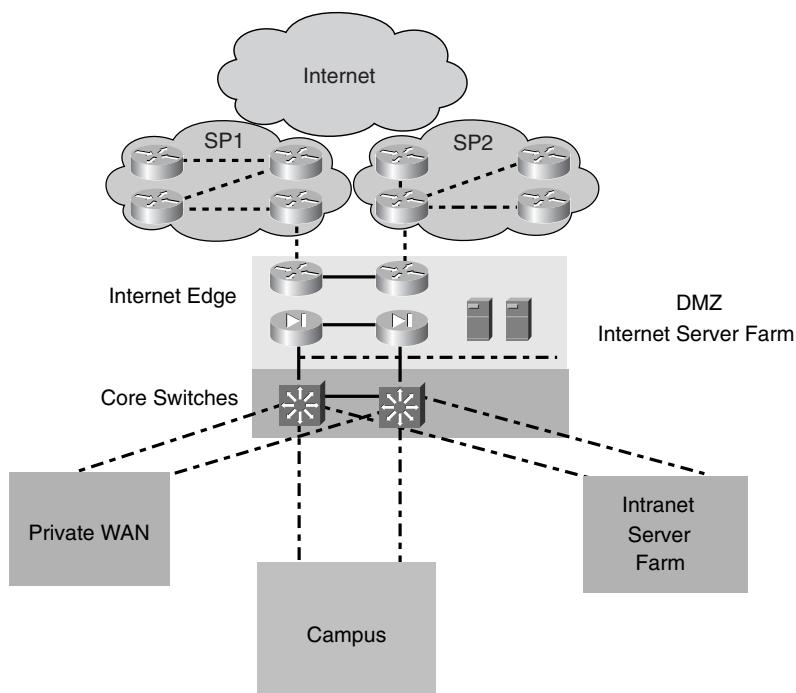
The next type of Internet server farm, shown in Figure 4-3, is built to support Internet-based applications in addition to Internet access from the enterprise. This means that the infrastructure supporting the server farms also is used to support Internet access from enterprise users. These server farms typically are located in the demilitarized zone (DMZ) because they are part of the enterprise network yet are accessible from the Internet. These server farms are referred to as DMZ server farms, to differentiate them from the dedicated Internet server farms.

Figure 4-2 Dedicated Internet Server Farms

These server farms support services such as e-commerce and are the access door to portals for more generic applications used by both Internet and intranet users. The scalability considerations depend on how large the expected user base is. Security requirements are also very stringent because the security policies are aimed at protecting the server farms from external users while keeping the enterprise's network safe. Note that, under this model, the enterprise network supports the campus, the private WAN, and the intranet server farm.

NOTE

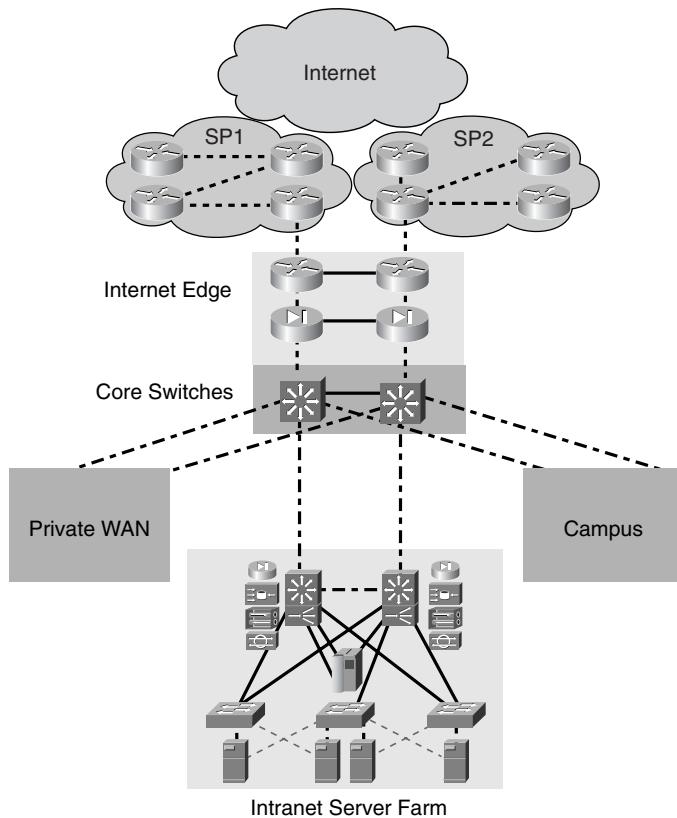
Notice that Figure 4-3 depicts a small number of servers located on a segment off the firewalls. Depending on the requirements, the small number of servers could become hundreds or thousands, which would change the topology to include a set of Layer 3 switches and as many Layers 2 switches for server connectivity as needed.

Figure 4-3 DMZ Server Farms

Intranet Server Farms

The evolution of the client/server model and the wide adoption of web-based applications on the Internet was the foundation for building intranets. Intranet server farms resemble the Internet server farms in their ease of access, yet they are available only to the enterprise's internal users. As described earlier in this chapter, intranet server farms include most of the enterprise-critical computing resources that support business processes and internal applications. This list of critical resources includes midrange and mainframe systems that support a wide variety of applications. Figure 4-4 illustrates the intranet server farm.

Notice that the intranet server farm module is connected to the core switches that form a portion of the enterprise backbone and provide connectivity between the private WAN and Internet Edge modules. The users accessing the intranet server farm are located in the campus and private WAN. Internet users typically are not permitted access to the intranet; however, internal users using the Internet as transport have access to the intranet using virtual private network (VPN) technology.

Figure 4-4 Intranet Server Farms

The Internet Edge module supports several functions that include the following:

- Securing the enterprise network
- Controlling Internet access from the intranet
- Controlling access to the Internet server farms

The Data Center provides additional security to further protect the data in the intranet server farm. This is accomplished by applying the security policies to the edge of the Data Center as well as to the applicable application tiers when attempting to harden communication between servers on different tiers. The security design applied to each tier depends on the architecture of the applications and the desired security level.

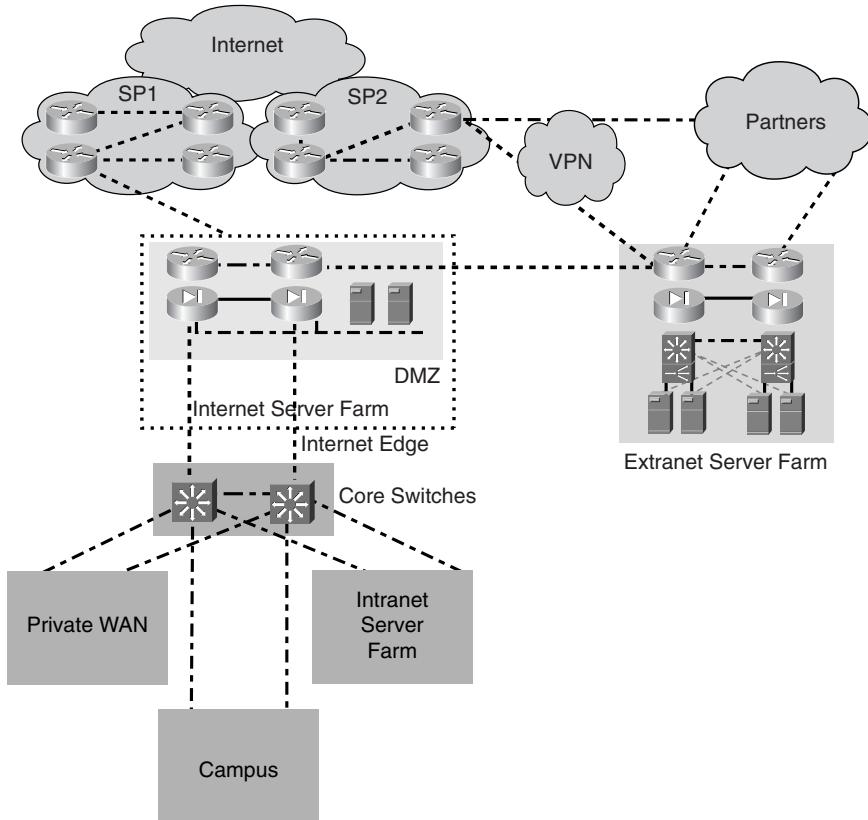
The access requirements of enterprise users dictate the size and architecture of the server farms. The growing number of users, as well as the higher load imposed by rich applications, increases the demand placed on the server farm. This demand forces scalability to become a critical design criterion, along with high availability, security, and management.

Extranet Server Farms

From a functional perspective, extranet server farms sit between Internet and intranet server farms. Extranet server farms continue the trend of using web-based applications, but, unlike Internet- or intranet-based server farms, they are accessed only by a selected group of users that are neither Internet- nor intranet-based. Extranet server farms are mainly available to business partners that are considered external yet trusted users. The main purpose for extranets is to improve business-to-business communication by allowing faster exchange of information in a user-friendly and secure environment. This reduces time to market and the cost of conducting business. The communication between the enterprise and its business partners, traditionally supported by dedicated links, rapidly is being migrated to a VPN infrastructure because of the ease of the setup, lower costs, and the support for concurrent voice, video, and data traffic over an IP network.

As explained previously, the concept of extranet is analogous to the IDC, in that the server farm is at the edge of the enterprise network. Because the purpose of the extranet is to provide server farm services to trusted external end users, there are special security considerations. These security considerations imply that the business partners have access to a subset of the business applications but are restricted from accessing the rest of the enterprise network. Figure 4-5 shows the extranet server farm. Notice that the extranet server farm is accessible to internal users, yet access from the extranet to the intranet is prevented or highly secured. Typically, access from the extranet to the intranet is restricted through the use of firewalls.

Many factors must be considered in the design of the extranet topology, including scalability, availability, and security. Dedicated firewalls and routers in the extranet are the result of a highly secure and scalable network infrastructure for partner connectivity, yet if there are only a small number of partners to deal with, you can leverage the existing Internet Edge infrastructure. Some partners require direct connectivity or dedicated private links, and others expect secure connections through VPN links. The architecture of the server farm does not change whether you are designing Internet or intranet server farms. The design guidelines apply equally to all types of server farms, yet the specifics of the design are dictated by the application environment requirements.

Figure 4-5 Extranet Server Farms

The following section discusses the types of Data Centers briefly mentioned in this section.

Internet Data Center

Internet Data Centers (IDCs) traditionally are built and operated by service providers, yet enterprises whose business model is based on Internet commerce also build and operate IDCs. The architecture of enterprise IDCs is very similar to that of the service provider IDCs, but the requirements for scalability are typically lower because the user base tends to be smaller and there are fewer services compared with those of SP IDCs hosting multiple customers.

In fact, the architecture of the IDC is the same as that presented in Figure 4-2. An interesting consideration of enterprise IDCs is that if the business model calls for it, the facilities used by the Data Center could be collocated in a service provider Data Center, but it remains under the control of the enterprise. This typically is done to lower the costs associated with building the server farm and reducing a product's time to market by avoiding building a Data Center internally from the ground up.

Corporate Data Center

Corporate or enterprise Data Centers support many different functions that enable various business models based on Internet services, intranet services, or both. As a result, support for Internet, intranet, and extranet server farms is not uncommon. This concept was depicted in Figure 4-1, where the Data Center facility supports every type of server farm and also is connected to the rest of the enterprise network—private WAN, campus, Internet Edge, and so on. The support of intranet server farms is still the primary target of corporate Data Centers.

Enterprise Data Centers are evolving, and this evolution is partly a result of new trends in application environments, such as the n-tier, web services, and grid computing, but it results mainly because of the criticality of the data held in Data Centers.

The following section discusses the typical topologies used in the architecture of the Data Center.

Data Center Topologies

This section discusses Data Center topologies and, in particular, the server farm topology. Initially, the discussion focuses on the traffic flow through the network infrastructure (on a generic topology) from a logical viewpoint and then from a physical viewpoint.

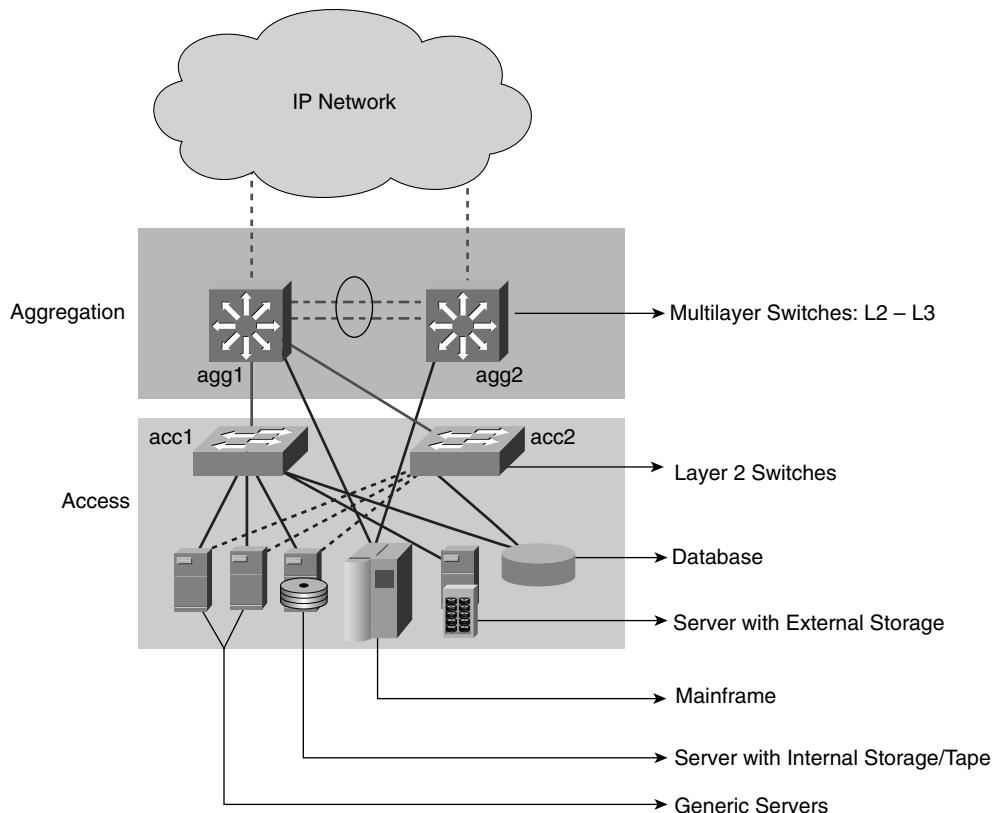
Generic Layer 3/Layer 2 Designs

The generic Layer 3/Layer 2 designs are based on the most common ways of deploying server farms. Figure 4-6 depicts a generic server farm topology that supports a number of servers.

NOTE

Notice that the distribution layer now is referred to as the aggregation layer resulting from becoming the aggregation point for most, if not all, services beyond the traditional Layer 2 and Layer 3.

Figure 4-6 Generic Server Farm Design



The highlights of the topology are the aggregation-layer switches that perform key Layer 3 and Layer 2 functions, the access-layer switches that provide connectivity to the servers in the server farm, and the connectivity between the aggregation and access layer switches.

The key Layer 3 functions performed by the aggregation switches are as follows:

- Forwarding packets based on Layer 3 information between the server farm and the rest of the network
- Maintaining a “view” of the routed network that is expected to change dynamically as network changes take place
- Supporting default gateways for the server farms

The key Layer 2 functions performed by the aggregation switches are as follows:

- Spanning Tree Protocol (STP) 802.1d between aggregation and access switches to build a loop-free forwarding topology.
- STP enhancements beyond 802.1d that improve the default spanning-tree behavior, such as 802.1s, 802.1w, Uplinkfast, Backbonefast, and Loopguard. For more information, refer to Chapter 12, “Layer 2 Protocol Essentials.”
- VLANs for logical separation of server farms.
- Other services, such as multicast and ACLs for services such as QoS, security, rate limiting, broadcast suppression, and so on.

The access-layer switches provide direct connectivity to the server farm. The types of servers in the server farm include generic servers such as DNS, DHCP, FTP, and Telnet; mainframes using SNA over IP or IP; and database servers. Notice that some servers have both internal disks (storage) and tape units, and others have the storage externally connected (typically SCSI).

The connectivity between the two aggregation switches and between aggregation and access switches is as follows:

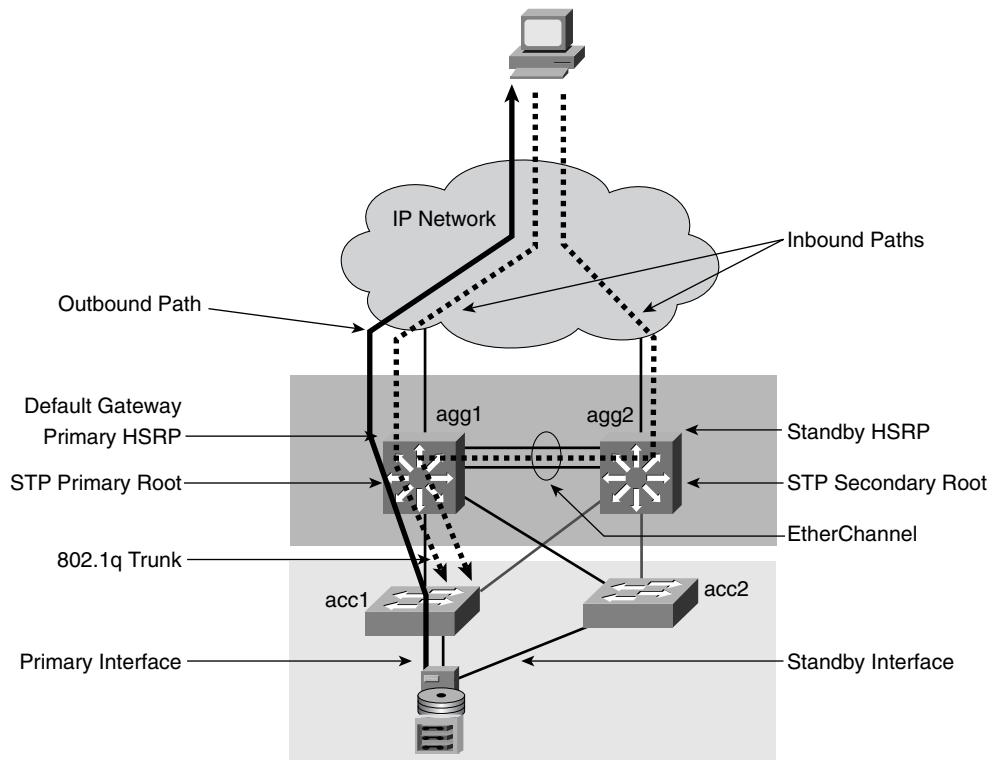
- EtherChannel between aggregation switches. The channel is in trunk mode, which allows the physical links to support as many VLANs as needed (limited to 4096 VLANs resulting from the 12-bit VLAN ID).
- Single or multiple links (EtherChannel, depending on how much oversubscription is expected in the links) from each access switch to each aggregation switch (uplinks). These links are also trunks, thus allowing multiple VLANs through a single physical path.
- Servers dual-homed to different access switches for redundancy. The NIC used by the server is presumed to have two ports in an active-standby configuration. When the primary port fails, the standby takes over, utilizing the same MAC and IP addresses that the active port was using. For more information about dual-homed servers, refer to Chapter 2, “Server Architecture Overview.”

The typical configuration for the server farm environment just described is presented in Figure 4-7.

Figure 4-7 shows the location for the critical services required by the server farm. These services are explicitly configured as follows:

- agg1 is explicitly configured as the STP root.
- agg2 is explicitly configured as the secondary root.
- agg1 is explicitly configured as the primary default gateway.
- agg2 is explicitly configured as the standby or secondary default gateway.

Figure 4-7 Common Server Farm Environment



NOTE

The explicit definition of these critical functions sets the primary and alternate paths to and from the server farm. Notice that there is no single point of failure in the architecture, and the paths are now deterministic.

Other STP services or protocols, such as UplinkFast, are also explicitly defined between the aggregation and access layers. These services/protocols are used to lower convergence time during failover conditions from the 802.d standard of roughly 50 seconds to 1 to 3 seconds.

In this topology, the servers are configured to use the agg1 switch as the primary default gateway, which means that outbound traffic from the servers follows the direct path to the agg1 switch. Inbound traffic can arrive at either aggregation switch, yet the traffic can reach

the server farm only through agg1 because the links from agg2 to the access switches are not forwarding (blocking). The inbound paths are represented by the dotted arrows, and the outbound path is represented by the solid arrow.

The next step is to have predictable failover and fallback behavior, which is much simpler when you have deterministic primary and alternate paths. This is achieved by failing every component in the primary path and recording and tuning the failover time to the backup component until the requirements are satisfied. The same process must be done for falling back to the original primary device. This is because the failover and fallback processes are not the same. In certain instances, the fallback can be done manually instead of automatically, to prevent certain undesirable conditions.

NOTE

When using 802.1d, if the primary STP root fails and the secondary takes over, when it comes back up, it automatically takes over because it has a lower priority. In an active server farm environment, you might not want to have the STP topology change automatically, particularly when the convergence time is in the range of 50 seconds. However, this behavior is not applicable when using 802.1w, in which the fallback process takes only a few seconds.

Whether using 802.1d or 802.1w, the process is automatic, unlike when using HSRP, in which the user can control the behavior of the primary HSRP peer when it becomes operational again through the use of preemption. If preemption is not used, the user has manual control over when to return mastership to the initial master HSRP peer.

The use of STP is the result of a Layer 2 topology, which might have loops that require an automatic mechanism to be detected and avoided. An important question is whether there is a need for Layer 2 in a server farm environment. This topic is discussed in the following section.

For more information about the details of the Layer 2 design, see Chapter 20, “Designing the Data Center Infrastructure.”

The Need for Layer 2 at the Access Layer

Access switches traditionally have been Layer 2 switches. This holds true also for the campus network wiring closet. This discussion is focused strictly on the Data Center because it has distinct and specific requirements, some similar to and some different than those for the wiring closets.

The reason access switches in the Data Center traditionally have been Layer 2 is the result of the following requirements:

- When they share specific properties, servers typically are grouped on the same VLAN. These properties could be as simple as ownership by the same department or performance of the same function (file and print services, FTP, and so on). Some servers that perform the same function might need to communicate with one another, whether as a result of a clustering protocol or simply as part of the application function. This communication exchange should be on the same subnet and sometimes is possible only on the same subnet if the clustering protocol heartbeats or the server-to-server application packets are not routable.
- Servers are typically dual-homed so that each leg connects to a different access switch for redundancy. If the adapter in use has a standby interface that uses the same MAC and IP addresses after a failure, the active and standby interfaces must be on the same VLAN (same default gateway).
- Server farm growth occurs horizontally, which means that new servers are added to the same VLANs or IP subnets where other servers that perform the same functions are located. If the Layer 2 switches hosting the servers run out of ports, the same VLANs or subnets must be supported on a new set of Layer 2 switches. This allows flexibility in growth and prevents having to connect two access switches.
- When using stateful devices that provide services to the server farms, such as load balancers and firewalls, these stateful devices expect to see both the inbound and outbound traffic use the same path. They also need to constantly exchange connection and session state information, which requires Layer 2 adjacency. More details on these requirements are discussed in the section, “Access Layer,” which is under the section, “Multiple Tier Designs.”

Using just Layer 3 at the access layer would prevent dual-homing, Layer 2 adjacency between servers on different access switches, and Layer 2 adjacency between service devices. Yet if these requirements are not common on your server farm, you could consider a Layer 3 environment in the access layer. Before you decide what is best, it is important that you read the section titled “Fully Redundant Layer 2 and Layer 3 Designs with Services,” later in the chapter. New service trends impose a new set of requirements in the architecture that must be considered before deciding which strategy works best for your Data Center.

The reasons for migrating away from a Layer 2 access switch design are motivated by the need to drift away from spanning tree because of the slow convergence time and the operation challenges of running a controlled loopless topology and troubleshooting loops when they occur. Although this is true when using 802.1d, environments that take advantage of 802.1w combined with Loopguard have the following characteristics: They do not suffer from the same problems, they are as stable as Layer 3 environments, and they support low convergence times.

NOTE

The STP standard 802.1d has limitations in addressing certain conditions in addition to its convergence time, yet a fair amount of spanning tree-related problems are the result of misconfiguration or rogue STP devices that appear on the network and “bridge” between Layer 2 domains. More information on this topic is presented in Chapter 12.

The next section discusses an alternate solution for a topology with spanning tree that does not present the STP problems or limitations.

Alternate Layer 3/Layer 2 Designs

Figure 4-8 presents an alternate Layer 3/Layer 2 design resulting from the need to address STP limitations.

Figure 4-8 Loopless Topology

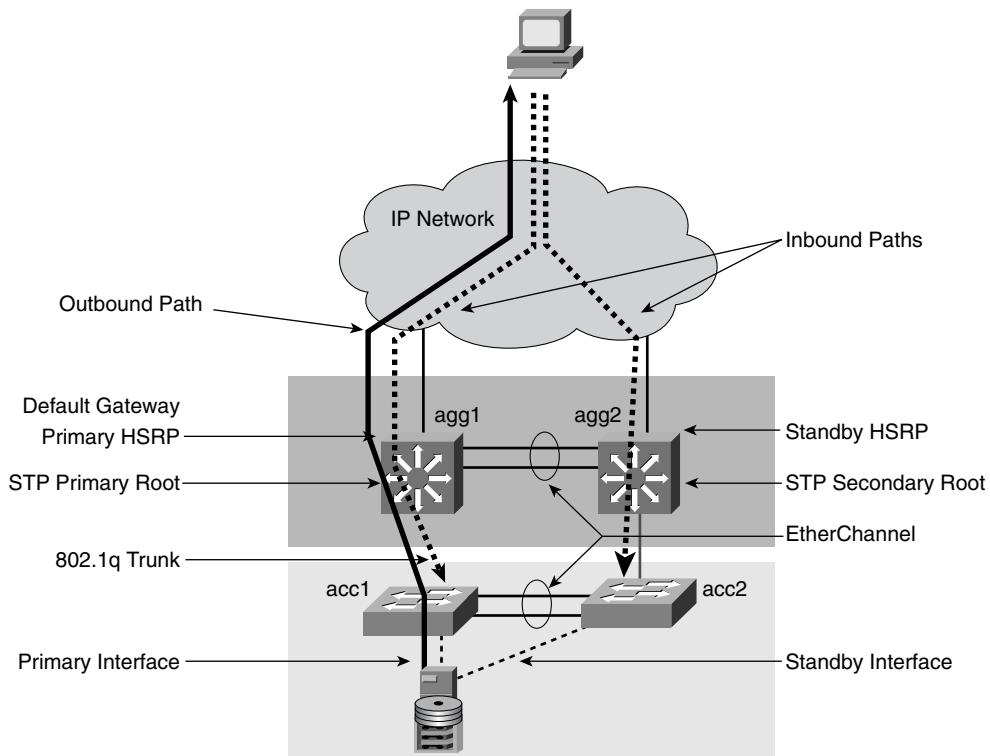


Figure 4-8 presents a topology in which the network purposely is designed not to have loops. Although STP is running, its limitations do not present a problem. This loopless topology is accomplished by removing or not allowing the VLAN(s), used at the access-layer switches, through the trunk between the two aggregation switches. This basically prevents a loop in the topology while it supports the requirements behind the need for Layer 2.

In this topology, the servers are configured to use the agg1 switch as the primary default gateway. This means that outbound traffic from the servers connected to acc2 traverses the link between the two access switches. Inbound traffic can use either aggregation switch because both have active (nonblocking) paths to the access switches. The inbound paths are represented by the dotted arrows, and the outbound path is represented by the solid arrows.

This topology is not without its own challenges. These challenges are discussed later in the chapter after other information related to the deployment of services becomes available.

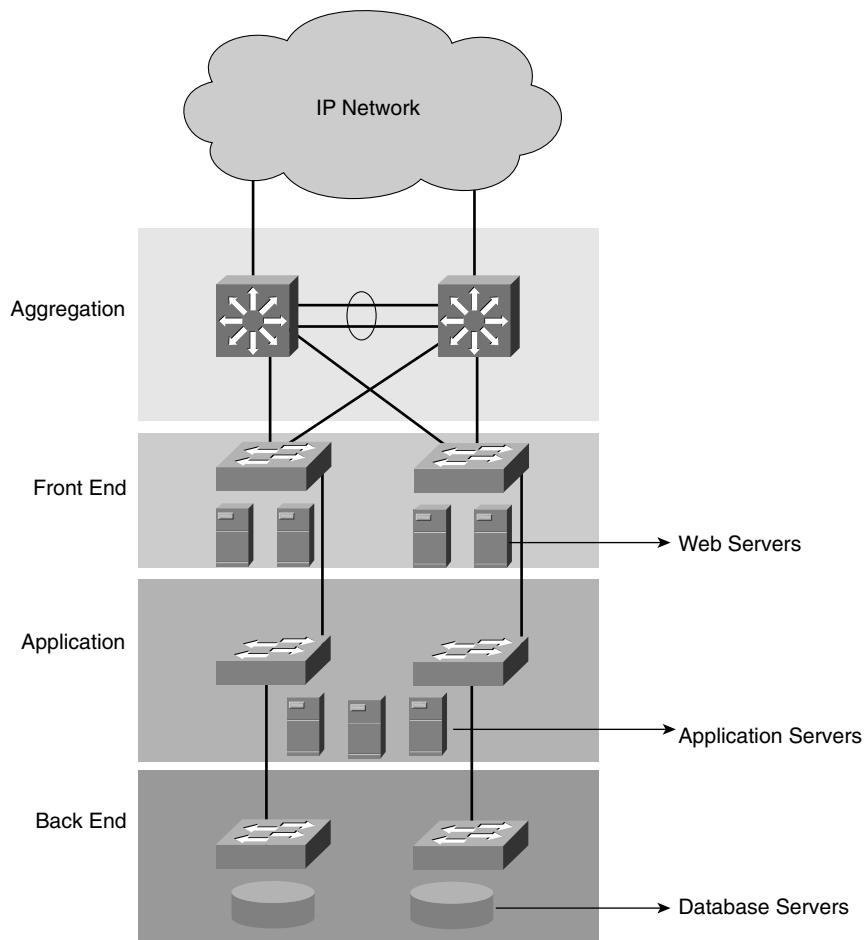
Multiple-Tier Designs

Most applications conform to either the client/server model or the n-tier model, which implies most networks, and server farms support these application environments. The tiers supported by the Data Center infrastructure are driven by the specific applications and could be any combination in the spectrum of applications from the client/server to the client/web server/application server/database server. When you identify the communication requirements between tiers, you can determine the needed specific network services. The communication requirements between tiers are typically higher scalability, performance, and security. These could translate to load balancing between tiers for scalability and performance, or SSL between tiers for encrypted transactions, or simply firewalling and intrusion detection between the web and application tier for more security.

Figure 4-9 introduces a topology that helps illustrate the previous discussion.

Notice that Figure 4-9 is a logical diagram that depicts layer-to-layer connectivity through the network infrastructure. This implies that the actual physical topology might be different. The separation between layers simply shows that the different server functions could be physically separated. The physical separation could be a design preference or the result of specific requirements that address communication between tiers.

For example, when dealing with web servers, the most common problem is scaling the web tier to serve many concurrent users. This translates into deploying more web servers that have similar characteristics and the same content so that user requests can be equally fulfilled by any of them. This, in turn, requires the use of a load balancer in front of the server farm that hides the number of servers and virtualizes their services. To the users, the specific service is still supported on a single server, yet the load balancer dynamically picks a server to fulfill the request.

Figure 4-9 Multiple-Tier Application Environments

Suppose that you have multiple types of web servers supporting different applications, and some of these applications follow the n-tier model. The server farm could be partitioned along the lines of applications or functions. All web servers, regardless of the application(s) they support, could be part of the same server farm on the same subnet, and the application servers could be part of a separate server farm on a different subnet and different VLAN.

Following the same logic used to scale the web tier, a load balancer logically could be placed between the web tier and the application tier to scale the application tier from the web tier perspective. A single web server now has multiple application servers to access.

The same set of arguments holds true for the need for security at the web tier and a separate set of security considerations at the application tier. This implies that firewall and intrusion-detection capabilities are distinct at each layer and, therefore, are customized for the requirements of the application and the database tiers. SSL offloading is another example of a function that the server farm infrastructure might support and can be deployed at the web tier, the application tier, and the database tier. However, its use depends upon the application environment using SSL to encrypt client-to-server and server-to-server traffic.

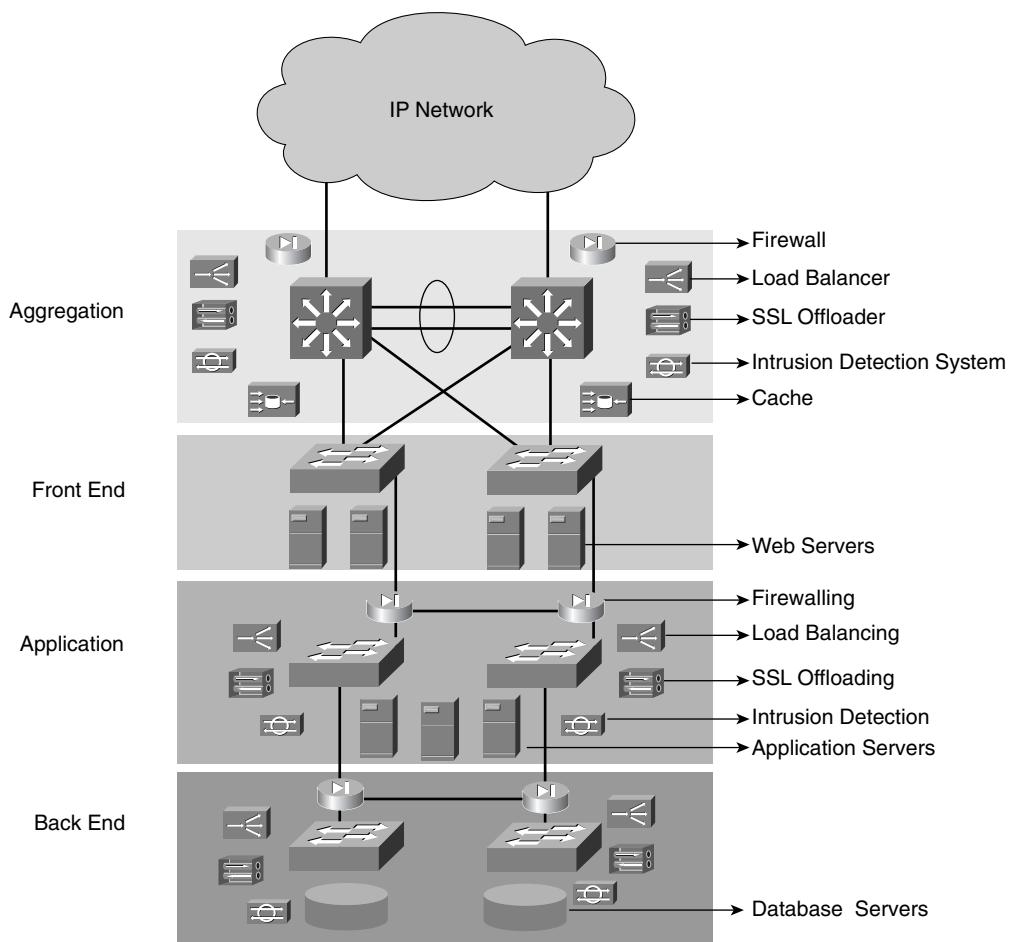
Expanded Multitier Design

The previous discussion leads to the concept of deploying multiple network-based services in the architecture. These services are introduced in Figure 4-10 through the use of icons that depict the function or service performed by the network device.

NOTE

Figure 4-10 introduces the icons used through this chapter to depict the services provided by network devices in the Data Center.

The different icons are placed in front of the servers for which they perform the functions. At the aggregation layer, you find the load balancer, firewall, SSL offloader, intrusion-detection system, and cache. These services are available through service modules (line cards that could be inserted into the aggregation switch) or appliances. An important point to consider when dealing with service devices is that they provide scalability and high availability beyond the capacity of the server farm, and that to maintain the basic premise of “no single point of failure,” at least two must be deployed. If you have more than one (and considering you are dealing with redundancy of application environments), the failover and fallback processes require special mechanisms to recover the connection context, in addition to the Layer 2 and Layer 3 paths. This simple concept of redundancy at the application layer has profound implications in the network design.

Figure 4-10 Network Service Icons

A number of these network service devices are replicated in front of the application layer to provide services to the application servers. Notice in Figure 4-10 that there is physical separation between the tiers of servers. This separation is one alternative to the server farm design. Physical separation is used to achieve greater control over the deployment and scalability of services. The expanded design is more costly because it uses more devices, yet it allows for more control and better scalability because the devices in the path handle only a portion of the traffic. For example, placing a firewall between tiers is regarded as a more secure approach because of the physical separation between the Layer 2 switches.

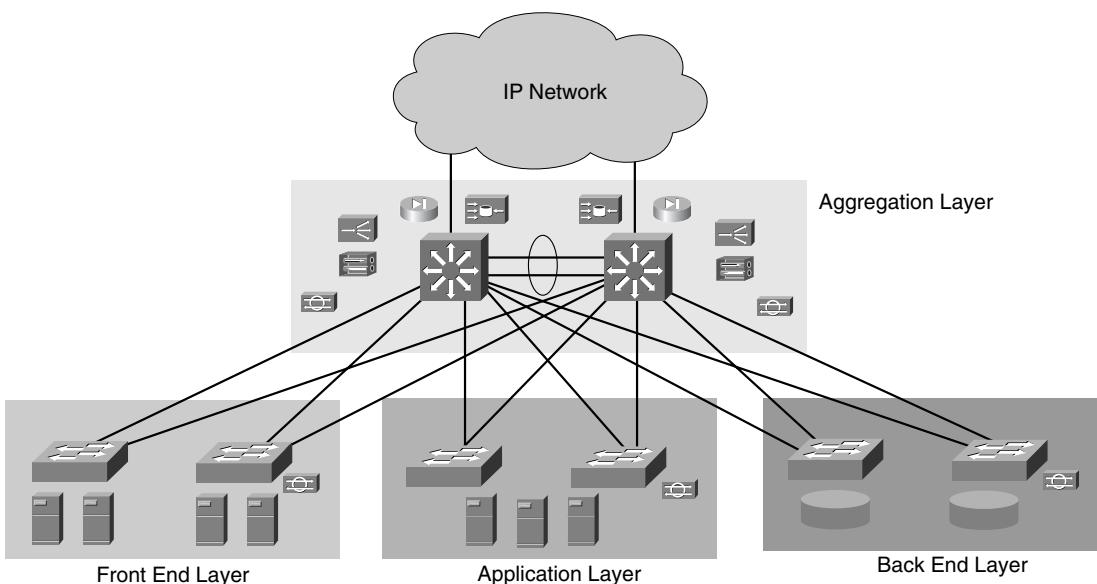
This argument is correct, yet it is likely to be much more related to an existing security policy than a real threat. Having logical instead of physical separation simply requires a consistent application of security policies to ensure that the expanded security zone is as secure logically as it is physically.

This brings the discussion to another alternative of designing the multitier server farm, an alternative in which there is no physical separation, but rather a logical separation between tiers, as presented in the next section.

Collapsed Multitier Design

A collapsed multitier design is one in which all the server farms are directly connected at the access layer to the aggregation switches, and there is no physical separation between the Layer 2 switches that support the different tiers. Figure 4-11 presents the collapsed design.

Figure 4-11 Collapsed Multiple-Tier Design



Notice that in this design, the services again are concentrated at the aggregation layer, and the service devices now are used by the front-end tier and between tiers. Using a collapsed model, there is no need to have a set of load balancers or SSL offloaders dedicated to a particular tier. This reduces cost, yet the management of devices is more challenging and the performance demands are higher. The service devices, such as the firewalls, protect all

server tiers from outside the Data Center, but also from each other. The load balancer also can be used concurrently to load-balance traffic from client to web servers, and traffic from web servers to application servers.

Notice that the design in Figure 4-11 shows each type of server farm on a different set of switches. Other collapsed designs might combine the same physical Layer 2 switches to house web applications and database servers concurrently. This implies merely that the servers logically are located on different IP subnets and VLANs, yet the service devices still are used concurrently for the front end and between tiers. Notice that the service devices are always in pairs. Pairing avoids the single point of failure throughout the architecture. However, both service devices in the pair communicate with each other, which falls into the discussion of whether you need Layer 2 or Layer 3 at the access layer.

The Need for Layer 2 at the Access Layer

Each pair of service devices must maintain state information about the connections the pair is handling. This requires a mechanism to determine the active device (master) and another mechanism to exchange connection state information on a regular basis. The goal of the dual-service device configuration is to ensure that, upon failure, the redundant device not only can continue service without interruption, but also seamlessly can failover without disrupting the current established connections.

In addition to the requirements brought up earlier about the need for Layer 2, this section discusses in depth the set of requirements related to the service devices:

- Service devices and the server farms that they serve are typically Layer 2-adjacent. This means that the service device has a leg sitting on the same subnet and VLAN used by the servers, which is used to communicate directly with them. Often, in fact, the service devices themselves provide default gateway support for the server farm.
- Service devices must exchange heartbeats as part of their redundancy protocol. The heartbeat packets might or might not be routable; if they are routable, you might not want the exchange to go through unnecessary Layer 3 hops.
- Service devices operating in stateful failover need to exchange connection and session state information. For the most part, this exchange is done over a VLAN common to the two devices. Much like the heartbeat packets, they might or might not be routable.
- If the service devices provide default gateway support for the server farm, they must be adjacent to the servers.

After considering all the requirements for Layer 2 at the access layer, it is important to note that although it is possible to have topologies such as the one presented in Figure 4-8, which supports Layer 2 in the access layer, the topology depicted in Figure 4-7 is preferred. Topologies with loops are also supportable if they take advantages of protocols such as 802.1w and features such as Loopguard.

NOTE

To date, most common implementations use Layer 2 at the access layer and rely on the Spanning Tree Protocols and Cisco enhancements to lower convergence times and achieve stability, as depicted in Figure 4-7. Few use the loopless topology. The main reasons relate to whether it is possible to have a loopless topology, given the restrictions imposed by the requirements, and, if possible, whether the setup is simple enough for support, maintenance, and management reasons. Dual-homing requires Layer 2 adjacency between access switches to carry the same VLANs, and redundant stateful service devices need Layer 2 adjacency to work properly. Therefore, it is important to carefully consider the requirements when designing the server farm network infrastructure.

The following section discusses topics related to the topology of the server farms.

Fully Redundant Layer 2 and Layer 3 Designs

Up to this point, all the topologies that have been presented are fully redundant. This section explains the various aspects of a redundant and scalable Data Center design by presenting multiple possible design alternatives, highlighting sound practices, and pointing out practices to be avoided.

The Need for Redundancy

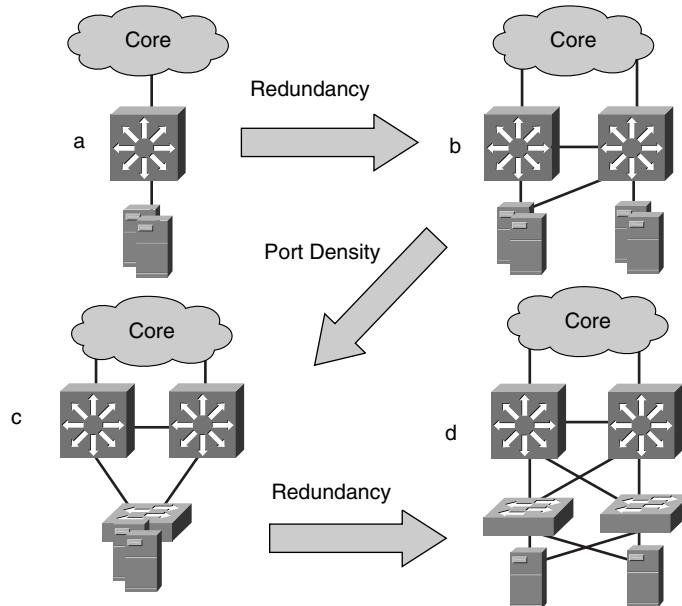
Figure 4-12 explains the steps in building a redundant topology.

Figure 4-12 depicts the logical steps in designing the server farm infrastructure. The process starts with a Layer 3 switch that provides ports for direct server connectivity and routing to the core. A Layer 2 switch could be used, but the Layer 3 switch limits the broadcasts and flooding to and from the server farms. This is option **a** in Figure 4-12. The main problem with the design labeled **a** is that there are multiple single point of failure problems: There is a single NIC and a single switch, and if the NIC or switch fails, the server and applications become unavailable.

The solution is twofold:

- Make the components of the single switch redundant, such as dual power supplies and dual supervisors.
- Add a second switch.

Redundant components make the single switch more tolerant, yet if the switch fails, the server farm is unavailable. Option **b** shows the next step, in which a redundant Layer 3 switch is added.

Figure 4-12 Multilayer Redundant Design

By having two Layer 3 switches and spreading servers on both of them, you achieve a higher level of redundancy in which the failure of one Layer 3 switch does not completely compromise the application environment. The environment is not completely compromised when the servers are dual-homed, so if one of the Layer 3 switches fails, the servers still can recover by using the connection to the second switch.

In options **a** and **b**, the port density is limited to the capacity of the two switches. As the demands for more ports increase for the server and other service devices, and when the maximum capacity has been reached, adding new ports becomes cumbersome, particularly when trying to maintain Layer 2 adjacency between servers.

The mechanism used to grow the server farm is presented in option **c**. You add Layer 2 access switches to the topology to provide direct server connectivity. Figure 4-12 depicts the Layer 2 switches connected to both Layer 3 aggregation switches. The two uplinks, one to each aggregation switch, provide redundancy from the access to the aggregation switches, giving the server farm an alternate path to reach the Layer 3 switches.

The design described in option **c** still has a problem. If the Layer 2 switch fails, the servers lose their only means of communication. The solution is to dual-home servers to two different Layer 2 switches, as depicted in option **d** of Figure 4-12.

NOTE

Throughout this book, the terms *access layer* and *access switches* refer to the switches used to provide port density. The terms *aggregation layer* and *aggregation switches* refer to the switches used both to aggregate the traffic to and from the access switches and to connect service devices (load balancers, SSL offloaders, firewalls, caches, and so on).

The *aggregation switches* are Layer 3 switches, which means that they have a built-in router that can forward traffic at wire speed.

The *access switches* are predominantly Layer 2 switches, yet they could be Layer 3 switches merely operating in Layer 2 mode for the server farms.

Layer 2 and Layer 3 in Access Layer

Option d in Figure 4-12 is detailed in option a of Figure 4-13.

Figure 4-13 Layer 3 and Layer 2 in the Data Center

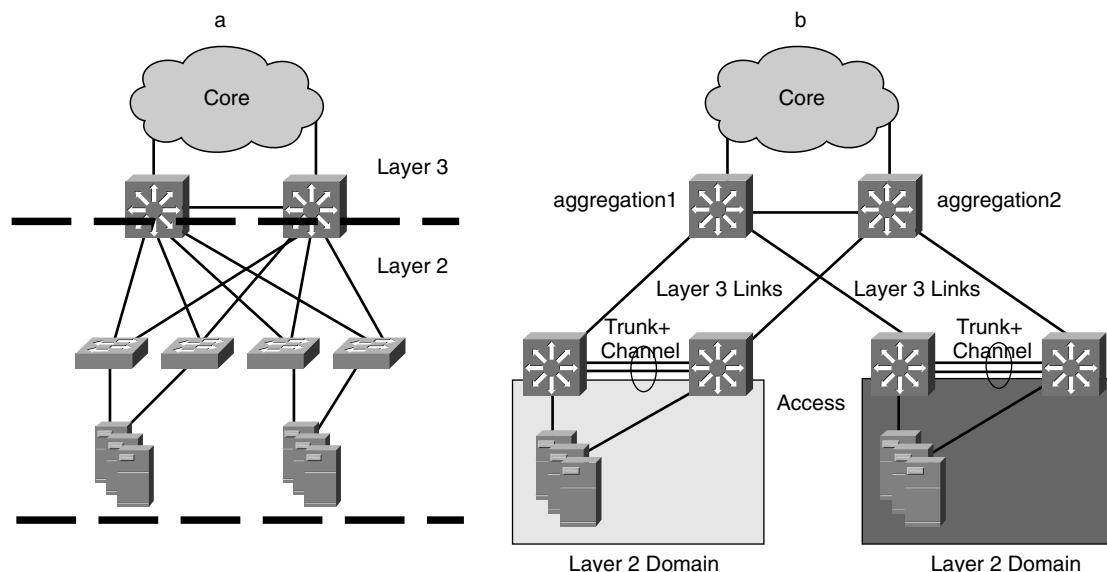


Figure 4-13 presents the scope of the Layer 2 domain(s) from the servers to the aggregation switches. Redundancy in the Layer 2 domain is achieved mainly by using spanning tree, whereas in Layer 3, redundancy is achieved through the use of routing protocols.

Historically, routing protocols have proven more stable than spanning tree, which makes one question the wisdom of using Layer 2 instead of Layer 3 at the access layer. This topic was discussed previously in the “Need for Layer 2 at the Access Layer” section. As shown

in option **b** in Figure 4-13, using Layer 2 at the access layer does not prevent the building of pure Layer 3 designs because of the routing between the access and distribution layer or the supporting Layer 2 between access switches.

The design depicted in option **a** of Figure 4-13 is the most generic design that provides redundancy, scalability, and flexibility. Flexibility relates to the fact that the design makes it easy to add service appliances at the aggregation layer with minimal changes to the rest of the design. A simpler design such as that depicted in option **b** of Figure 4-13 might better suit the requirements of a small server farm.

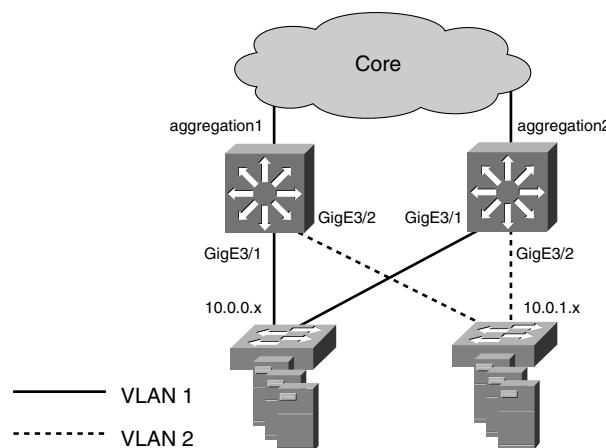
Layer 2, Loops, and Spanning Tree

The Layer 2 domains should make you think immediately of loops. Every network designer has experienced Layer 2 loops in the network. When Layer 2 loops occur, packets are replicated an infinite number of times, bringing down the network. Under normal conditions, the Spanning Tree Protocol keeps the logical topology free of loops. Unfortunately, physical failures such as unidirectional links, incorrect wiring, rogue bridging devices, or bugs can cause loops to occur.

Fortunately, the introduction of 802.1w has addressed many of the limitations of the original spanning tree algorithm, and features such as Loopguard fix the issue of malfunctioning transceivers or bugs.

Still, the experience of deploying legacy spanning tree drives network designers to try to design the Layer 2 topology free of loops. In the Data Center, this is sometimes possible. An example of this type of design is depicted in Figure 4-14. As you can see, the Layer 2 domain (VLAN) that hosts the subnet 10.0.0.x is not trunked between the two aggregation switches, and neither is 10.0.1.x. Notice that GigE3/1 and GigE3/2 are not bridged together.

Figure 4-14 Loop-Free Layer 2 Design



TIP

It is possible to build a loop-free access layer if you manage to keep subnets specific to a single access switch. If subnets must span multiple access switches, you should have a “looped” topology. This is the case when you have dual-attached servers because NIC cards configured for “teaming” typically use a floating IP and MAC address, which means that both interfaces belong to the same subnet.

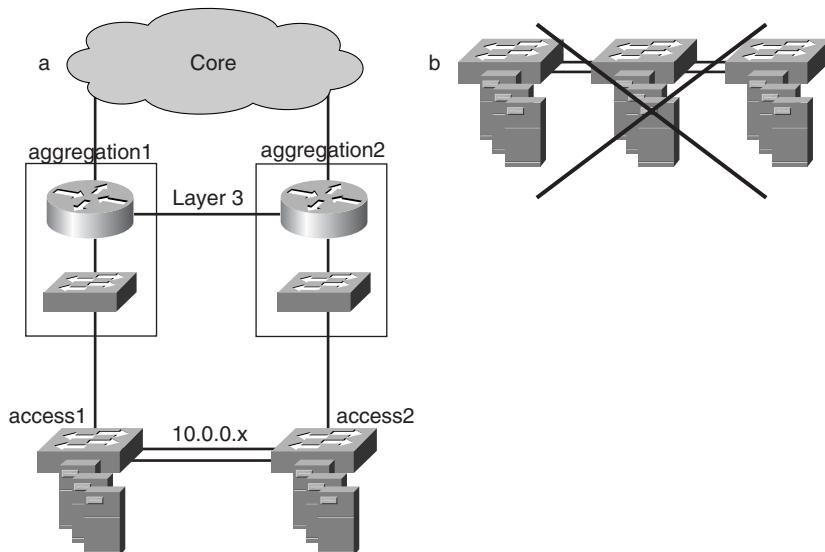
Keep in mind that a “loop-free” topology is not necessarily better. Specific requirements such as those mandated by content switches actually might require the additional path provided by a “looped” topology.

Also notice that a “looped” topology simply means that any Layer 2 device can reach any other Layer 2 device from at least two different physical paths. This does not mean that you have a “forwarding loop,” in which packets are replicated infinite times: Spanning tree prevents this from happening.

In a “looped” topology, malfunctioning switches can cause Layer 2 loops. In a loop-free topology, there is no chance for a Layer 2 loop because there are no redundant Layer 2 paths.

If the number of ports must increase for any reason (dual-attached servers, more servers, and so forth), you could follow the approach of daisy-chaining Layer 2 switches, as shown in Figure 4-15.

Figure 4-15 Alternate Loop-Free Layer 2 Design



To help you visualize a Layer 2 loop-free topology, Figure 4-15 shows each aggregation switch broken up as a router and a Layer 2 switch.

The problem with topology **a** is that breaking the links between the two access switches would create a discontinuous subnet—this problem can be fixed with an EtherChannel between the access switches.

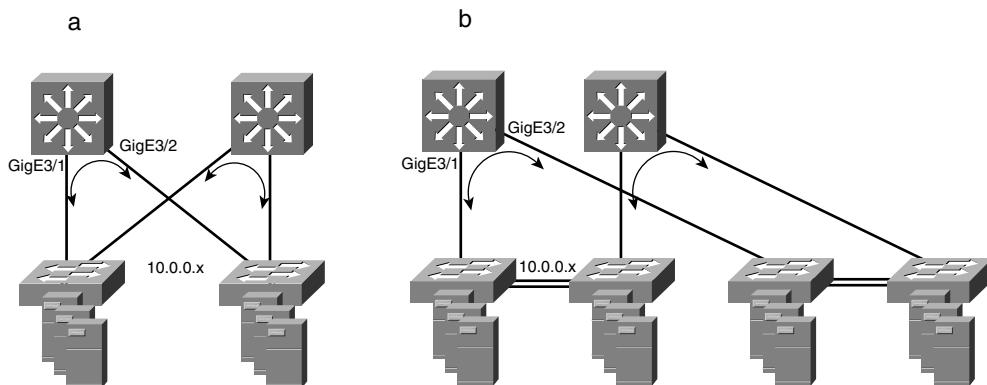
The other problem occurs when there are not enough ports for servers. If a number of servers need to be inserted into the same subnet 10.0.0.x, you cannot add a switch between the two existing servers, as presented in option **b** of Figure 4-15. This is because there is no workaround to the failure of the middle switch, which would create a split subnet. This design is not intrinsically wrong, but it is not optimal.

Both the topologies depicted in Figures 4-14 and 4-15 should migrate to a looped topology as soon as you have any of the following requirements:

- An increase in the number of servers on a given subnet
- Dual-attached NIC cards
- The spread of existing servers for a given subnet on a number of different access switches
- The insertion of stateful network service devices (such as load balancers) that operate in active/standby mode

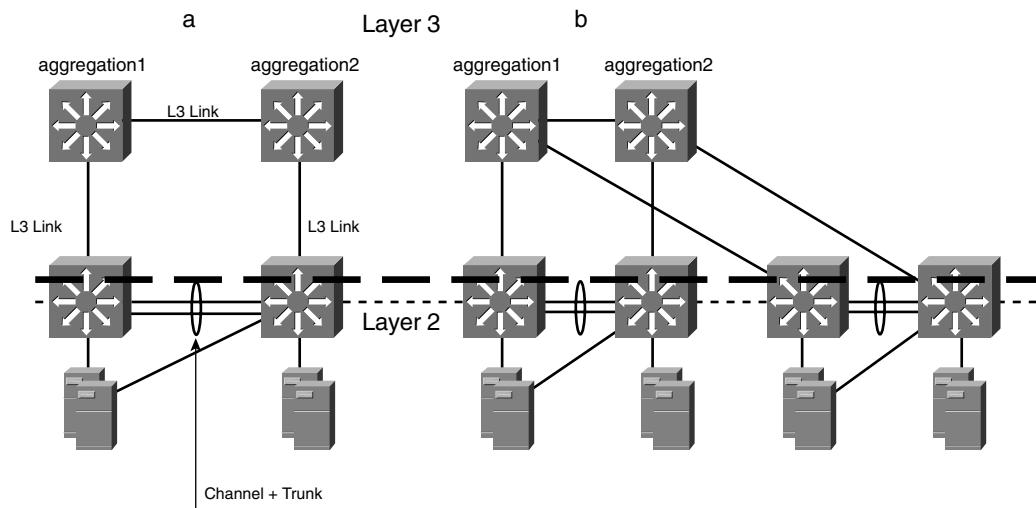
Options **a** and **b** in Figure 4-16 show how introducing additional access switches on the existing subnet creates “looped topologies.” In both **a** and **b**, GigE3/1 and GigE3/2 are bridged together.

Figure 4-16 Redundant Topologies with Physical Layer 2 Loops



If the requirement is to implement a topology that brings Layer 3 to the access layer, the topology that addresses the requirements of dual-attached servers is pictured in Figure 4-17.

Figure 4-17 Redundant Topology with Layer 3 to the Access Switches



Notice in option **a** of Figure 4-17, almost all the links are Layer 3 links, whereas the access switches have a trunk (on a channel) to provide the same subnet on two different switches. This trunk also carries a Layer 3 VLAN, which basically is used merely to make the two switches neighbors from a routing point of view. The dashed line in Figure 4-17 shows the scope of the Layer 2 domain.

Option **b** in Figure 4-17 shows how to grow the size of the server farm with this type of design. Notice that when deploying pairs of access switches, each pair has a set of subnets disjointed from the subnets of any other pair. For example, one pair of access switches hosts subnets 10.0.1.x and 10.0.2.x; the other pair cannot host the same subnets simply because it connects to the aggregation layer with Layer 3 links.

NOTE

If you compare the design in Figure 4-17 with option **b** in Figure 4-12, the natural questions are these: Why is there an aggregation layer, and are the access switches not directly connected to the core? These are valid points, and the answer actually depends on the size of the Data Center. Remember that the access layer is added for reasons of port density, whereas the aggregation layer is used mainly to attach appliances, such as load-balancing devices, firewalls, caches, and so on.

So far, the discussions have centered on redundant Layer 2 and Layer 3 designs. The Layer 3 switch provides the default gateway for the server farms in all the topologies introduced thus far. Default gateway support, however, could also be provided by other service devices, such as load balancers and firewalls. The next section explores the alternatives.

Fully Redundant Layer 2 and Layer 3 Designs with Services

After discussing the build-out of a fully redundant Layer 2 and Layer 3 topology and considering the foundation of the Data Center, the focus becomes the design issues related to other Data Center services. These services are aimed at improving security and scaling the performance of application services by offloading processing away from the server farm to the network. These services include security, load balancing, SSL offloading, and caching; they are supported by a number of networking devices that must be integrated into the infrastructure following the design requirements.

Additionally, this section discusses application environment trends brought about by technology advancements in either applications, the application infrastructure, or the network infrastructure.

Additional Services

At the aggregation layer, in addition to Layer 2 and Layer 3, the Data Center might need to support the following devices:

- Firewalls
- Intrusion Detection Systems (IDSs)
- Load balancers
- SSL offloaders
- Caches

It is important to discuss design issues when supporting some of these devices.

Service devices bring their own requirements that could change certain aspects of the design—for instance, the exchange state or status information, the NAT function that they perform on the source or destination IP addresses that forces them to be in the inbound and outbound path, and so on.

Service devices can be deployed using service modules integrated in the aggregation switches or as appliances connected to the aggregation switches. Both deployments require network connectivity and forethought about the actual traffic path.

Firewalls and load balancers may support the default gateway function on behalf of the server farms. Default gateway support traditionally has been provided by the router, so with two additional alternatives, you need to decide which is the default gateway and in which order traffic is processed through the multiple devices. Firewalls and load balancers are capable of providing stateful failover, which is supported by specific redundancy protocols. The protocols, which are specific to the firewalls or load balancers, must be supported by the design. SSL offloaders are typically used with load balancers and require the same considerations, with one exception: They do not support default gateway services.

IDSs are transparent to the design, which means that they integrate well with any existing design. The main consideration with regard to IDSs is their placement, which depends on selecting the location to analyze traffic and the traffic types to be monitored.

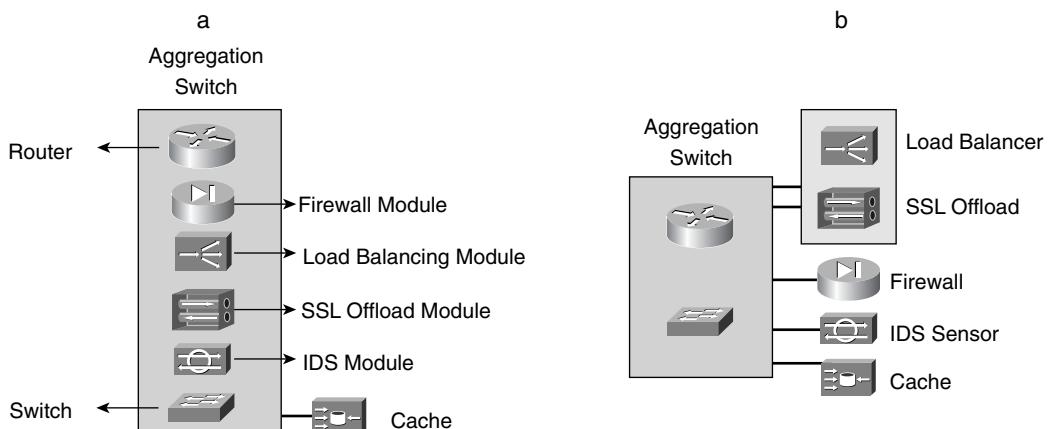
Caches, on the other hand, are deployed in reverse proxy cache mode. The placement of the caches and the mechanism for directing traffic to them impact the Data Center design. The options for traffic redirection are the Web Cache Communication Protocol (WCCP) on the Layer 2 or Layer 3 switches, and load balancers to distribute the load among the cache cluster. In either case, the cache or cache cluster changes the basic traffic path to the server farm when in use.

The following section presents the multiple deployment options.

Service Deployment Options

Two options exist when deploying Data Center services: using service modules integrated into the aggregation switch and using appliances connected to the aggregation switch. Figure 4-18 shows the two options.

Figure 4-18 Service Deployment Options



Option **a** shows the integrated design. The aggregation switch is represented by a router (Layer 3) and a switch (Layer 2) as the key components of the foundation (shown to the left) and by a firewall, load balancer, SSL module, and IDS module (shown to the right as add-on services). The service modules communicate with the routing and switching components in the chassis through the backplane.

Option **b** shows the appliance-based design. The aggregation switch provides the routing and switching functions. Other services are provided by appliances that are connected directly to the aggregation switches.

NOTE

Designs that use both modules and appliances are also possible. The most common case is when using caches, which are appliances, in both design options. Current trends on Data Center services lean toward integrated services. Evidence of this integration trend is the proliferation of services modules in the Catalyst 6500 family and the use of blade servers and blade chassis to collapse multiple services in one device.

A thoughtful approach to the design issues in selecting the traffic flow across different devices is required whether you are considering option **a**, option **b**, or any combination of the options in Figure 4-18. This means that you should explicitly select the default gateway and the order in which the packets from the client to the server are processed. The designs that use appliances require more care because you must be concerned with physical connectivity issues, interoperability, and the compatibility of protocols.

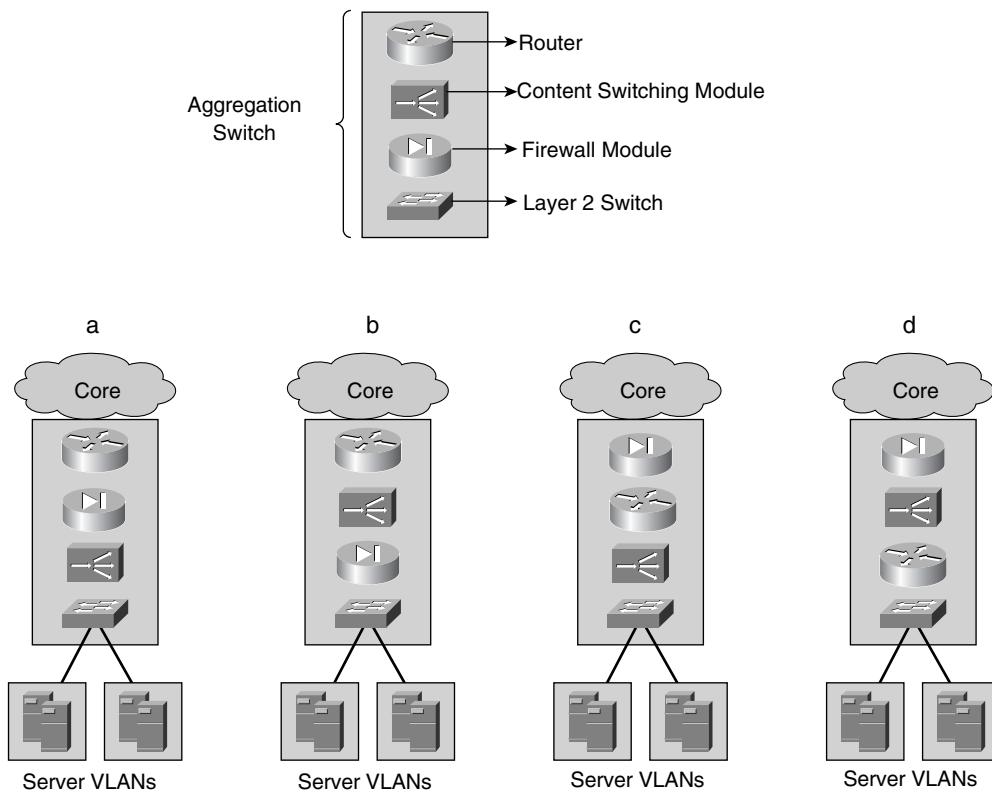
Design Considerations with Service Devices

Up to this point, several issues related to integrating service devices in the Data Center design have been mentioned. They are related to whether you run Layer 2 or Layer 3 at the access layer, whether you use appliance or modules, whether they are stateful or stateless, and whether they require you to change the default gateway location away from the router. Changing the default gateway location forces you to determine the order in which the packet needs to be processed through the aggregation switch and service devices.

Figure 4-19 presents the possible alternatives for default gateway support using service modules. The design implications of each alternative are discussed next.

Figure 4-19 shows the aggregation switch, a Catalyst 6500 using a firewall service module, and a content-switching module, in addition to the routing and switching functions provided by the Multilayer Switch Feature Card (MSFC) and the Supervisor Module.

The one constant factor in the design is the location of the switch providing server connectivity; it is adjacent to the server farm.

Figure 4-19 Service Module Interoperability Alternatives

Option **a** presents the router facing the core IP network, the content-switching module facing the server farm, and the firewall module between them firewalls all server farms. If the content switch operates as a router (route mode), it becomes the default gateway for the server farm. However, if it operates as a bridge (bridge mode), the default gateway would be the firewall. This configuration facilitates the creation of multiple instances of the firewall and content switch combination for the segregation and load balancing of each server farm independently.

Option **b** has the firewall facing the server farm and the content switch between the router and the firewall. Whether operating in router mode or bridge mode, the firewall configuration must enable server health-management (health probes) traffic from the content-switching module to the server farm; this adds management and configuration tasks to the design. Note that, in this design, the firewall provides the default gateway support for the server farm.

Option **c** shows the firewall facing the core IP network, the content switch facing the server farm, and the firewall module between the router and the content-switching module. Placing a firewall at the edge of the intranet server farms requires the firewall to have “router-like” routing capabilities, to ease the integration with the routed network while segregating all the server farms concurrently. This makes the capability to secure each server farm independently more difficult because the content switch and the router could route packets between the server farm without going through the firewall. Depending on whether the content-switching module operates in router or bridge mode, the default gateway could be the content switch or the router, respectively.

Option **d** displays the firewall module facing the core IP network, the router facing the server farm, and the content-switching module in between. This option presents some of the same challenges as option **c** in terms of the firewall supporting IGPs and the inability to segregate each server farm independently. The design, however, has one key advantage: The router is the default gateway for the server farm. Using the router as the default gateway allows the server farms to take advantage of some key protocols, such as HSRP, and features, such as HSRP tracking, QoS, the DHCP relay function, and so on, that are only available on routers.

All the previous design options are possible—some are more flexible, some are more secure, and some are more complex. The choice should be based on knowing the requirements as well as the advantages and restrictions of each. The different design issues associated with the viable options are discussed in the different chapters in Part V. Chapter 21, “Integrating Security into the Infrastructure,” addresses the network design in the context of firewalls.

Application Environment Trends

Undoubtedly, the most critical trends are those related to how applications are being developed and are expected to work on the network. These trends can be classified arbitrarily into two major areas:

- Application architectures
- Network infrastructure

Application Architecture Trends

Application architecture trends include the evolution of the classic client/server model to the more specialized n-tier model, web services, specific application architectures, the server and client software (operating systems), application clients, the server and client hardware, and middleware used to integrate distributed applications in heterogeneous environments.

The more visible trends of application architectures are the wide adoption of web technology in conjunction with the use of the n-tier model to functionally segment distinct server

types. Currently, web, application, and database servers are the basic types, yet they are combined in many ways (depending on the vendor of the application and how the buyer wants to implement it).

This functional partitioning demands that the network be smarter about securing and scaling the tiers independently. For instance, the n-tier model's web tier layer created the need for smaller and faster servers used to scale up the front-end function. This resulted in 1RU (rack unit) servers, which offer adequate performance for web servers at a low cost and minimal infrastructure requirements (power and rack space).

Web services are bringing a service-oriented approach to the use of different and distinct distributed applications that are accessible using standard messages over Internet protocols. Web services rely initially on the transport functions of the network and eventually on using the network as an extension to provide computing capacity to the distributed application environments by offloading tasks to network hardware.

NOTE

The World Wide Web consortium (W3C) defines a web service as “a software application identified by a URI, whose interfaces and binding are capable of being defined, described, and discovered by XML artifacts and supports direct interactions with other software applications using XML-based messages via Internet-based protocols.” For more information on web services and its architecture, consult the W3C at www.w3.org.

Grid computing is another trend that actually brings the applications and the network closer together by treating the servers as a network of CPUs in which the applications use the most available CPU on the network. Other trends related to grid computing include blade servers as an alternative to 1RU servers, to provide higher CPU density per RU, lower power consumption per server, and an additional benefit of lower cabling requirements. Blade servers are servers on blades (or modules) that are inserted into a chassis, much like network modules or line cards are inserted on a switch chassis. Using blade servers in blade chassis enables you to centralize the server-management functions (one chassis instead of however many servers are in the chassis), requires less cables (one set per chassis instead of one set per server), and provides higher computing and memory capacity per rack unit.

However, the blade server technology is still young, which explains the variety of flavors, architectures, connectivity options, and features.

An instance of middleware is the software used in the management and control of distributed CPUs in a grid of computers that can be 1RU or blade servers. This specific middleware virtualizes the use of CPUs so that the applications are given a CPU cycle from CPUs on the network instead of through the traditional manner.

Network Infrastructure Trends

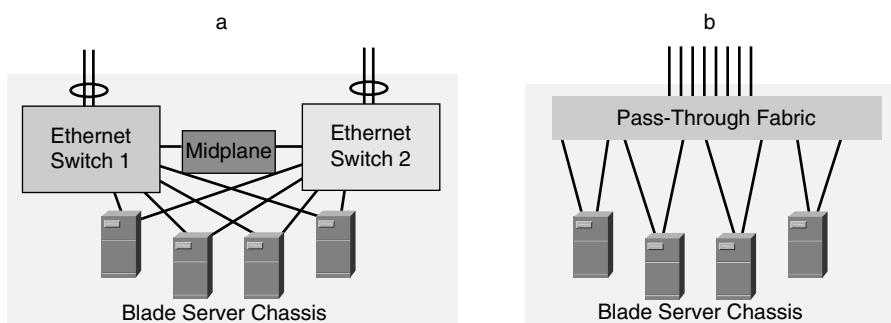
The network infrastructure is growing smarter and more application-aware, and it thereby supports application environments both by offloading some computationally intense tasks to the network (typically hardware-based) and by replacing some functions performed by servers that could be better handled by networking devices.

Load balancing is a good example of a function performed by the network that replaces clustering protocols used by servers for high availability. Clustering protocols tend to be software-based, hard to manage, and not very scalable in providing a function that the network performs well using hardware.

Trends such as blade servers bring new design considerations. Most blade server chassis (blade chassis, for short) in the market support both an option to provide redundant Ethernet switches inside the chassis and as an option to connect the blade servers to the network using pass-through links, with the chassis simply providing at least twice as many uplinks as servers in the chassis, to allow dual-homing.

Figure 4-20 presents both connectivity alternatives for a blade chassis.

Figure 4-20 *Blade Server Chassis Server Connectivity*



Option **a** in Figure 4-20 shows a blade server chassis in which each blade server is connected to each of the blade chassis's redundant Layer 2 Ethernet switches. Each blade chassis's Ethernet switch provides a number of uplinks that can be channeled to the IP network. The number of uplinks is typically smaller than the combined number of links per server, which requires planning for oversubscription, particularly if the servers are Gigabit Ethernet–attached. The midplane is the fabric used for management tasks, that is, control plane traffic such as switch status.

Option **b** in Figure 4-20 presents the pass-through option in which the servers are dual-homed and preconnected to a pass-through fabric that provides the connectivity to the IP network. This option does not use Ethernet switches inside the chassis. The pass-through fabric is as simple as a patch panel that conserves the properties of the server NICs, but it

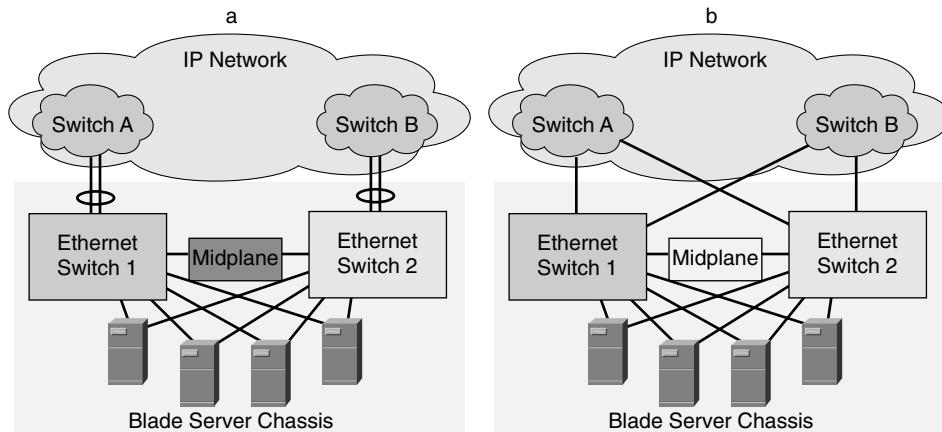
also could become a more intelligent fabric, adding new features and allowing blade server vendors to differentiate their products. Either approach you take to connect blade servers to your network requires careful consideration on short- and long-term design implications.

For instance, if the choice is to utilize the redundant Ethernet switches in the blade chassis, you have the following design alternatives to consider:

- How to use the redundant Ethernet switches' uplinks for connectivity
- Whether to connect the blade chassis to the access or aggregation switches
- What level of oversubscription is tolerable

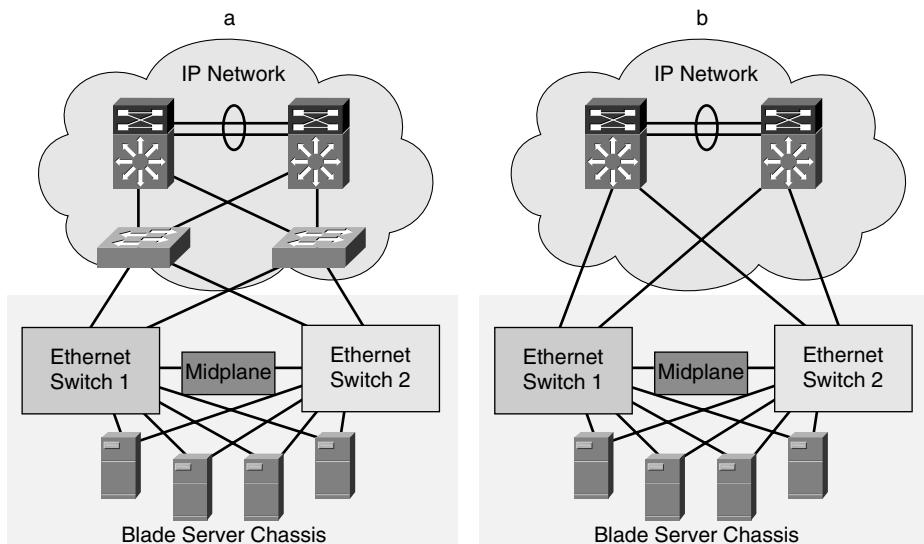
Figure 4-21 displays two connectivity choices utilizing the uplinks on the redundant Ethernet switches. For redundancy, two switches are used to connect the uplinks from the blade chassis. Switches A and B, the small clouds in the IP network cloud, provide a redundant network fabric to the blade chassis to avoid single point of failure issues.

Figure 4-21 Blade Chassis Uplink Connectivity



Option **a** in Figure 4-21 shows all the uplinks from both blade chassis' Ethernet switches connected to a single switch in the IP network. This allows the uplinks to be channeled. In contrast, option **b** in Figure 4-21 shows each blade chassis Ethernet switch connected to each IP network switch, also avoiding a single point of failure. This presents the advantage of having a direct link to either switch A or switch B, thus avoiding unnecessary hops. Additionally, if each blade chassis Ethernet switch supports more than two uplinks, they can also be channeled to switches A and B for greater redundancy and higher bandwidth.

The next step is to determine whether to connect the blade chassis to the access-layer switches, as is traditionally done with servers, or to the aggregation layer switches. Figure 4-22 displays the connectivity options for the next-hop switches from the blade chassis.

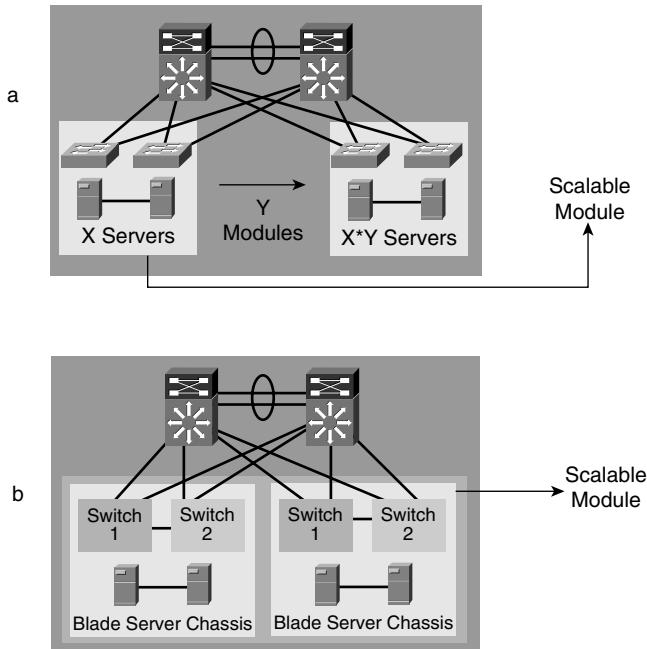
Figure 4-22 Blade Chassis Next-Hop Switch

Option **a** in Figure 4-22 shows the blade chassis connected to the access-layer switches. This particular design choice is equivalent to connecting Layer 2 access switches to Layer 2 access switches. The design must take into account spanning tree recommendations, which, based on the topology of option **a** in Figure 4-22, are aimed at determining a loop-free topology given the number of Layer 2 switches and the amount of available paths to the STP root and the secondary root from each leaf node. If the blade chassis Ethernet switches support 802.1w, the convergence time stays within two to three seconds; however, if the support is strictly 802.1d, the convergence time goes back to the typical range of 30 to 50 seconds. Other design considerations have to do with whether the midplane is used for more than management and switch-to-switch control traffic communication functions. If for some reason the midplane also is used to bridge VLANs (forward Bridge Protocol Data Units, or BPDUs) the STP topology needs to be considered carefully. The design goals remain making the topology predictable and deterministic. This implies that you need to explicitly set up root and bridge priorities and analyze the possible failure scenarios to make sure they support the requirements of the applications.

Option **b** in Figure 4-22 shows the blade chassis Ethernet switches directly connected to the aggregation switches. This is the preferred alternative because it lends itself to being more deterministic and supporting lower convergence times. Much like in the previous option, if the blade chassis Ethernet switches do not support 802.1w or some of the STP enhancements such as Uplinkfast and Loopguard, the convergence time would be in the range of 30 to 50 seconds. The topology still needs to be made deterministic and predictable by explicitly setting up root and bridge priorities and testing the failures scenarios.

How to scale the blade server farm is another consideration. Scalability on server environments is done simply by adding pairs of access switches for redundancy and connecting them to the aggregation switches, as shown in option **a** in Figure 4-23.

Figure 4-23 Server Farm Scalability



If a single scalable server module supports X servers (limited by port density), higher scalability is achieved by replicating the scalable module Y times (limited by slot density in the aggregation switch). The total number of servers could be $X * Y$. Depending on the access switch port density and the aggregation switch slot density, this could grow to thousands of servers. Scaling the number of blade servers might require a slightly different strategy. Because blade chassis with Ethernet switches are the access layer, the amount of blade server is limited to the number of slots and ports per slot at the aggregation switches. Option **b** in Figure 4-23 shows this alternative.

Notice that the scalable module is now the aggregation switch along with a set number of blade chassis. This is because the aggregation switch has a limit to the number of slots that can be used for blade chassis. In addition, line cards used to support blade server uplinks now receive aggregate server traffic, thus requiring less oversubscription. This leads to fewer ports used per line card. So, the total number of blade servers is limited somewhat by the slot and port density. Even though this design alternative is likely to support hundreds

of blade servers and satisfy the requirements for a fast-growing server farm environment, you must have a plan for what to do if you need to increase your server farm beyond what the current design supports. Figure 4-24 shows this alternative.

Figure 4-24 Core Layer Within the Data Center

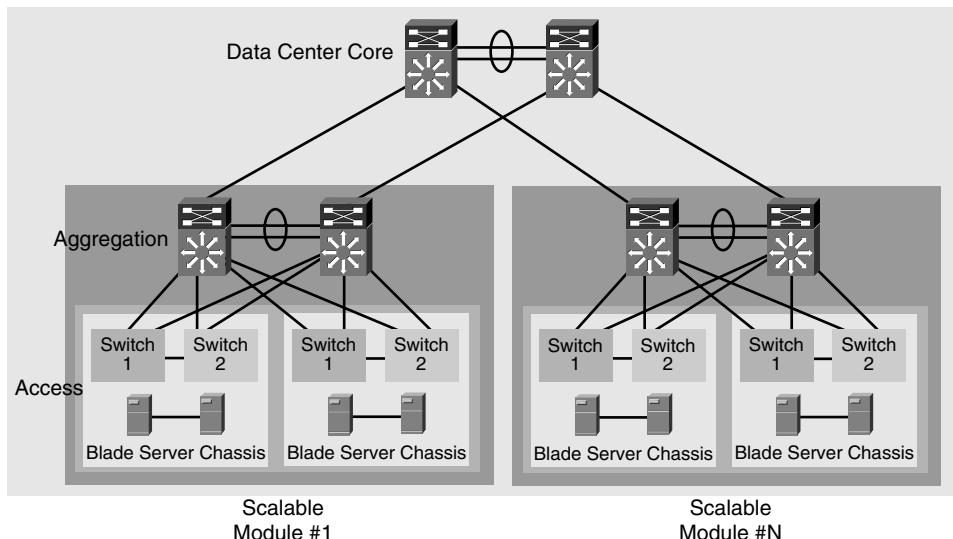


Figure 4-24 introduces a new layer in the Data Center: the core layer. The core layer is used to aggregate as many server blade modules as needed, but the number is limited to the port and slot capacity to the aggregation switches. The pass-through option might not require as much planning because the blade chassis do not have redundant Ethernet switches. The uplinks are connected to the access layer, which is equivalent to current designs in which servers are dual-homed to a redundant set of access switches.

Setting aside the connectivity, port density, slot density, and scalability considerations, other areas, such as oversubscription, uplink capacity, and service deployment options, might require design and testing before the Data Center architecture is established.

Additional trends include the dual-homing of servers, the migration from Fast Ethernet to Gigabit Ethernet, application firewalls, and the use of transparent network service devices. Application firewalls are firewalls that are more in tune with application behavior than ordinary firewalls, thus making the firewalling process more granular to application information in addition to just network or transport layer information. For instance, an application firewall might be capable of identifying not only that a packet is TCP and that the information in the TCP payload is HTTP, but also that the request comes from a specific high-priority user and is a SQL request for sensitive payroll information, which requires a higher security service level.

Transparent network services include firewalling, load balancing, SSL offloading, and so on. These services are provided by network devices with minimal interoperability issues that leave the existing designs unchanged. These transparent services could apply to traditional network services such as load balancing and firewalling, yet they are implemented to minimize disruption and changes in the application environment. This approach might include using physical devices as if they were distinct logical entities providing services to different server farms concurrently. This implies that the administration of those services, such as configuration changes or troubleshooting efforts, is isolated to the specific logical service. Think of it as a single physical firewall that is deployed to support many server farms concurrently where access to the CLI and configuration commands is available only to users who have been granted access to the specific server farm firewall service. This would appear to the user as a completely separate firewall.

Some of these trends are ongoing, and some are barely starting. Some will require special design and architectural considerations, and some will be adopted seamlessly. Others will not exist long enough for concern.

Summary

Data Centers are very dynamic environments hosting multiple types of server farms that all support key business applications. The design of the Data Center involves a variety of aspects related to how applications are architected, how they are deployed, and their network infrastructure.

A sound approach to design involves using a combination of architectural principles, such as scalability, flexibility, and high availability, as well as applying those principles to the requirements of the application environment. The result should be an architecture that meets the current needs but that is flexible enough to evolve to meet the needs of short- and long-term trends.

A solid foundation for Data Center design is based on a redundant, scalable, and flexible Layer 2 and Layer 3 infrastructure in which the behavior is both predictable and deterministic. The infrastructure also should accommodate service devices that perform key functions aimed at scaling or securing application environments. The deployment of service devices such as firewalls, load balancers, SSL offloaders, and caches requires careful planning.

The planning efforts must ensure that the desired behavior is achieved in the following areas: redundancy protocols between service devices, the exchange of connection and session information between stateful devices, the location of default gateway services, and the traffic path through the Data Center infrastructure from device to device.

Additional considerations require an architectural approach to deal with the application environment trends and the requirements that are imposed on the network infrastructure. Subsequent chapters in the book dig deeper into the specifics of Data Center and server farm designs.

INDEX

Numerics

3DES encryption, 600
10-GigE (10-Gigabit Ethernet), 492
10GBASE-ER, 493
10GBASE-EW, 493
10GBASE-LR, 493
10GBASE-LW, 493
10GBASE-LX4, 493
10GBASE-SR, 493
10GBASE-SW, 493
100BASE-FX, 489
100BASE-T. *See* Fast Ethernet, 489
100BASE-TX, 489
400 status codes (HTTP), 56
500 error codes (HTTP), 56
802.1Q tag all, 187
802.1s, 516
 configuring, 519–520
802.3ad, 33
1000BASE-LX, 491
1000BASE-SX, 491
1000BASE-T, 491
4096 VLANs, 514

A

A (Address) records, 403
AAA (Authentication, Authorization, and Accounting)
 RADIUS, 646
 security, 197
 TACACS+, 645
ABRs (Area Border Routers), 543
 summarizatoin, 550
absolute URIs, 312
absolute URLs, 316
Accept field (HTTP request header), 353
Accept-Charset field (HTTP request header), 353
Accept-Encoding field (HTTP request header), 354
Accept-Language field (HTTP header), 979–980
access layer
 application segment, 17
 back-end segment, 18
 front-end segment, 16
access ports, 32, 520, 839–840
access switches, 141
acknowledgment number field (TCP), 263
ACKs (TCP), 48, 666
ACLs (access control lists), 25, 170, 873
 dynamic, 171
 extended, 170
 reflexive, 172–173
 router, 170
 standard, 169
active-active firewall configuration, 906
active-active load balancing, 229–230
active-backup algorithm, 437
active-standby firewalls, 904
active-standby load balancing, 228
ActiveX controls, 86, 1017–1018
 server-side, 89
addresses
 formatting (Ethernet), 485–487
 MAC address table, 499
advertisement interval, 535
advertising local subnets (OSPF), 854
AES-Rijndael, 601
aggregation layer, 15
aggregation routers, connecting to core routers, 846–849
aggregation switches, 141
algorithms
 cache farm load-balancing, 683–685
 hashing, 607
 message digest, 607
 SHA, 608

- load balancing, 673
 - fastest, 680
 - hash address, 681
 - least connections, 678
 - round-robin, 676
 - server farm, 673–675
 - source IP, 681
 - URL and hash URL, 681
 - weighted least connections, 679
 - weighted round-robin, 677
- alternate Layer 3/Layer 2 designs, 133
- alternate ports, 829
- analog video streaming, 447
 - codecs, 448
- analyzing SSL traces, 391–393
- anomaly-based IDSs, 181
- antireplay protection, 190
- antispoofing filtering, 870
 - uRPF, 873
- Apache web servers, 330
 - virtual hosting configuration, 58–59
 - IP-based, 59
 - name-based, 61
 - port-based, 60
- APIs (application programming interfaces),
 - server-specific, 88
- applets, 86
 - Java, 1014–1015
- application architecture trends, 150–151
- application layer, 244
 - probes, 713
 - DNS probes, 717
 - FTP probes, 717
 - HTTP probes, 714
 - IMAP4 probes, 718
 - POP3 probes, 718
 - SMTP probes, 718
 - SSL probes, 715
 - security, 21
- application segment (access layer), 17
- application services, 24
- application tier, 77
- applications
 - Data Center architecture models
 - client/server, 9–10
 - multitier, 12
 - n-Tier, 11
 - enterprise, 71
 - integration, 75
 - EAI, 75–77
 - multitier design (case study), 108–111
 - network architecture implications, 97
 - clustering, 99–102, 104
 - load balancing, 97–98
 - security, 104–105, 107
 - n-Tier model, 77
 - database access, 95–96
 - markup languages, 79–83
 - middleware, 91–95
 - server-side programming, 87–91
 - user agents, 84–85
 - web servers, 86
 - portal, 72
 - TCP, 41
 - ACKs, 48
 - data processing, 41
 - HTTP, 47, 55–56
 - maximum burst size on high-speed networks, 49–50
 - segments, 42
 - Telnet, 43–46
 - windows, 47–48
 - UDP, 50–51
 - upgrades, 71
- APPN (advanced peer-to-peer networking), 572
 - node types, 579–580

- architectures
 - MLS, 809
 - of Data Centers
 - flexibility, 118
 - high availability, 118
 - scalability, 117
 - of load balancers, 232–235
 - critical components, 234–235
 - generic components, 232–234
- Area Border Routers (ABRs), 543
- ARP (Address Resolution Protocol), 525–526
 - ARP inspection, 184, 895
 - ARP spoofing, 167
 - timeout values compared with CAM tables, 526
- ASBR (autonomous system border router)
 - summarization, 550
- ASCII character sets
 - extended, 965–966
 - nonprintable, 963–964
 - printable, 964–965
- ASPs (active server pages), 88, 1022
- asymmetric cryptography, 602
 - D-H, 606
 - DSS, 605
 - RSA, 603–604
- asymmetric encryption, 191
- attachment options for mainframes, 573
 - channel attachments, 573–574
 - LAN attachments, 575
- attacks
 - buffer overflow, 167
 - DDoS, 164
 - DoS, 163
 - eavesdropping, 165
 - Internet infrastructure attacks, 166
 - Layer 2, 167–168
 - mitigation, 202
- scanning/probing, 162
- session hijacking, 167
- smurf, 163
- trust exploitation, 166
- unauthorized access, 165
- viruses and worms, 165
- attributes of cookies, 729–731
- audio streaming
 - transport formats, 442, 454
 - RTCP, 457–459
 - RTP, 454
- authentication, 640, 876
 - AAA protocols
 - RADIUS, 646
 - TACACS+, 645
 - challenge/response schemes, 642
 - digital certificates, 642
 - HTTP, 364
 - Kerberos, 644
 - management network, 911–913
 - OTPs, 641
 - SSL, 385–387
 - PKI, 388–389
- authenticity tags, 194
- authoritative name servers, zone transfers, 418–420
- Authorization field (HTTP request header), 354
- autonegotiation
 - Gigabit Ethernet, 492
 - NICs, 490
- autostate, 810, 814
- auto-unfail, 706
- availability, optimizing with load balancing, 65

B

baby giant frames, 496
BackboneFast, 827–828
back-end segment (access layer), 18
backup designated routers (BDRs), 542
backup ports, 829
bandwidth, 444
 scaling with Etherchannels, 815
baseline testing, performance metrics, 950
basic data transfer, 256
BDP (Bandwidth Delay Product), 50
BDRs (backup designated routers), 542
B-frames, 451, 991
BIND (Berkeley Internet Name Domain), 408
binding, 94
black-hole problem, 287–288
blade chassis, 152–156
blade servers, 21
bottlenecks, performance metrics, 933
BPDUs (bridge protocol data units), TCN (Topology Change Notification), 527
bridge identifiers, 510
bridging, 654
broadcast suppression, 487
browsers, 84
 cookies, 731–732
 multiple, 733
 session cookies, 769
 storage of, 734–735
 HTTP compression, 343
buffer overflow attacks, 167
bulk transfer traffic, 47
 ACKs, 48
 maximum burst size on high-speed networks, 49–50
 TCP windows, 47–48
bus and tag technology, 574

bus architecture

 PCI, 34
 PCI-X, 35
business continuance infrastructure services, 26
business continuance services, 27
BXN (branch extender node), 583

C

CA servers, 74
cabling, Ethernet, 481
cache load balancing, 210–211
 server farms, 683–685
Cache-Control field (HTTP general header), 344–345
caching, 25
 cache hits, 681
 cacheable objects, 683
 DNS, 420
 client applications, 422–423
 TTL values, 421
 hit rate, 673
 in site-selection architecture, 436–437
 on-demand, 472
 RPC, 683
 transparent, 684
caching-only servers (DNS), 411
campus core, security, 884
CAs (certificate authorities), 619
 certificates, 621
 deployment options, 623
 enrollment, 624
 key exchange, 620
 revocation, 625
CC metric, 933
 load balancers, 942–943
 SSL offloaders, 948
CDP (Cisco Discovery Protocol), 500

- CEF (Cisco Express Forwarding), 807–809
MLS, 821
certificates, SSL, 629
CF channels, 585
CGI, 88–89, 1018–1019
challenge response schemes, 642
channel link-layer protocols, 576
channeling, 507
channel-protocol lacp command, 508
channels, 569
connecting mainframes to peripheral devices, 573–574
character sets, 326
 ASCII
 extended, 965–966
 nonprintable, 963–964
 printable, 964–965
 ISO-8859-1, 969
checksum field (TCP), 266
chroma subsampling, 989
ciphers, 188
 export-grade, 611
 overview, 608
 RCs, 602
 SSL cipher suites, 632–633
ciphersuites, 371, 389–390
Cisco IOS Software
 internal redundancy, 835
 switching paths, 807–808
Cisco IPTV, 442
CISCO-SLB-MIB, 698–699
CLASSID, 1017
client error status codes (HTTP response header), 360
client NAT (load balancers), 662
 performance, 672
client tier, 77
client/server application model, 9–10
client/server architecture
 network attachment options, 32
 NICs, 32–33
PCI, 34
PCI-X, 35
server multihoming, 33
NICs, Ethernet driver, 36
packet processing, 35–36
sockets, 39
 system calls (UNIX), 39–40
TCP/IP processing, 37–39
clients
 browsers, 84
 thick, 83
 thin, 83
client-side programming, 85
 ActiveX controls, 1017–1018
 Java applets, 1014–1015
 JavaScript, 1013
cluster controllers, 570
clustered proxy servers, persistence, 759
clustering, 97, 382
 cluster modules, 100
 geographical, 101
 implications for application integration, 99–104
 Sysplex, 585–589
CNAME (Canonical Name) records, 404
codecs, 441, 448, 473
 comparison of, 452
 video encoding, 987
coded character sets, 327
CodeRed, 165
collaborative applications, 72
collapsed multitier design, 137–138
collapsed server-farm design, 898–900
collision domains, diameter, 487
commands, netstat -a, 37
communications controller, 570
components of IBM Data Centers, 570–573
compression
 HTTP, 342–343
 redundancy (video), 448

- confidentiality, 189
 - configuring
 - 802.1s, 519–520
 - cookie active, 775
 - cookie match, 772
 - cookie passive, 770
 - HTTP redirection stickiness, 783
 - Layer 2 features
 - access ports, 839–840
 - overview, 844
 - spanning trees, 841–843
 - trunks, 840
 - VLANs, 837–839
 - Layer 3 features, 846
 - default gateway redundancy, 849–851
 - EIGRP, 858–862
 - OSPF, 852–857
 - routing options, 846–849
 - load balancers for given applications, 98
 - loopback interfaces, 995
 - Linux, 1005–1006
 - Windows 2000, 996–998
 - Windows NT, 1002
 - NAT on routers and firewalls, 558
 - preemption, 851
 - rapid PVST+, 518
 - routing on servers, 524
 - server farms on a load balancer, 691
 - source IP stickiness, 765
 - mega proxies, 766–767
 - source IP hash, 768
 - SSL stickiness, 786
 - URL cookies, 779
 - web servers, 57
 - directories, 58
 - inserting cookies, 1010
 - server processes, 57
 - TCP parameters, 57
 - virtual hosting, 58–60
- congestion avoidance, 279
 - congestion control, 278
 - congestion window (TCP), 47
 - CONNECT method (request header), 351
 - connection establishment, Telnet sessions, 43–44
 - Connection field (HTTP general header), 345
 - connections, 257
 - embryonic, 564
 - failover, 231
 - HTTP, 335, 337
 - persistent connections, 339
 - pipelining, 340
 - load balancing, 674
 - long-lived, 929–931
 - maxconns, 678
 - performance metrics, 935
 - persistence, 219
 - reassigning, 704
 - remapping, 667
 - short-lived, 925–927
 - spoofing (load balancers), 664–667
 - connection remapping, 667–669
 - performance, 672
 - TCP, 267
 - establishment phase, 268–270
 - monitoring, 67
 - termination phase, 46, 272, 275
 - TCP/UDP, stickiness, 674
 - tracking, 219
 - connectivity, blade chassis options, 152–156
 - content switching, 205. *See also* server
 - load balancing
 - horizontal scaling, 206
 - versus DNS round-robin, 207–209
 - vertical scaling, 206
 - Content-Encoding headers, 343
 - control flags (TCP), 264–266
 - control protocols, 466
 - RTSP, 467–470
 - control units, 574

controllers, 493
convergence, 827
 MST, 831
 OSPF, 856
 PVST+, 828
 Rapid PVST+, 829–830
cookies, 221–223, 728
 browser storage, 734–735
 browser treatment of, 731–732
 browser treatment of multiple cookies, 733
 format, 729–730
 inserting, 1010
 load balancers
 cookie active, 775
 cookie match, 771–773
 cookie passive, 769
 persistent, 728–729
 session, 728–729
 specifications and standards, 735
 stickiness, 222
 tracking user sessions, 739
 URL, 776–778
CORBA, 92, 95
core routers, connecting to aggregation routers, 846–849
corporate Data Centers, 126
CPS metric, 933
 load balancers, 942
 SSL offloaders, 948
cryptography, 188–189
 asymmetric, 602
 D-H, 606
 DSS, 605
 RSA, 603
 RSA key exchange, 604
 asymmetric encryption, 191
 ciphers, 608
 export-grade, 611
 digital signatures, 195
 FIPS, 609

hashing algorithms, 193, 607
 message digests, 607
 SHA, 608
 HMACs, 194
 NIST, 609
 PKI, 612
 CAs, 619–625
 digital certificates, 615–619
 standards, 614
 symmetric, 190, 597
 3DES, 600
 DES, 598–600
 RCs, 602

D

dark fiber, 104
data, 452
 encoding, 448–451
 multimedia transport formats, 454
 RTP, 454, 457–459
 UDP versus TCP, 445–446
 packetization, 453
 replication, 103
 TCP, 463
 transport security, 626
 IPSec, 633–634, 637–638
 SGC, 631
 SSL, 626, 628–629
 SSL cipher suites, 632–633
 VPNs, 639
Data Centers
 application architecture
 client/server model, 9–10
 multitier, 12
 n-Tier model, 11
 applications
 EAI, 75–77
 integration, 75

- multitier design (case study), 108, 111
- network architecture implications, 97–107
- n-Tier model, 77–96
- portal, 72
- architecture, 13–14
 - access layer, 16–18
 - aggregation layer, 15
 - layers, 14
 - storage layer, 19
 - transport layer, 20–21
- design criteria, 6
- facilities, 7
- goals, 6
- high availability, 109
- infrastructures, 801–805
 - spanning trees, 822
 - virtualizing with VLANs, 804, 810, 813–814
- Layer 2 design
 - access ports, 839–840
 - configuration overview, 844
 - spanning trees, 841–843
 - trunk configuration, 840
 - VLAN configuration, 837–839
- Layer 3 design, 846
 - default gateway redundancy, 849–851
 - EIGRP, 858–862
 - OSPF, 852–857
 - routing considerations, 846–849
- overview, 5
- performance metrics, 934–935
 - firewalls, 938
 - load balancers, 939–945
 - multilayer switches, 936–937
 - SSL offloaders, 946–949
 - testing, 950–957
- redundancy, 833
 - NSF, 835–837
 - supervisor redundancy, 834–835
 - redundant links, 815–817
- roles
 - enterprise, 7
 - SP environment, 9
- security framework
 - incident response and attack mitigation, 202
 - secure management framework, 200–201
 - security life cycle, 198
 - security policies, 198
 - zones, 866
- server failure detection, 700
 - probes, 701
 - SNMP, 701
- server management, 689–690
 - CISCO-SLB-MIB, 698–699
 - DFP, 708
 - graceful shutdown feature, 691
 - HTTP and HTTPS (case study), 722–723
 - in-band probes, 703–706
 - load balancing overview, 690
 - Max/Min Connections, 694–695
 - out-of-band probes, 707–708, 711, 713–714, 716–718
 - probe comparison, 709
 - slowstart feature, 693
 - SNMP, 697–698
 - virtual hosting (case study), 718–720
 - XML, 696–697
- services, 22
 - application, 24
 - business continuance, 26–27
 - IP infrastructure, 23
 - security, 25
 - storage, 26
 - static routing, 527
 - traffic patterns, 924
 - long-lived traffic, 929–931
 - performance metrics, 933
 - short-lived traffic, 925–927
 - VLANs, 502

- data processing on TCP applications, 41
- database access, 95–96
- database middleware, 91
- database servers, 73
- database tier, 77
- datagrams, 245
- Date field (HTTP general header), 346
- DBMSs (database management systems), 96
- DCOM objects, 93–95
 - passing through firewalls, 95, 106
- DCT (discrete cosine transform), 988
- DDoS (distributed denial-of-service) attacks, 164
- debounce feature, 831
- decryption, 188
- dedicated Internet server farms, 120
- defining
 - security zones, 865–868
 - VTP domains, 504
- delayed ACKs, 45, 280
- delegated name servers, 428
- DELETE method (request header), 351
- deploying
 - antispooing filtering, 870
 - services in redundant Layer 2/Layer 3 Data Centers, 148
- DES encryption, 598–600
- designated ports (DPs), 512, 829
- designing
 - Data Centers
 - bus architecture, 34–35
 - client/server architecture, 35–39
 - criteria, 6
 - flexibility, 118
 - fully redundant Layer 2/Layer 3 designs, 139–157
 - high availability, 118
 - optimizing performance, 62–67
 - scalability, 117
 - server multihoming, 33
 - EAI networks, 76–77
- high availability, 51
- management network security, 914
- NICs, 32–33
- server farms
 - alternate Layer 2/Layer 3 designs, 133
 - collapsed server-farm design, 898–900
 - expanded server-farm design, 900–902
 - generic Layer 2/Layer 3 designs, 126–131
 - multiple-tier designs, 133–138
 - redundant firewall designs, 904–906
 - VLANs, 505–506
- devices, codecs, 987
- DFP (Dynamic Feedback Protocol), 675, 708
- D-H, 606
- DHCP servers, 74
- diameter, 487
- diffusing DUAL, 553
- digital certificates, 615, 642
 - extensions, 619
 - formats, 617
 - generating, 616
 - SSL authentication, 385–387
- digital signatures, 195
- Digital Video Compression (DVC), 450
- digital video streaming, 447
- Direct Server Return (DSR), 669–670
- directed mode (load balancers), 654, 660–661
 - performance, 672
- directories, configuring on web servers, 58
- directory servers, 74
- directory services (APPN), 579
- discarding ports, 511
- disk replication, 102
- dispatch mode (load balancers), 654, 657–659
 - performance, 672
- distributed DVIPAs, 588
- distributing multiple records
 - A records, 425
 - client applications, 426
 - NS records, 423–424

- distribution servers, 471
DivX, 451
DLSw (Data Link Switching), 580–581
DLUR/DLUS (dependent LU requesters/dependent LU servers), 583
DMA (Directe Memory Access), 33
DMZ server farms, 120
DNS (domain naming system), 397
 A records, 425
 caching, 420
 client applications, 422–423
 TTL values, 421
 forwarders, placement of, 427–428
 FQDNs, 399–400
 hierarchical name structure, 398–399
 name resolution process, 404–406
 name servers, 418
 NS records, 423–424
 probes, 713
 queries, communication flows, 420
 resolution process, 411
 iterative queries, 417
 queries, 412
 recursive queries, 417
 referrals, 414–417
 root hints, 413–414
 resource records, 402–403
 servers, 74, 407
 signatures, 881
 site-selection architecture, 430–433
 caching, 436–437
 proximity, 435
 referrals to site selectors, 433–435
 stickiness, 437–438
 split namespace, 428–430
 TLDs, 399
 zone transfers, 418–420
 zones, 400–402
DNS proxy, 409
 caching-only servers, 411
 forwarders, 410
- DNS round-robin, 207–209
domain hash predictor, 685
DoS attacks, 163
 preventing with traffic rate limiting, 874
 smurf, 163
download-and-play, 442–444
download rate (streaming traffic), 466
DP (designated port), 512
DSR (Direct Server Return), 669–672
DSS (Digital Signature Standard), 605
DTP (Dynamic Trunking Protocol), 501
dual-attached servers, 821
dummy unicast MAC addresses, 98
DV (Digital Video Compression), 450
DVIPA (dynamic VIPA), 587
 distributed DVIPAs, 588
dynamic ACLs, 171
Dynamic Feedback Protocol (DFP), 675, 708

E

- EAI, 75
 network design implications, 76–77
eavesdropping, 165
ECB (electronic code book), 600
e-commerce applications, 727
 session persistence, 757, 790
e-commerce applications, 72
edge ports, 829, 840
EEs (enterprise extenders), 582–583
EIGRP (Enhanced IGRP), 551, 858
 configuration overview, 862
 default advertisement, 555
 default routers, 860
 failure detection, 552
 metric tuning, 553–554
 redistribution, 554
 summarization, 860

- summarization and filtering, 555
- topology, 859
- EJBs, 93
- electronic code book (ECB), 600
- e-mail servers, 73
- e-mail signatures, 881
- embryonic connections, 564
- encoding, 448, 473
 - formats, 450–451
 - HTTP, MIME comparison, 326
 - MIME, 323–324
 - character sets, 326
 - HTTP comparison, 326
 - media types, 327–328
 - transport rate, 452
 - URLs, 316
 - reserved characters, 318
 - unsafe characters, 318
 - URNs, 320
- encoding video, 987
- encryption, 910
 - 3DES, 600
 - asymmetric, 191
 - control data, 201
 - cryptography, 188–189
 - DES, 598, 600
 - symmetric, 190
- ENs (end nodes), 572
- enterprise networks
 - applications, 71
 - Data Center roles, 7
 - Data Centers, 126
 - architecture, 13–21
 - services, 22–27
- entity header, 365
- Entity header fields (HTTP), 985
- ephemeral RSA, 631
- ESCD (ESCON directors), 574
- ESCON (enterprise system connections), 574
- establishing TCP connections, 268, 270
- establishment controllers, 570
- Etherchannels, 507
 - creating channels, 507
 - scaling bandwidth, 815
- Ethernet
 - 10-GigE, 492
 - physical layers, 495
 - 10GBASE-ER, 493
 - 10GBASE-EW, 493
 - 10GBASE-LR, 493
 - 10GBASE-LW, 493
 - 10GBASE-LX4, 493
 - 10GBASE-SR, 493
 - 10GBASE-SW, 493
 - 100BASE-FX, 489
 - 100BASE-TX, 489
 - 1000BASE-LX, 491
 - 1000BASE-SX, 491
 - 1000BASE-T, 491
 - address format, 485–487
 - EtherChannels, 507
 - creating channels, 507
 - Fast Ethernet, 489
 - autonegotiation, 490
 - frame size, 488
 - physical layers, 494
 - frames
 - baby giant, 496
 - format, 482–484
 - jumbo, 496
 - size, 487–488
 - Gigabit Ethernet, 491
 - autonegotiation, 492
 - flow control, 492
 - physical layers, 495

Layer 2 protocols, 500–501
overview, 481
physical layers, 493
switching, 498–500
examples of SSL applications
 HTTPS, 372–374
expanded multilayer design, 135–136
expanded server-farm design, 900–902
Expect field (HTTP header), 980
export-grade ciphers, 611
extended ACLs, 170
extended ASCII character sets, 965–966
Extensible Markup Language. *See* XML
external redundancy, 833
extranet server farms, 124

F

failure detection
 EIGRP, 552
 HSRP, 531
 OSPF, 545
 redundant firewalls, 906
failure recovery, spanning trees, 842
Fast Ethernet, 489
 autonegotiation, 490
 frame size, 488
 transceivers, 495
fast paths, 933
fast recovery, 280
fast retransmission, 446
fast switching, 807
FastCGI, 89
fastest predictor, 680
FCIP (Fibre Channel over IP), 103
FEPs (front-end processors), 570
FICON (fiber connectivity), 574
fields
 HTTP entity headers, 365
 HTTP general headers, 344–347

HTTP messages, 334
HTTP response headers, 362–363
IP headers
 flags field, 251
 fragment offset field, 251
 header checksum field, 254
 header length field, 248
 identifier field, 250–251
 options field, 255–256
 protocol field, 252–254
 TOS field, 248–250
 total length field, 250
 TTL field, 251–252
 Version field, 247
request headers, 352
 Accept field, 353
 Accept-Charset field, 353
 Accept-Encoding field, 354
 Authorization field, 354
 Host field, 354
 If-Modified-Since field, 355
 Max-Forwards field, 355
 Range field, 355
 Referer field, 355
 User-Agent field, 356
TCP headers
 acknowledgment number field, 263
 checksum field, 266
 control flags, 264, 266
 options field, 266–267
 sequence number field, 262
 TCP header length field, 264
 urgent pointer field, 266
 window size field, 266
UDP headers, 299–301
file servers, 73
filtering
 ACLs, 873
 antispoofing, 870
 EIGRP, 555

- OSPF, 550
- packet filters, 890
- RFC 1918, 870
- RFC 2817, 870
- route filters, 876
- final permutation, 600
- FIPS, 609
- firewall load balancing, 212–213
- Firewall Service Module (FWSM), 887
- firewalls, 173
 - hybrid, 176–177
 - Internet traffic patterns, 921
 - limitations, 178
 - NAT, 557
 - packet-filtering, 174
 - passing DCOM through, 95, 106
 - performance metrics, 938
 - PIX, NAT, 563–564
 - proxy, 175
 - redundant
 - active-active (clusters), 906
 - redundant firewall server-farm design, 904
 - server farm design, 905–906
 - stateful, 175, 878–879
 - flags field, 251
 - flexibility in Data Center design, 118
 - flooding, 98, 831
 - unicast, 499
 - flow control, 257
 - congestion avoidance, 279
 - congestion control, 278
 - delayed ACKs, 280
 - fast recovery, 280
 - immediate ACKs, 280
 - Nagle algorithm, 281–282
 - retransmission, 276
 - sliding windows, 277
 - slow start, 279
 - flow-based forwarding, 809
 - flow-based MSL, 820
 - forking servers, 51
 - versus threaded servers, 53
 - form fields, 91
 - form hidden fields, 737
 - formal namespaces, 322
 - forward zones, 402
 - forwarders (DNS), 410
 - placement of, 427–428
 - forwarding delay, 520
 - forwarding links, failure, 843
 - forwarding ports, 511
 - FQDN (fully qualified domain name), 399–400
 - fragment offset field, 251
 - frame/packet loss, 937
 - frames, 42, 245, 487–488
 - defining nonstandard size, 497
 - Ethernet
 - baby giant frames, 496
 - jumbo frames, 496
 - formatting (Ethernet), 482–484
 - jumbo, 33
 - From field (HTTP header), 980
 - front-end segment (access layer), 16
 - FTP (File Transfer Protocol)
 - probes, 717
 - session persistence, 755–756
 - full NAT, 662
 - full URIs, 312
 - fully switched topology, 804
 - FWSM (Firewall Service Module), 887
 - election process, 905
 - failure detection, 906

G

-
- gateway redundancy, 849–851
 - GDPS (geographically dispersed parallel Sysplex), 589

-
- general header (HTTP), 344
 - Cache-Control field, 344–345
 - Connection field, 345
 - Date field, 346
 - Pragma field, 346
 - Transfer-Encoding field, 347
 - generic Layer 3/Layer 2 designs, 126–130
 - Layer 2 access switches, 130–131
 - geographical clustering, 101
 - GET method (request header), 349
 - Gigabit Ethernet, 491
 - 10-GigE, 492
 - autonegotiation, 492
 - flow control, 492
 - GLBP, 527, 536, 818
 - active/standby election, 537
 - failure detection, 538–539
 - load distribution, 540
 - glean adjacencies, 808
 - glue records, 415
 - GOP (Group of Pictures), 450
 - graceful shutdown feature, 691
 - gratuitous ARP, 526
 - grid computing, 151
 - Group of Pictures (GOP), 450

 - H**

 - H.261, 450
 - H.263, 450
 - half-closed connections, 282
 - handshakes (SSL), 374–375
 - session negotiation phases, 376–378
 - session resumption, 380–382
 - hard failures, 117
 - hardware
 - load balancing, 98
 - performance metric testing, 953
 - hash address predictor, 681
 - hashing algorithms, 607
 - message digests, 607
 - SHA, 608
 - HEAD method (request header), 349
 - header checksum field, 254
 - header compression, 296–298
 - UDP, 305
 - header fields of IPv4, 246
 - flags field, 251
 - fragment offset field, 251
 - header checksum field, 254
 - header length field, 248
 - identifier field, 250–251
 - options field, 255–256
 - protocol field, 252–254
 - TOS field, 248–250
 - total length field, 250
 - TTL field, 251–252
 - Version field, 247
 - header length field, 248
 - health checks, 690
 - hierarchical DNS name structure, 398–399
 - FQDN, 400
 - resource records, 402–403
 - zones, 400–402
 - high availability, 51, 109, 227
 - in Data Center design, 118
 - redundancy protocol, 226, 228
 - active-active environments, 229–230
 - active-standby environments, 228
 - server failures, 54
 - SYN retransmission, 55
 - TCP timeouts, 54
 - hint-tracks, 453
 - hit rate, 673
 - HMACs (hash method authentication codes), cryptographic, 194
 - horizontal scaling, 206

- Host field (HTTP request header), 354
host replication, 102
host-based IDSSs, 180, 880–882, 893
host-route adjacencies, 808
HSRP (Hot Standby Routing Protocol), 527–528
 failure detection, 531
 groups, 530
 preempt option, 529
 tracking, 533
HTML (Hypertext Markup Language), 79–80
 form fields, 91
HTTP (HyperText Transfer Protocol), 47
 applications, 55–56
 authentication, 364
 character sets, 327
 configuring on web servers, 57
 connection remapping, 667–669
 connections, 335–337
 cookies, 728
 entity header, 365
 Entity header fields, 985
 functionality, 329–330
 general header, 344
 Cache-Control field, 344–345
 Connection field, 345
 Date field, 346
 Pragma field, 346
 Transfer-Encoding field, 347
 header fields
 Accept-Language, 979–980
 Expect, 980
 From, 980
 If-Match, 981
 If-Modified-Since, 982
 If-None-Match, 981
 If-Range, 981
 Proxy-Authorization, 982
 TE, 982
 Trailer, 977
 Upgrade, 978
 Via, 978
 Warning, 978
 HTTP redirection, 782–784, 792
 message format, 332
 components, 334
 fields, 333
 methods, 309
 MIME comparison, 326
 overview, 328
 performance
 attribute comparision, 341
 compression, 342–343
 version differences, 340
 persistent connections, 339
 pipelining, 340
 probes, 714
 RDT, 466
 redirection, 782–784, 792
 request header, 347
 CONNECT method, 351
 DELETE method, 351
 fields, 352–356
 GET method, 349
 HEAD method, 349
 methods, 348
 OPTION method, 348
 POST method, 349
 PUT method, 350
 request URI, 351
 TRACE method, 351
 request/response, 333
 fields, 362–363
 Status-Codes, 356–362
 servers, 87
 health management (case study), 722–723
 virtual hosting, 58–61
 session persistence, 754–755, 757
 signatures, FTP signatures, 881
 status codes, 983–985
 streaming, 442–444

- tunneling, 461, 466
- URIs, 310
- versions, 330
- HTTPS (HTTP over SSL), 372–374
 - server health (case study), 722–723
- hybrid firewalls, 176–177
- hybrid servers, 53

- I/O handling, 35–36
- IANA, language tags, 980
- IBM Data Centers, 570–573, 590–591
- IBM networking, 577
 - APPN, 572
 - mainframes, 569–575
 - SNA
 - APPN, 579–580
 - over TCP/IP, 580–585
 - subnetwork SNA, 577–579
 - VTAM, 571
 - Sysplex, 585–588
 - GDPS, 589
- ICANN (Internet Corporation for Assigned Names and Numbers), 399
- ICMP (Internet Control Message Protocol)
 - probes, 711
- IDCs (Internet Data Centers), 9, 125
- identifier field, 250–251
- IDSs (intrusion detection systems), 178
 - anomaly-based versus signature-based, 181
 - host-based, 180
 - Internet edge, 880–882
 - intranet server farms, 891–893
 - network-based, 179, 891
 - responses, 182
 - signatures, 107, 181, 891
- IEEE 802.1D, 501
- IEEE 802, 479

- IEEE 802.1Q, 501
- IEEE 802.3ad, 33, 501
- If-Match field (HTTP header), 981
- If-Modified-Since field (HTTP header), 982
- If-Modified-Since field (HTTP request header), 355
- If-None-Match field (HTTP header), 981
- I-frames, 450, 990
- If-Range field (HTTP header), 981
- IKE (Internet Key Exchange), 637
- IMAP4 probes, 718
- immediate ACKs, 280
- in-band health verification, 67
- in-band probes, 703–705
 - HTTP return code checks, 706
 - server recovery, 706
- incomplete adjacencies, 808
- informational status codes (HTTP response header), 357
- infrastructure (Data Centers), 801–805
- inserting cookies, 1010
- inside global addresses, 558
- inside local addresses, 558
- integrating applications, 75
 - EAI, 75–77
 - network architecture implications, 97
- integrity, 189
- Internet infrastructure security attacks, 166
- interactive traffic, 41–43
 - connection termination, 46
 - delayed ACKs, 45
 - MSS, 44
 - Nagle algorithm, 46
 - TCP retransmission, 44
- interfaces
 - database access, 96
 - SVIs, 813
- interleaving, 470
- internal redundancy, 833
 - NSF, 835, 837
 - supervisor redundancy, 834–835

- Internet
 - HTTP, 328
 - traffic patterns, 919–921
 - long-lived traffic, 931
 - protocols, 922
 - short-lived traffic, 926
 - Internet Data Centers, 9, 125
 - Internet edge security, 869
 - ACLs, 873
 - antispoofing filtering, 870
 - IDSs, 880–882
 - Internet edge design, 882
 - securing routing protocols, 875–876
 - stateful firewalls, 878–879
 - traffic rate limiting, 874
 - uRPF, 872–873
 - Internet server farms, 120
 - dedicated, 120
 - DMZ server farms, 120
 - interrupt coalescing, 33, 63
 - interrupt processing, optimizing, 62–63
 - intranet server farms, 122–124
 - security, 885–886
 - ARP inspection, 895
 - IDSs, 891–893
 - packet filters, 890
 - port security, 894
 - server-farm design alternatives, 896–906
 - stateful firewalls, 887–888
 - VLAN features, 895
 - intrranets
 - traffic patterns, 919–920, 923
 - long-lived, 931
 - short-lived, 926
 - IOS NAT, 561–562
 - IP addressing, DVIPAs, 587–588
 - IP header compression, enabling on Cisco routers, 298
 - IP infrastructure services, 23
 - IP spoofing, 167
 - IP-based virtual web hosting, 59
 - IPSec, 633
 - IKE, 637
 - security parameters, 638
 - TCP/IP layers, 634
 - VPNs, 639
 - IPTV, 442
 - IPv4 header, 246
 - flags field, 251
 - fragment offset field, 251
 - header checksum field, 254
 - header length field, 248
 - identifier field, 250–251
 - options field, 255–256
 - protocol field, 252–254
 - TOS field, 248–250
 - total length field, 250
 - TTL field, 251–252
 - Version field, 247
 - ISAPI, 88
 - iSCSI, 103
 - ISL (InterSwitch Link), 501–503
 - ISO-8859-1 character set, 969
 - isolation, 910
 - iterative queries (DNS), resolution process, 417

J

-
- J2EE (Java 2 Enterprise Edition), 92
 - Java
 - applets, 86, 1014–1015
 - database access, 96
 - J2EE, 92
 - servlets
 - case study, 90–91
 - user session tracking, 743
 - Java Virtual Machine (JVM), 1014–1015
 - JavaScript, 86, 1013
 - server-side, 88

JSPs, 88, 1021
 jumbo frames, 33, 496
 optimizing interrupt processing, 63
 JVM (Java Virtual Machine), 86, 1014–1015

K–L

Keep-Alive field (HTTP messages), 334
 keepalives, TCP, 55
 Kerberos, 644
 kernel mode, 35–36
 language tags, IANA, 980
 LANs
 10-GigE, 492
 physical layers, 495
 connecting mainframes to peripheral devices, 575
 Ethernet
 addresses, 485–487
 baby giant frames, 496
 frame size, 487–488
 frames, 482–484
 jumbo frames, 496
 Layer 2 protocols, 500–501
 overview, 481
 physical layers, 493
 switching, 498–500
 Fast Ethernet, 489
 autonegotiation, 490
 physical layers, 494
 Gigabit Ethernet, 491
 autonegotiation, 492
 flow control, 492
 physical layers, 495
 IEEE 802, 479
 VLANs. *See* VLANs

latency
 load balancers, 942, 944
 multilayer switch metrics, 937
 SSL offloaders, 949
 Layer 2
 access ports, 839–840
 attacks, 167–168
 configuration overview, 844
 convergence, 827
 MST, 831
 PVST+, 828
 Rapid PVST+, 829–830
 dual-attached servers, 821
 Ethernet. *See* Ethernet
 security, 183
 802.1Q tag all, 187
 ARP inspection, 184
 port security, 183
 private VLANs, 185–187
 spanning trees, 841, 843
 STP, 508–520
 traffic distribution, 818
 trunk configuration, 840
 VLAN configuration, 837–839
 Layer 2/Layer 3 designs, redundancy, 139
 access layer, 141–146
 application architecture trends, 150–151
 network infrastructure trends, 152–157
 services, 146–150
 Layer 3
 design options, 846
 default gateway redundancy, 849–851
 EIGRP, 858, 860, 862
 OSPF, 852–854, 856–857
 routing considerations, 846, 849
 links, 805
 protocols, 523
 ARP, 525–526
 EIGRP, 551–555
 GLBP, 536–540

- HSRP, 528–533
- NAT, 556–566
- OSPF, 541–551
- VRRP, 534–535
- redundant paths, 814
- switches, 807
- traffic distribution, 819–820
- Layer 4 load balancing, 216
- Layer 5 load balancing, 217
 - persistence, 754
- layers of OSI reference model, 241–243
 - application layer, 244
- learning ports, 511
- least connections predictor, 678
- LEN (low-entry networking) nodes, 579
- links
 - EtherChannels, 816
 - Layer 3, 805
 - load distribution, 815, 817
 - Layer 2, 818
 - Layer 3, 819–820
 - redundant, 815–817
- Linux
 - configuring loopback interfaces, 1005–1006
 - enabling PMTUD, 291–292
- load balancers
 - HTTP redirection, 782–784
 - Internet traffic patterns, 921
 - NAT, 557
 - performance metrics, 939–941
 - CC metric, 943
 - CPS metric, 942
 - latency, 942, 944
 - PPS metric, 944
 - response time, 945
 - persistence, 754
 - comparing mechanisms, 789
 - cookies, 769–775
 - predictors, 761
 - SSL persistence, 791
 - sticky groups, 764
 - sticky methods, 762
 - URL cookies, 794
 - PIX, NAT, 565–566
 - reassigning connections, 705
 - server failure detection, 700
 - probes, 701
 - SNMP, 701
 - server health management, 690
 - CISCO-SLB-MIB, 698–699
 - DFP, 708
 - graceful shutdown feature, 691
 - in-band probes, 703–706
 - Max/Min Connections, 694–695
 - out-of-band probes, 707–708, 711–718
 - probe comparison, 709
 - slowstart feature, 693
 - SNMP, 697–698
 - XML, 696–697
 - source IP hash, 768
 - source IP stickiness, 765–767
 - SSL stickiness, 785
 - challenges and concerns, 787–788
 - configuring, 786
 - traffic patterns, 939
 - URL cookies, 776–778
 - URL hash, 780–781
 - URL match, 779
- load balancing, 24, 97, 205
 - algorithms, 673
 - cache farm load-balancing, 683–685
 - fastest, 680
 - hash address, 681
 - least connections, 678
 - round-robin, 676
 - server farm, 673–675
 - source IP, 681
 - URL and hash URL, 681
 - weighted least connections, 679
 - weighted round-robin, 677

architecture, 232–235
critical components, 234–235
generic components, 232–234
cache load balancing, 210–211
client NAT, 662
connection failover, 231
connection persistence, 219
connection spoofing, 664–669
connection tracking, 219
directed mode, 660–661
dispatch mode, 657–659
DSR, 669–670
firewall load balancing, 212–213
flexibility, 659
hardware, 98
high availability, redundancy protocol, 226, 228–230
horizontal scaling, 206
implications for application integration, 97–98
Layer 4 load balancing, 216
Layer 5 load balancing, 217
modes of operation overview, 653
optimizing server availability, 65
overview, 690
performance, 671–672
process description, 215–216
proxy servers, 760
RTP, 472
server health, 224
 in-band server health tracking, 224
 out-of-band server health tracking, 225
server load balancing, 209–210
server-selection mechanism, 654
session persistence, 219
 cookies, 222–223
 session-sharing servers, 761
SSL traffic, 382, 384
stateful failover, 231
stateless failover, 231
sticky failover, 231
unicast streaming, 472
versus DNS round-robin, 207–209
vertical scaling, 206
VPN/IPSec load balancing, 211
load distribution, 815, 817
dual-attached servers, 821
EtherChannels, 816
Layer 2, 818
Layer 3, 819–820
looped topologies, 819
 loop-free topologies, 818
load-share adjacencies, 808
local DUAL, 552
lock and key, 171
logical ports, 517–518
long-lived traffic, 929–931
 performance metrics, 933
loop-free topology, 818
loopback interfaces. configuring, 995
 Linux, 1005–1006
 Windows 2000, 996–998
 Windows NT, 1002
looped topologies, 818–819
loop-free topologies, 832–833
 load distribution, 818
 spanning trees, 822–825
Loopguard, 832–833
LPAR (logical partitions), 570, 576
LSAs, 544
LU Type 6.2, 579
LUs (logical units), 571

M

MAC address tables, 499
MAC addresses, 486
 flooding, 168
 Layer 2 protocols, 501
 reducing, 514
 redundant firewalls, 905

- mac-address-table aging-time command, 500
- macroblocks, 991
- mainframes, 569
 - attachment options, 573
 - channel attachments, 573–574
 - LAN attachments, 575
- FEP, 570
- LPAR, IP addressing, 576
- operating systems, 570
- Management Information Bases. *See* MIBs
- management networks, security, 908
 - authentication, 911–913
 - encryption, 910
 - isolation, 908–910
 - secure design, 914
- man-in-the-middle attacks, 184
- MANs, IEEE 802, 479
- markup languages
 - HTML, 79–80
 - WML, 83
 - XML, 79, 82–83
- master-down interval, 535
- Max Connections parameter, 694–695
- maxconns, 678
- Max-Forwards field (HTTP request header), 355
- maximum connections, 682
- Maximum Transmission Unit, 488
- MD5 (Message Digest-5), 607
- media types, 327–328
- mega proxies, 766–767
- messages
 - HTTP, 309, 332
 - components, 334
 - fields, 333
 - MIME, 323–324
 - character sets, 326
 - HTTP comparison, 326
 - media types, 327–328
- messaging middleware, 91
- META tag, configuring web servers to insert cookies, 1010
- methods
 - HTTP, 309
 - request header, 347–348
 - CONNECT, 351
 - DELETE, 351
 - GET, 349
 - HEAD, 349
 - OPTION, 348
 - POST, 349
 - PUT, 350
 - TRACE, 351
 - URLs, 316
- metrics
 - performance, 934–935
 - firewalls, 938
 - load balancers, 939–945
 - multilayer switches, 936–937
 - SSL offloaders, 946, 948–949
 - testing, 950–957
 - tuning
 - EIGRP, 553–554
 - OSPF, 547, 856
- MHSRP, 818
- MIBs (Management Information Bases), 698
 - platform flexibility, 702
 - RMON, 699
- Microsoft .NET, 92
- middleware, 76, 91–92
 - components, 93
 - traffic patterns, 94–95
- MIME format, 323–324
 - character sets, 326
 - HTTP comparison, 326
 - media types, 327–328
- Min Connections parameter, 694–695
- MJPEG (Motion JPEG), 450

MLS (Multilayer Switching)

architectures, 809
 CEF-based, 821
 flow-based, 820
 switching paths, 809
mod_session, session-tracking case study, 740–741
 modifying TCP keepalive defaults, 55
 monitoring TCP connections, 67
 motion estimation, 989
MPEG (Motion Pictures Experts Group), 990–991
 macroblocks, 991
 MPEG1, 450
 MPEG2, 450
 slices, 991
MSS (maximum segment size), 44, 283–284
MST, 823, 831
MTU (Maximum Transmission Unit), 488
 mtu command, defining nonstandard frame size, 497
 multicast addresses, mapping, 486
 multicast packets, 471
 multicast streaming, 24
 multilayer switches, performance metrics, 936–937
 latency, 937
 throughput, 936
 multimedia streaming, TCP versus UDP, 445–446
 multimedia transport formats, 454
 RTCP, 457–459
 RTP, 454
 multiple-tier designs, 133, 135
 collapsed multitier design, 137–138
 expanded multitier design, 135–136
 multiplexing, 257
 multiprocess application servers, 53
 multiprocess servers, 53
 multitier architecture application environment, 12
MX (Mail Exchange) records, 404

N

Nagle algorithm, 46, 281–282
 name servers, 418
 name-based virtual hosting, 61
 namespace, URNs, 321
 naming relative URIs, 314–315
NAT (Network Address Translation), 556–558, 663
 application support, 559–560
 IOS NAT, 561–562
 load balancers, 565–566
 PIX firewalls, 563–564
 native VLANs, 503
 NAU (network addressable unit), 571
NBMA (nonbroadcast multiaccess), 542
NCP (Network Control Program), 570
 negative caching, 421
 neighbor router authentication, 876
Netscape
 introduction of cookies, 735
 JavaScript, 1013
netstat –a command, 37
 network infrastructure trends, 152–157
 network management security
 SNMPv3, 649
 SSH, 647
 network security infrastructure, 169
 ACLs, 169–171
 firewalls, 173
 hybrid, 176–177
 limitations, 178
 packet-filtering, 174
 proxy, 175
 statefull, 175
 IDSs, 178
 anomaly-based versus
 signature-based, 181

- host-based, 180
- network-based, 179
- responses, 182
- signatures, 181
- Layer 2, 183
 - 802.1Q tag all, 187
 - ARP inspection, 184
 - port security, 183
 - private VLANs, 185
 - private VLANs with firewalls, 187
- network-based IDSSs, 179, 891
- networks
 - campus core, security, 884
 - Data Centers
 - roles of, 7
 - SP environment, 9
 - designing, multitier design (case study), 108–111
 - Internet edge security, 869–882
 - intranet server farms
 - design alternatives, 896–906
 - security, 885–895
 - management network security, 908–914
 - security, implications for application
 - integration, 104–107
 - traffic patterns, 923
 - VLANs, 502
 - access ports, 520
 - creating trunks, 505–506
 - designing, 505
 - PVIDs, 503
 - trunks, 503
 - NICs (network interface cards), 32–33
 - autonegotiation, 490
 - Ethernet driver, 36
 - interrupt coalescing, 63
 - server multihoming, 33
 - NIDs (namespace IDs), 321
 - Nimda, 165
- NIST (National Institute of Standards and Technology), 609
- NNs (network nodes), 572
- node types (APPN), 579–580
- nonbroadcast multiaccess (NBMA), 542
- nonedge ports, 829, 840
- nonprintable ASCII character sets, 963–964
- nonrepudiation, 189
- NS (Name Server) records, 403, 423–425
- NSAPI, 88
- NSF, 835–837
- NSSAs (not-so-stubby areas), 543
- n-Tier model, 11, 77
 - database access, 95–96
 - Java, 96
 - markup languages, 79–83
 - middleware, 91–92
 - components, 93
 - traffic patterns, 94–95
 - server-side programming, 87–89
 - case study, 90–91
 - user agents, 84
 - browsers, 84
 - client-side programming, 85
 - helpers and plug-ins, 85
 - web servers, 86

O

- object middleware, 91
- OIDs (object identifiers), 697
- on-demand caching, 472
- operating systems
 - LPAR, 570
 - mainframe-based, 570
 - UNIX, system calls, 39–40

optimizing server performance, 62
interrupt processing, 62–63
load balancing, 65
preventing server overload, 65–67
reverse proxy caching, 63
SSL, 384–385

OPTION method (request header), 348

options field, 255–256
TCP header, 266–267

OSA (open system adapters), 576

OSI reference model, 241–243
application layer, 244

OSPF, 541–542, 852
advertising the local subnets, 854
area assignment and summarization, 853
areas, 543
convergence time, 856
default advertisement, 551
failure detection, 545
LSAs, 544
metric tuning, 547, 856
neighbor states, 542
redistribution, 547–549
stub areas, 854
summarization and filtering, 550
topology, 852

OTPs (one-time passwords), 641

OUI (organizationally unique identifier) format, 486

out-of-band probes, 707–708
application layer, 713
DNS probes, 717
FTP probes, 717
HTTP probes, 714
IMAP4 probes, 718
POP3 probes, 718
SMTP probes, 718
SSL probes, 716

ICMP, 711

TCP, 711

UDP, 712

outside global addresses, 558
outside local addresses, 558
overloaded servers, 65–67

P

packet filters, 890
packet processing, 35–36
packet-filtering firewalls, 174
packetization, 453
packets
directed mode processing, 661
Ethernet, 482
filtering, ACLs, 25
header rewrites, 656
multicast, 471
RMI, 94
unicast, 471

PAgP (Port Aggregation Protocol), 501

parallel Sysplex, 585–588

partial URIs, 311

passive state, 552

passwords, OTPs, 641

paths, switching, 806

PAUSE frames, 492

PAWS, 295

PCI (Peripheral Component Interface), 34

PCI-X bus architecture, 35

performance metrics, 934–935
firewalls, 938

HTTP
attribute comparison, 341
compression, 342–343
version differences, 340
implications of SSL, 379–380
improving in SSL transactions, 384–385
load balancers, 671–672, 939–941
CC metric, 943
CPS metric, 942

- latency, 942–944
- PPS metric, 944
- response time, 945
- multilayer switches, 936–937
- SSL offloaders, 946
 - CPS metric, 948
 - latency, 949
 - PPS metric, 949
- testing, 950
 - hardware, 953
 - selecting data mix, 956–957
 - software, 952
- test environment, 954–955
- tools, 951
- persistence, 749
 - cookies, 728–729
 - active, 775
 - match, 771–773
 - passive, 769
 - HTTP sessions, 339, 374, 754, 757
 - redirection, 784
 - load balancers, 754, 789
 - multi-port protocols, 755–756
 - proxy servers, 758
 - clustered proxies, 759
 - session sharing servers, 761
 - source IP hash, 768
 - source IP stickiness, 765
 - mega proxies, 766–767
 - SSL, 755, 790–791
 - stickiness, 785–789
 - streaming protocols, 757
 - URL cookies, 776–778, 794–796
 - URL hash, 780–781
 - URL match, 779
- P-frames, 451, 991
- PHP, 88
- physical layers
 - 10-GigE, 495
 - Ethernet, 493
- Fast Ethernet, 494
- Gigabit Ethernet, 495
- ping of death (PoD) attacks, 163
- pipelining, 340
- PIX Firewalls
 - election process, 905
 - failure detection, 906
 - NAT, 563–564
- pixels
 - chroma subsampling, 989
 - DCT, 988
- PKCS (Public Key Cryptography Standards), 388
- PKI (public key infrastructure), 388–389, 612
 - CAs, 619
 - certificates, 621
 - deployment options, 623
 - enrollment, 624
 - key exchange, 620
 - revocation, 625
 - digital certificates, 615
 - extensions, 619
 - formats, 617
 - standards, 614
- placement
 - of DNS servers
 - forwarders, 427–428
 - split namespace, 428, 430
 - of switches in redundant Data Centers with services, 148
- plug-ins, 85
- PMTUD (path MTU discovery), 284–287
 - black-hole problem, 287–288
 - enabling on Linux, 291–292
 - enabling on Solaris 2, 291
 - enabling on Windows 2000/Windows NT, 289–290
 - enabling on Windows 95/98, 290
- point-to-point links, 829
- POP3 probes, 718
- port mappings, 105

- port remapping, 667
- port security, 183, 894
- Port VLAN IDs (PVIDs), 503
- portal applications, 72
- port-based virtual hosting, 60
- PortFast, 828, 839
- portmapper, 94
- ports
 - 802.1w, 829
 - logical ports, 517–518
 - putting into a permanent trunk, 841
 - roles and states, 511
 - switch ports, 505
- POST method (request header), 349
- PPS metric, 933
 - load balancers, 944
 - SSL offloaders, 949
- Pragma field (HTTP general header), 346
- precedence bits, 249
- predictors, 761
 - cache farm load-balancing, 683–685
 - fastest, 680
 - hash address, 681
 - least connections, 678
 - round-robin, 676
 - source IP, 681
 - URL and hash URL, 681
 - URL hash, 780–781
 - weighted least connections, 679
 - weighted round-robin, 677
- preemption, 851
- presentation tier, 77
- preventing server overload, 65–67
- printable ASCII character sets, 964–965
- private VLANs
 - in conjunction with firewalls, 187
 - security, 185
- probes, 690
 - comparing and selecting, 709
- DNS, 713
- in-band health checks, 703–705
- HTTP return code checks, 706
- server recovery, 706
- out-of-band probes, 707–708
 - application layer, 713–718
 - ICMP, 711
 - TCP, 711
 - UDP, 712
- server failure detection, 700–701
- probing, 162
- process switching, 807
- processes, 51–53
 - channels, 569
 - configuring on web servers, 57
 - multiprocess application servers, 53
- programming
 - client-side, 85
 - server-side, 87–91
- progressive playback. *See* HTTP streaming
- protocol field, 252–254
- protocols
 - ARP, 525–526
 - authentication, 640
 - control, 466
 - EIGRP, 551
 - default advertisement, 555
 - failure detection, 552
 - metric tuning, 553–554
 - redistribution, 554
 - summarization and filtering, 555
 - GLBP, 536
 - active/standby election, 537
 - failure detection, 538–539
 - load distribution, 540
 - HSRP, 528
 - failure detection, 531
 - groups, 530
 - preempt option, 529
 - tracking, 533
 - Intenet traffic patterns, 922

Layer 2, STP, 508–520
 NAT, 556–558
 application support, 559–560
 IOS NAT on routers, 561–562
 load balancers, 565–566
 PIX firewalls, 563–564
 OSPF, 541–542
 areas, 543
 default advertisement, 551
 failure detection, 545
 LSAs, 544
 metric tuning, 547
 neighbor states, 542
 redistribution, 547–549
 summarizatoin and filtering, 550
 routing, securing, 875–876
 streaming, 441–442
 VRRP
 failure detection, 535
 master/backup election, 534
 wire format, 474
 proxies, mega proxies, 766
 proxy firewalls, 175
 proxy servers
 load balancing, 760
 persistence, 758–759
 Proxy-Authorization field (HTTP header), 982
 PTR (Pointer Resource Records) records, 404, 408
 PU Type 2.1, 579
 public key encryption, 191, 379. *See also*
 asymmetric cryptography
 punt adjacencies, 808
 PUs (physical units), 571
 PUT method (request header), 350
 PVID (Port VLAN ID), 503
 PVST+ (Per VLAN Spanning-Tree Plus), 501
 convergence, 828
 rapid PVST+, 514
 configuring, 518
 VLAN support, 518

Q

QoS policies, 24
 quantization, 988
 queries (DNS)
 communication flows, 420
 resolution process, 412
 QuickTime, 460, 474
 Real Video, 451

R

RADIUS servers, 74
 Range field (HTTP request header), 355
 Rapid PVST+, 823–825
 convergence, 829–830
 rapid PVST+, 514
 configuring, 518
 RCs, 602
 RDT stream delivered on HTTP, 466
 real-time streaming, 442–444
 bandwidth, 444
 HTTP tunneling, 461
 Real-Time Streaming Protocol, 467–470
 realtime-streaming, 443
 RealVideo, 460, 474
 reassembler module, 453
 reassigned connections, 704
 receive window (TCP), 47
 records, 375
 A records, 425
 glue records, 415
 NS records, 423–424
 recursive queries, 404, 409
 recursive queries (DNS), resolution process, 417
 redirection status codes (HTTP response header), 359

- redistribution
 - EIGRP, 554
 - OSPF, 547–549
- redundancy, 448, 833
 - EtherChannels, 507
 - gateways, 849–851
 - high availability, 226–228
 - active-active load balancing, 229–230
 - active-standby load balancing, 228
 - NSF, 835–837
 - spanning trees, 842
 - supervisor redundancy, 834–835
- redundant firewall server-farm design, 905–906
- redundant Layer 2/Layer 3 designs, 139
 - access layer, 141–146
 - application architecture trends, 150–151
 - network infrastructure trends, 152–157
 - services, 146–150
- redundant links, 815–817
- Referer field (HTTP request header), 355
- referrals (DNS), resolution process, 414–417
- reflexive ACLs, 172–173
- registered informal namespaces, 321
- registries, 94
- relative URIs, 311
 - naming, 314–315
- relative URLs, 316
- reliability, 257
- Remote Network Monitoring, 699
- removing, temporal redundancy, 987, 989
- request header, 347
 - fields, 352
 - Accept field, 353
 - Accept-Charset field, 353
 - Accept-Encoding field, 354
 - Authorization field, 354
 - Host field, 354
 - If-Modified-Since field, 355
 - Max-Forwardst field, 355
 - Range field, 355
 - Referer field, 355
 - User-Agent field, 356
- methods, 348
 - CONNECT, 351
 - DELETE, 351
 - GET, 349
 - HEAD, 349
 - OPTION, 348
 - POST, 349
 - PUT, 350
 - TRACE, 351
- request URI, 351
- request URI, 351
- Rescorla, Eric, 379
- reserved characters, 318
- residual macroblock, 989
- resolving DNS names, 404–406, 411
 - caching, 420
 - client applications, 422–423
 - TTL values, 421
 - DNS proxy, 409
 - caching-only servers, 411
 - forwarders, 410
 - DNS servers, 407
 - iterative queries, 417
 - queries, 412
 - recursive queries, 417
 - referrals, 414–417
 - root hints, 413–414
- resources (HTTP), 309
 - URNs, 320
- response header
 - fields, 362–363
 - Status-Codes, 356
 - client error status codes, 360
 - informational status codes, 357
 - redirection status codes, 359
 - server error status codes, 362
 - success status codes, 358

- response time
 load balancers, 945
 SSL offloaders, 949
- retransmission, 276
- reverse proxy caching. *See* RPC
- reverse zones, 402
- RFC 1738, 315
- RFC 1918 filtering, 870
- RFC 2827 filtering, 870
- RFCs (requests for comments), 310
- RHI (Route Healt Injection), 846
- RMI, passing through firewalls, 106
- RMON (Remote Network Monitoring), 699
- root DNS servers, 407
- root hints (DNS), resolution process, 413–414
- root port (RP), 512
- root ports (RPs), 829
- root switches, setting priority, 511
- round-robin predictors, 676
- route filters, 876
- Route Health Injection (RHI), 846
- router ACLs (RAACLs), 170
- routing, 655
 between core and aggregation routers, 846, 849
 NAT, 557, 561–562
 neighbor router authentication, 876
 OSPF, 853
 passive states, 552
 process overview, 655
- routing protocol security, 875–876
- RP (root port), 512
- RPC (reverse proxy caching), 683
 optimizing server performance, 63
- RPR+, 835
- RRs (resource records), 402–403
 TTL values, 421
- RSA (Rivest, Shamir, and Adelman), ephemeral
 RSA, 631
- RTP (Real-time Transport Protocol), 454
 load balancing, 472
 payload types, 455
 QuickTime, 460
- RTSP (Real-Time Streaming Protocol), 467–470
-
- ## S
- SACK, 292–293
- SANs (storage-area networks), connecting storage devices to servers, 19
- scalability
 in Data Center design, 117
 EtherChannels, 815
 spanning-tree algorithm, 824
- scanning, 162
- scripting, 88
 ASP, 1022
 CGI, 1019
- secondary root switches, 511
- secret-key algorithms, 190
 SSL, 378
- security
 AAA, 197
 attacks
 buffer overflow, 167
 DDoS, 164
 DoS, 163
 eavesdropping, 165
 Internet infrastructure attacks, 166
 Layer 2, 167–168
 scanning/probing, 162
 session hijacking, 167
 trust exploitation, 166
 unauthorized access, 165
 viruses and worms, 165

- authentication, 640
 - AAA protocols, 645–646
 - challenge/response schemes, 642
 - digital certificates, 642
 - HTTP, 364
 - Kerberos, 644
 - OTPs, 641
- campus core, 884
- cryptography, 188–189
 - asymmetric, 602–606
 - asymmetric encryption, 191
 - CAs, 619–625
 - ciphers, 608
 - cryptographic hashing algorithms, 193–194
 - digital signatures, 195
 - export-grade ciphers, 611
 - FIPS, 609
 - hashing algorithms, 607–608
 - NIST, 609
 - PKI, 612–619
 - symmetric, 190, 597–602
- Data Center framework
 - incident response and attack mitigation, 202
 - secure management framework, 200–201
 - security life cycle, 198
 - security policies, 198
- defining security zones, 865–868
- implications for application integration, 104–107
- Internet edge, 869
 - ACLs, 873
 - antispoofting filtering, 870
 - auRPF, 872
 - IDSs, 880–882
 - Internet edge design, 882
 - securing routing protocols, 875–876
 - stateful firewalls, 878–879
- traffic rate limiting, 874
- uRPF, 873
- intranet server farms, 885–886
 - ARP inspection, 895
 - design alternatives, 896–906
 - IDSs, 891–893
 - packet filters, 890
 - port security, 894
 - stateful firewalls, 887–888
 - VLAN features, 895
- isolation of management infrastructure, 200
- management network, 908
 - authentication, 911–913
 - encryption, 910
 - isolation, 908–910
 - secure design, 914
- need overview, 159
- network management
 - SNMPv3, 649
 - SSH, 647
- network security infrastructure, 169
 - ACLS, 169–171
 - firewalls, 173–178
 - IDSs, 178–182
 - Layer 2, 183–187
- services, 25
- terminology, 160
- threats, 160
- transport security, 626
 - IPSec, 633–634, 637–639
 - SGC, 631
 - SSL, 626–629
 - SSL cipher suites, 632–633
- VLANs, 506
- VPNs, 196
- vulnerability, 161
 - out-of-date software, 161
 - software default settings, 162

- segments, 41, 245
- MSS, 44
 - small segments, 46
- SEQ numbers, 666
- sequence number field (TCP), 262
- sequence numbers, 257
- sequence states (TCP), 38
- server adapters, 33
- server applications, processes, 51–53
- server error status codes (HTTP response header), 362
- server failures, 54–55
- server farms
 - aggregation layer, 15
 - alternate Layer 2/Layer 3 designs, 133
 - ARP inspection, 895
 - creating, 749
 - design alternatives, 896–906
 - extranet server farms, 124
 - generic Layer 2/Layer 3 designs, 126–130
 - Layer 2 access switches, 130–131
 - Internet server farms, 120
 - dedicated, 120
 - DMZ server farms, 120
 - Intranet server farms, 122–124
 - load-balancing algorithms, 673–675
 - multiple-tier designs, 133–135
 - collapsed multitier design, 137–138
 - expanded multitier design, 135–136
 - port security, 894
 - security, 885–886
 - IDSs, 891–893
 - packet filters, 890
 - stateful firewalls, 887–888
 - VLAN features, 895
 - signatures, 892
 - server markdowns, 704
 - server recovery, 706
 - servers, 73, 690
 - clustering
 - geographical, 101
 - implications for application integration, 99–104
 - session persistence, 749
 - cookies, 728, 732
 - database, 96
 - DNS
 - forwarders, placement of, 427–428
 - site selectors, 431
 - split namespace, 428–430
 - dual-attached, 821
 - failure detection, 700
 - probes, 701
 - SNMP, 701
 - health management, 224, 689
 - CISCO-SLB-MIB, 698–699
 - DFP, 708
 - graceful shutdown feature, 691
 - in-band server health, 224, 703–706
 - load balancing overview, 690
 - HTTP and HTTPS (case study), 722–723
 - Max/Min Connections, 694–695
 - out-of-band server health, 225, 707–718
 - probe comparison, 709
 - slowstart feature, 693
 - SNMP, 697–698
 - virtual hosting environment (case study), 718–720
 - XML, 696–697
 - HTTP, 442
 - load balancing, 205, 209–210
 - application integration implications, 98
 - maximum connections, 682
 - multihoming, 33

sessions, 727
 persistence, 761
 session tracking, 728
 tracking, 736–740
streaming, 442
URL cookies, 776, 778
vservers, 690
web, 86

server-selection mechanism (load balancers), 653–654

server-side ActiveX, 89

server-side JavaScript, 88

server-side programming, 87–89
 ASP, 1022
 case study, 90–91
 CGI, 1018–1019
 servlets and JSP, 1021

server-specific APIs, 88

services
 Data Centers, 22
 application, 24
 business continuance, 27
 business continuance infrastructure, 26
 IP infrastructure, 23
 security, 25
 storage, 26
 web services, 151

servlet APIs, session-tracking case study, 743–748

servlets, 88, 1021

session affinity, 53

session keys, 616

session negotiation phases (SSL), 376–378

session sharing servers, 761

sessions, 727
 APPN service, 579
 hijacking, 167
 persistence, 219, 673–674, 749
 cookies, 222–223, 769
 e-commerce applications, 790

HTTP, 754, 757
multi-port protocols, 755–756

predictors, 761

proxy servers, 758–759

SSL, 755
sticky groups, 764
sticky methods, 762
resuming, 380, 382, 785

session cookies, 728–729, 769
 matching predictable strings, 773

session tracking, 728

SSL, persistence challenges, 787

tracking, 736
 Apache mod_session (case study), 740–741
 combining methods, 740
 cookies, 731, 739
 form hidden fields, 737
 HTTP sessions with servlets (case study), 743–748
 URL rewriting, 738

SGC, 631

SHA (Secure Hash Algorithm), 608

shared links, 829

short-lived traffic, 925–927
 performance metrics, 933
 SSL connections, 947

show cdp command, output, 501

show spanning-tree vlan 10 command, 516

show spanning-tree vlan command, 514

signature-based IDSs, 181

signatures, 105, 181
 digital, 195
 IDSs, 107, 891
 Internet edge IDSs, 881
 Solaris, 894
 Windows, 893

- site selection architecture, 430–433
 - caching, 436–437
 - proximity, 435
 - referrals to site selectors, 433–435
 - stickiness, 437–438
- size (Ethernet), 487–488
- slave name servers, 418
- slices, 991
- sliding windows, 277
- slow paths, 933
- slow start, 279
- slowstart feature, 693
- small segments, 46
- SMTP probes, 718
- smurf attacks, 163
- SNA (Systems Network Architecture), 577
 - APPN, 579–580
 - over TCP/IP, 580
 - DLSw, 580–581
 - SNAsw, 581–585
 - subnetwork SNA, 577–579
 - VTAM, 571
- SNAsw (SNA switching), 581
 - BXN, 583
 - DLUR/DLUS, 583
 - EES, 582–583
 - TN3270, 584–585
- SNMP (Simple Network Management Protocol)
 - Management Stations, 697
 - OIDs, 697–698
 - server failure detection, 700–701
 - TRAPs, 700
- SNMPv3, 649
- SOA (Start of Authority) records, 403
- sockets, 39–40
- software
 - clustering, 100
 - default settings (security risk), 162
 - load balancing, 98
- middleware, 76, 91–92
 - components, 93
 - traffic patterns, 94–95
- out-of-date (security risk), 161
- performance metric testing, 952
- Solaris signatures, 894
- Solaris 2, enabling PMTUD, 291
- source IP hash, 768
- source IP predictor, 681
- source IP stickiness, 221, 765–76, 792
- spanning trees, 822, 841–843
 - client-side VLANs, 826
 - selecting algorithms, 823–825
- spatial redundancy, 448
 - removing, 987–989
- specifications, MIME, 323
- speed negotiate command, 492
- split namespace, 428–430
- splitting (stream), 471
- spoofing
 - ARP, 167
 - connection spoofing, 664–667
- SQL (Structured Query Language), 96
- SQL Slammer, 165
- SSH, 647
- SSL (Secure Sockets Layer), 626
 - authentication, 385–387
 - PKI, 388–389
 - certificates, 629
 - ciphersuites, 371, 389–390, 632–633
 - client authentication, 642
 - connections, 371–372
 - connection, 371
 - data encryption, 378
 - example applications of, 370–371
 - handshakes, 374–378
 - HTTPS, 372–374
 - load balancing, 382–384

- offloading, 794–796
 - CPS metric, 948
 - latency, 949
 - performance metrics, 946
 - PPS metric, 949
- performance, 379–380
 - optimizing, 384–385
- persistence, 755, 791
- probes, 715
- secret keys, 378
- sessions, 372, 380–382
- stickiness, 785
 - challenges and concerns, 787–788
 - configuring, 786
- TCP/IP layers, 627–628
- traces, analyzing, 391–393
- VPNs, 639
- SSLv2, 627
- SSLv3, TLS 1.0, 627
- SSO, 835
- standalone servers, processes, 52
- standard ACLs, 170
- standard retransmission, 446
- standards
 - cookies, 735
 - PKI, 614
- stateful devices, 803
- stateful failover, 227, 231
- stateful firewalls, 175, 878–879
 - intranet server farms, 887–888
- stateless failover, 231
- static routing, 527
- Status-Codes (HTTP response header), 56, 356, 983–985
 - client error status codes, 360
 - informational status codes, 357
 - redirection status codes, 359
 - server error status codes, 362
 - success status codes, 358
- Step-Up, 631
- stickiness, 219, 674
 - in site-selection architecture, 437–438
- sticky failover, 231
- sticky groups, 764
- sticky methods, 761–762
- sticky tables, 221
- storage layer, 19
- storage services, 26
- storing cookies, 734–735
- STP (Spanning-Tree Protocol), 508
 - 802.1s configuration, 519–520
 - bridge identifiers, 510
 - convergence, 827
 - failure detection, 513
 - logical ports, 517–518
 - loop prevention, 832–833
 - multiple VLANs, 513–517
 - port roles and states, 510–512
 - rapid PVST+ configuration, 518
 - versions, 509
- stream splitting, 471
- streaming, 441–442
 - applications, session persistence, 757
 - congestion, 463
 - download rate, 466
 - HTTP tunneling, 466
 - real-time streaming, 443
 - RTSP, 467, 469
 - selecting protocol, 445
 - servers, 74, 442
 - packetizer module, 453
 - unicast/multicast packets, 471
 - software products, 473
 - streaming rate, 466
 - TCP, 462
 - transport formats, 454
 - RTCP, 457–459
 - RTP, 454

UDP, 464–465
 video, 447
 stub areas, 543, 854
 stub resolver (DNS), 405
 subdomains, 398
 subnetwork SNA, 577–579
 success status codes (HTTP response header), 358
 summarization
 EIGRP, 555, 860
 OSPF, 550, 853
 supervisor redundancy, 834–835
 suppressing broadcasts, 487
 SVIs, VLANs, 813
 switch fabric, 233–234
 switch ports, 505
 switching
 debounce feature, 831
 Ethernet, 498–499
 frame size support, 497
 MAC address table, 500
 failure detection, 513
 Layer 3, 807
 multilayer, performance metrics, 936–937
 operation overview, 654
 root, setting priority, 511
 switching paths, 806, 933
 Cisco IOS, 807–808
 MLS, 809
 switchport mode trunk command, 506, 841
 switchport trunk allowed vlan 10,20 command, 506
 switchport trunk encapsulation dot1q command, 506
 symmetric cryptography, 597
 3DES, 600
 DES, 598–600
 RCs, 602
 symmetric encryption, 190
 SYN floods, 163
 SYN retransmission mechanism, 55
 Sysplex, 585–589
 system jumbomtu command, 497

T

tables
 ARP, 526
 CAM, 526
 TACACS+, 645
 tagging traffic, 504
 TCN (Topology Change Notification) BPDUs, 527
 TCP (Transport Control Protocol), 256, 461
 ACKs, 48
 applications, 41
 HTTP, 47
 Telnet, 43–46
 configuring on web servers, 57
 connections, 257, 267
 establishment, 268–270
 termination, 272–275
 data processing, 41
 flow control, 257
 congestion avoidance, 279
 congestion control, 278
 delayed ACKs, 280
 fast recovery, 280
 immediate ACKs, 280
 Nagle algorithm, 281–282
 retransmission, 276
 sliding windows, 277
 slow start, 279
 half close, 282
 header compression, 296–298
 header fields, 258–259
 acknowledgment number field, 263
 checksum field, 266
 control flags, 264–266
 options field, 266–267
 sequence number field, 262
 TCP header length field, 264
 urgent pointer field, 266
 window size field, 266

- keepalives, 55
 - maximum burst size on high-speed networks, 49–50
 - monitoring connections, 67
 - MSS, 283–284
 - multiplexing, 257
 - offloading, 33
 - PAWS, 295
 - PMTUD, 284–287
 - black-hole problem, 287–288
 - enabling on Linux, 291–292
 - enabling on Solaris 2, 291
 - enabling on Windows 2000/Windows NT, 289–290
 - enabling on Windows 95/Windows 98, 290
 - probes, 711
 - Real Player, 463
 - reliability, 257
 - retransmission, 44
 - SACK, 292–293
 - segments, 42
 - sequence numbers, 257
 - server failure handling, 54
 - SYN retransmission, 55
 - TCP timeouts, 54
 - server failures, 54
 - streaming, 462
 - timestamps, 294
 - versus UDP, 445–446
 - well-known port numbers, 260–261
 - window scale, 295
 - windows, 47–50
- TCP/IP protocol suite, 243
 - client/server architectures, 37–39
 - TE field (HTTP header), 982
 - Telnet
 - connection establishment, 43
 - connection termination, 46
 - delayed ACKs, 45
- interactive traffic, 41–43
 - MSS, 44
 - Nagle algorithm, 46
 - TCP retransmission, 44
 - temporal redundancy, 448
 - removing, 987–989
 - temporary cookies, 729
 - terminating TCP connections, 272, 275
 - testing performance metrics, 950
 - hardware, 953
 - selecting data mix, 956–957
 - software, 952
 - test environment, 954–955
 - tools, 951
 - thick clients, 9, 83
 - thin clients, 83
 - threaded servers versus forking servers, 51, 53
 - threats (security), 160
 - three-way handshakes, 268
 - thresholds, reassigning connections, 705
 - throughput, multilayer switch metrics, 936
 - timestamps, 294
 - TLDs (top-level domains), 399
 - TN3270 servers, 74, 584–585
 - topologies
 - Data Center architecture, 13–14
 - access layer, 16–18
 - aggregation layer, 15
 - layers, 14
 - storage layer, 19
 - transport layer, 20–21
 - EIGRP, 859
 - fully switched, 804
 - Layer 2, 818
 - minimizing changes, 831
 - OSPF, 852
 - redundant Layer 2/Layer 3 designs, 139
 - access layer, 141–146
 - application architecture trends, 150–151

- network infrastructure trends, 152–157
- services, 146–150
- VLANs, 804
- TOS field, 248–250
- total length field, 250
- totally stubby area, 543
- TRACE method (request header), 351
- traceroute, 252
- tracking
 - server health, 224
 - in-band server health tracking, 224
 - out-of-band server health tracking, 225
- user sessions, 736
 - Apache mod_session (case study), 740–741
 - combining methods, 740
 - cookies, 739
 - form hidden fields, 737
 - HTTP sessions with servlets (case study), 743–748
 - URL rewriting, 738
- traffic
 - channeling, 507
 - client NAT, 663
 - encoding formats, 450–451
 - Internet, HTTP, 328
 - load balancing
 - architecture, 232–235
 - cache load balancing, 210–211
 - connection failover, 231
 - connection persistence, 219
 - connection tracking, 219
 - firewall load balancing, 212–213
 - flexibility, 659
 - high availability, 226–230
 - implications for application integration, 97–98
 - Layer 4 load balancing, 216
 - Layer 5 load balancing, 217
 - process description, 215–216
 - server health, 224–225
 - server load balancing, 209–210
 - session persistence, 219, 222–223
 - stateful failover, 231
 - stateless failover, 231
 - sticky failover, 231
 - VPN/IPSec load balancing, 211
- multimedia transport formats, 454
 - RTCP, 457–459
 - RTP, 454
 - packetization, 453
- patterns, 919–920
 - Data Centers, 924–933
 - Internet, 920–921
 - intranets, 923
 - load balancers, 939
 - protocols, 922
 - rate limiting, 874
 - SSL, load balancing, 382–384
 - switching paths, 806
 - tagging, ISL, 503
 - transport rate, 448
- traffic mix, 920
- Trailer field (HTTP header), 977
- Transfer-Encoding field (HTTP general header), 347
- transactions
 - middleware, 91
 - UDP, 301–302, 305
- transceivers, 493
 - Fast Ethernet, 495
- Transfer-Encoding headers, 343
- transparent caching, 684
- transparent devices, 824–825
- transport layer (Data Centers), 20–21
- transport protocols, UDP system calls, 40
- transport rate, 448
- transport security, 626
 - IPSec, 633
 - IKE, 637
 - security parameters, 638–639
 - TCP/IP layers, 634

- SGC, 631
 - SSL, 626
 - certificates, 629
 - cipher suites, 632–633
 - TCP/IP layers, 628
 - TRAPs, 700
 - troubleshooting
 - DoS attacks, traffic rate limiting, 874
 - Ethernet networks, frame size issues, 488
 - firewall limitations, 178
 - flooding, 831
 - loops, 832–833
 - server failure detection, 700, 704
 - probes, 701
 - SNMP, 701
 - STP, failure detection, 513
 - trunks, 503
 - configuring, 840
 - creating, 505–506
 - TTL field, 251–252, 421
 - TPP response header, Status-Codes, 359
-
- ## U
-
- UDLD (Unidirectional Link Detection), 501, 832–833
 - UDP (User Datagram Protocol), 50–51, 299, 461.
 - See also* TCP
 - header compression, 305
 - header fields, 299–301
 - probes, 712
 - server failure handling, 54–55
 - server failures, 54
 - streaming, 464–465
 - system calls, 40
 - transactions, 301–302, 305
 - versus TCP, 445–446
 - unicast, 499
 - unicast flooding, 499
 - unicast MAC addresses, dummy, 98
 - unicast packets, 471
 - Uniform Record Locators. *See URLs*
 - Uniform Resource Identifiers. *See URIs*
 - Universal Resource Names. *See URNs*
 - UNIX, system calls, 39–40
 - unsafe characters, 318
 - Upgrade field (HTTP header), 978
 - upgrading applications, 71
 - UplinkFast, 827–828
 - urgent pointer field (TCP), 266
 - URIs (Uniform Resource Identifiers), 310
 - absolute/full, 312
 - naming rules, 314–315
 - relative/partial, 311
 - request URI, 351
 - URNs and URLs, 322
 - URL match, 779
 - URLs (Uniform Record Locators), 311, 315
 - cookies, 776–778, 794–796
 - encoding, 316
 - hashing, 780–781
 - relative and absolute, 316
 - reserved characters, 318
 - rewriting, 738, 776
 - schemes, 316, 319
 - stickiness, 776
 - unsafe characters, 318
 - URIs and URNs, 322
 - URNs (Universal Resource Names), 311, 320
 - encoding, 320
 - namespace, 321
 - URIs and URLs, 322
 - uRPF, 872–873
 - user agents, 84
 - browsers, 84
 - client-side programming, 85
 - helpers and plug-ins, 85
 - user mode, 35–36
 - User-Agent field (HTTP request header), 356

V

valuation, 197
 Version field, 247
 vertical scaling, 206
 Via field (HTTP header), 978
 video encoding, 987
 video on demand (VoD), 445
 video streaming, 442, 447
 codecs
 analog, 448
 comparison of, 452
 MPEG, 990–991
 popular encoding formats, 450–451
 removing spatial and temporal redundancy, 987–989
 slices, 991
 redundancy, 448
 transport rate, 452
 transport formats, 454
 RTCP, 457–459
 RTP, 454
 VIPAs (virtual IP addresses), DVIPAs, 587–588
 virtual hosting
 configuring on web servers, 58–59
 IP-based, 59
 name-baseds, 61
 port-based, 60
 server health (case study), 718–720
 virtual servers (virtual servers), 690
 viruses, 165
 VLAN ACLs (VACLs), 170
 vlan dot1q tag native command, 506
 VLANs, 170, 502, 802
 4096 VLANs, 514
 802.1s, 516
 access ports, 520
 autostate, 814
 designing, 505–506
 PVIDs, 503

SVIs, 813
 topologies, 804
 trunks, 503
 virtualizing Data Center infrastructures, 810
 VoD (video on demand), 445
 VPN/IPSec load balancing, 211
 VPNs (Virtual Private Networks)
 IPSec versus SSL, 639
 security, 196
 VRRP, 527
 failure detection, 535
 master/backup election, 534
 vservers (virtual servers), 690
 VTAM (virtual telecommunications access method), 571
 VTP (VLAN Trunking Protocol), 500, 504
 domains, defining, 504
 modes, 839

W

W3C (World Wide Web Consortium), 151
 Warning field (HTTP header), 978
 WCCP (Web Cache Control Protocol), 685
 Web Cache Control Protocol (WCCP), 685
 web servers, 57, 86
 directories, 58
 HTTP applications, 55–56
 inserting cookies, 1010
 server processes, configuring, 57
 TCP parameters, configuring, 57
 virtual hosting, configuring, 58–61
 web services, 151
 weighted least connections predictor, 679
 weighted round-robin predictors, 677
 well-known port numbers, 260–261
 window scale, 295
 window size field (TCP), 266
 windows, BDP, 50

windows (TCP), 47–48
Windows 95, enabling PMTUD, 290
Windows 98, enabling PMTUD, 290
Windows 2000
 configuring loopback interfaces, 996–998
 enabling PMTUD, 289–290
Windows Media Video, 451, 461
Windows NT
 configuring loopback interfaces, 1002
 enabling PMTUD, 289–290
wire format, 474
WML (Wireless Markup Language), 83
worms, 165
WSA (Web Services Architecture), 21

X–Z

XML (Extensible Markup Language), 79–83,
696–697

zones (DNS), 400–402
 name servers, 418
 zone transfers, 418–420