



Getting Started with Networking, Wireless, and Security

FOR DUMMIES[®]

Learn

- Types of networks and how they process data
- What Ethernet protocols are and how different devices support them
- Ways to make your network better support evolving data center needs
- Strategies for creating a wireless network that is efficient and secure



Brian Underdahl



***Getting Started with
Networking, Wireless,
and Security***

FOR
DUMMIES®

by Brian Underdahl



WILEY

John Wiley & Sons, Inc.

These materials are the copyright of John Wiley & Sons, Inc. and any dissemination, distribution, or unauthorized use is strictly prohibited.

Getting Started with Networking, Wireless, and Security For Dummies®

Published by
John Wiley & Sons, Inc.
111 River Street
Hoboken, NJ 07030-5774
www.wiley.com

Copyright © 2011 by John Wiley & Sons, Inc., Hoboken, New Jersey

Published by John Wiley & Sons, Inc., Hoboken, New Jersey

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

Trademarks: Wiley, the Wiley logo, For Dummies, the Dummies Man logo, A Reference for the Rest of Us!, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, please contact our Business Development Department in the U.S. at 317-572-3205. For details on how to create a custom *For Dummies* book for your business or organization, contact info@dummies.biz. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

ISBN: 978-1-118-12343-0

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1



These materials are the copyright of John Wiley & Sons, Inc. and any dissemination, distribution, or unauthorized use is strictly prohibited.

Publisher's Acknowledgments

We're proud of this book and of the people who worked on it. For details on how to create a custom *For Dummies* book for your business or organization, contact info@dummies.biz. For details on licensing the *For Dummies* brand for products or services, contact BrandedRights&Licenses@Wiley.com.

Some of the people who helped bring this book to market include the following:

Acquisitions, Editorial, and Media Development

Project Editor: Carrie A. Burchfield

Editorial Manager: Rev Mengle

Business Development Representative:
Kimberley Schumacker

Custom Publishing Project Specialist:
Michael Sullivan

Composition Services

Senior Project Coordinator: Kristie Rees

Layout and Graphics: Carl Byers,
Carrie A. Cesavice, Joyce Haughey,
Laura Westhuis

Proofreader: Jessica Kramer

Special Help: Lyndon (LJ) Miller,
Liz Stine

Publishing and Editorial for Technology Dummies

Richard Swadley, Vice President and Executive Group Publisher

Andy Cummings, Vice President and Publisher

Mary Bednarek, Executive Director, Acquisitions

Mary C. Corder, Editorial Director

Publishing and Editorial for Consumer Dummies

Kathleen Nebenhaus, Vice President and Executive Publisher

Composition Services

Debbie Stailey, Director of Composition Services

Business Development

Lisa Coleman, Director, New Market and Brand Development

Table of Contents

.....

Introduction	1
About This Book	1
How This Book Is Organized	1
Icons Used in This Book.....	2
Chapter 1: Understanding Networks	3
Knowing Why Networking Matters	3
Getting to Know the Network Types	4
Figuring out the Requirements of Networking	5
High availability	5
Scalability	8
Manageability	8
Taking a Look at the OSI Networking Model	8
Layers and functionalities	8
Following a packet through the layers.....	10
Chapter 2: Understanding Network Protocols and Components	11
The Ethernet Protocol.....	11
Evolution of Ethernet	12
Data Center Bridging.....	12
Network devices	13
Switches and hubs	13
Routers	13
Storage area networks	14
Standards versus Proprietary Protocols	15
Chapter 3: Looking into Next Generation Data Centers (NGDC)	17
Understanding Today's Data Centers	17
Identifying Opportunities for Networking in the NGDC	18
Bringing in New Protocols and Standards	18
Flattening Your Network.....	19
Stepping into Network Virtualization	20



Chapter 4: Identifying Network Endpoints	21
Serving as the Network Centers: Servers.....	21
Accessing the Network with NICs	22
Using a Host Bus Adapter	22
What Is an Array?.....	23
Chapter 5: Security 101.	25
Physical Security: Locking Your Doors.....	25
Securing User Accounts	25
Hardening Your Network.....	26
Using a firewall.....	26
Disabling unnecessary services.....	26
Patching your servers.....	26
Using network appliances.....	27
Securing Your Users	27
Understanding Antivirus Programs	27
Firewalls — Network Traffic Cops	29
Looking at the Basic Types of Firewalls.....	29
Packet filtering	30
Stateful packet inspection (SPI).....	30
Deep packet inspection.....	30
Chapter 6: Setting Up and Securing a Wireless Network.	31
Diving into Wireless Networking.....	31
Wireless Access Points	32
Infrastructure mode	33
Roaming	33
Wireless bridging.....	34
Ad-hoc networks.....	34
Securing Your Wireless Network.....	34
Using MAC address filtering.....	35
Role-based access.....	36
Using encryption.....	36
Chapter 7: Ten (Okay, Nine) Networking Strategies to Consider	37
What Type of Network Topology Should I Consider?	37
What's the Best Way to Connect Workstations to My Existing Network?.....	38

What Types of Products Should I Consider?.....	38
What Can I Do about Viruses and Poor Performance?.....	38
What Should I Look for to Secure My Wireless Network?....	39
How Can I Detect and Contain Rogue APs and Other Wireless Threats?.....	39
What Considerations Help Me Plan for Future Expansion?.....	39
What Do I Need to Ask?.....	40
How Can I Identify the Actual Applications Consuming Bandwidth?.....	40
Case Study A: Aruba’s Approach: Taking the Campus Wireless	41
Analyzing the Situation	41
Understanding the Requirements.....	42
Making a Business Case for Wireless	43
WLAN Deployment Details.....	43
Broadcast Television over 802.11n.....	44
Case Study B: F5’s Approach: Achieving Application Availability	45
Understanding the Challenge	45
Defining a Solution	46
Realizing the Benefits	47
Zero downtime for e-mail.....	48
Superior support.....	48
Cost-effective platform for growth	48
BIG-IP LTM special benefits.....	49
Plan for growth and avoid downtime	49
Accelerate your applications up to 3x	49
Secure your applications and data	49
Reduce servers, bandwidth, and management costs.....	49
Take control over application delivery.....	50
Case Study C: Dell’s Approach: Reducing Costs and Streamlining Administration.	51
Facing the Challenges.....	51
Accelerating the Network	52
Increasing Network Performance	53

Case Study D: NetScout's Approach: Managing Performance	55
Understanding J-Flow	55
Getting to Know nGenius Service Assurance Solution for J-Flow	56
Analyzing J-Flow	57
Identifying Complex Applications from J-Flow Conversations	58
Benefits to the nGenius Service Assurance Solution for J-Flow	59
Case Study E: Riverbed Technology, Inc.'s Approach: Making the Network Deliver	61
Defining the Problem	61
Examining the Options	62
Creating a Solution	63
Gauging the Benefits	63
Drawing a Conclusion	64

Introduction

Are you trying to figure out how to set up a network that handles the needs of your business? Or are you trying to get up to speed so that you have some idea what your networking consultant is recommending? If so, this book is designed to help.

About This Book

Getting Started with Networking, Wireless, and Security For Dummies, shows you the basics of business networking along with case studies that show how Dell and its partners have been able to help people just like you to implement solutions that met the needs, saved the day, and made the world a better place.

How This Book Is Organized

This book is divided into seven chapters. Chapter 1 shows you the basics of networks and introduces a standard network model. In Chapter 2 you find out about the different flavors of Ethernet and the meat (or more accurately, the hardware) of networking. Here you discover pieces of stuff that you can actually pick up and handle, and you see what they do for you.

Chapter 3 gets into data centers. If you jump into Chapter 4, you get information on network endpoints.

In Chapter 5, you see important ways to keep your network secure. This chapter also shows you what firewalls do and also explains the different types of firewalls that are available. Wireless networks are covered in Chapter 6. You get the basics of wireless networking and see what it takes to make a wireless network secure. And Chapter 7 gives you the top ten (okay, nine) things to consider when planning your business network.

2 Getting Started with Networking, Wireless, and Security For Dummies

We also include some case studies in this book for you to see real-world applications. Head to the case studies in the back of this book for more information.

Icons Used in This Book

This book uses the following icons to call your attention to information that you may find helpful in particular ways.



The information in paragraphs marked by the Remember icon is important and therefore repeated for emphasis. This way, you can easily spot the information when you refer to the book later.



The Tip icon indicates extra-helpful information.



This icon marks places where technical matters, such as pixels and whatnot, are discussed. Sorry, it can't be helped, but it's intended to be helpful.



Paragraphs marked with the Warning icon call attention to common pitfalls that you may encounter.

Chapter 1

Understanding Networks

.....

In This Chapter

- ▶ Discovering why networks exist
 - ▶ Understanding the various types of networks
 - ▶ Looking at the requirements of network connections
 - ▶ Viewing the OSI networking model
-

Rumor has it that the first computer network was invented when ancient mathematicians connected their abacuses (or is it *abaci*?) together with kite string so they could instantly share their abacus answers with each other. Over the years, computer networks became more and more sophisticated. Now, instead of string, networks use electrical cables, fiber-optic cables, or wireless radio signals to connect computers to each other. The purpose, however, has remained the same: sharing information and getting work done faster.

This chapter describes the basics of what computer networking is and how it works.

Knowing Why Networking Matters

In today's fast-paced world, businesses need reliable and efficient access to information in order to be competitive. Rather than relying on mounds of paper stored in file cabinets, businesses rely on computers to store and manage information

electronically. This electronic storage means that workers throughout an organization can have instant access to the information they need.

Computer networking enables this access and provides the operational efficiencies organizations need. In fact, the largest network on the planet, the Internet, enables businesses to share information not only within the organization but also with customers worldwide.

These days, networking is everywhere and you may not always realize you're using it. In addition to the Internet, people use point-of-sale systems when they make a retail purchase in a store, they view flight information on network screens in an airport, and they search for titles on the terminals in a bookstore. In each of these cases, networking provides shared access to important information.

Getting to Know the Network Types

Networks are often classified as either a local area network (LAN) or as a wide area network (WAN). The two are different:

- ✔ LANs connect many devices that are relatively close to each other, such as in the same building. For example, library terminals that display book information connect over a LAN.
- ✔ WANs connect a smaller number of devices that can be a great distance apart. For example, if two libraries at the opposite ends of a city wanted to share their book catalog information, they most likely make use of a WAN.

Improvements in technology continue to blur the line between LANs and WANs. For example, fiber optic cables have allowed LAN technologies to connect devices over longer distances, while at the same time greatly improving the speed and reliability of WANs.

While LAN and WAN are by far the most popular network types mentioned, you may also commonly see references to these other types of Networks:

- ✔ **Wireless Local Area Network (WLAN):** A LAN based on WiFi (wireless fidelity) network technology
- ✔ **Metropolitan Area Network (MAN):** A network spanning a physical area larger than a LAN but smaller than a WAN, such as a city

A MAN usually interconnects a number of local area networks (LANs) by using a high-capacity backbone technology, such as fiber-optical links, and provides up-link services to wide area networks. A MAN is typically owned and operated by a single entity, such as a government body or large corporation.
- ✔ **Campus Area Network (CAN):** A network spanning multiple LANs but smaller than a MAN, such as on a university or local business campus
- ✔ **Storage Area Network (SAN):** A network used to connect servers to data storage devices through technologies such as Fibre Channel or iSCSI

Figuring out the Requirements of Networking

The more businesses rely on their networks, the more importance they place on several requirements. In the following sections, you look at some of these important requirements in today's networks.

High availability

One critical requirement of a network is often high availability. This requirement translates into a network that's accessible at all times. High availability is critical when your network hosts mission-critical applications. For example, an ATM network must be always available or a bank can suffer a tremendous amount of financial loss.

High availability can prevent financial loss, improve productivity, and improve customer satisfaction. High availability must be built into multiple layers of the network in order to prevent different types of failures. For example, here are some of the factors which can affect high-availability:

- ✔ **Layer 1: Physical Layer:** Redundant links and switching hardware can be used to provide an application with multiple paths for traveling across the network so if one path becomes unavailable, there's another path that the application can use.
- ✔ **Layers 2 and 3: Data-Link and IP Layer:** Ethernet protocols, such as Spanning Tree Protocol (STP), Hot Standby Routing Protocol (HSRP), Transparent Interconnect of Lots of Links (TRILL), and Shortest Path Bridging (SPB), provide built-in redundancy in an Ethernet environment.
- ✔ **Layers 4 through 6: Transport, Session, and Presentation Layers:** These layers are more and more prevalent because, with all the data being transported, optimization and cacheing really help data flow in these layers.
- ✔ **Layer 7: Application Layer:** Mission-critical applications should be deployed in a redundant fashion and support fast failover in order to avoid downtime.

Redundancy and high availability are partners in network design. Redundancy is a key part of building and designing highly available networks and can be realized within hardware at the link level, switch level, and physical appliance level — or even within device software or applications.

Figure 1-1 shows an example of hardware high availability, which is a switch that has two supervisor modules. If one module fails, traffic is still forwarded through the second module.

Figure 1-2 shows how multiple links between switches help provide high availability because a single failed link doesn't disrupt network traffic.

A mesh network, as shown in Figure 1-3, provides many different redundant paths and therefore high availability.




Figure 1-1: A switch with dual supervisors.




Figure 1-2: Link-level availability between network devices.




Figure 1-3: A full mesh network.

Scalability

Networks need *scalability* — the ability to grow to meet increasing business demands. Room for growth is needed so your network meets your future needs. Typically, experts recommend that your design allows for at least 20 percent growth. This added capacity enables your network to handle additional applications and traffic without accelerating the hardware refresh cycle.

Manageability

Manageability is a key network requirement because poor manageability means that your network may never be functioning as efficiently as you would like. Manageability includes everything from initial deployment, configuration, monitoring, and troubleshooting.



Each of these areas is vital in making sure that applications are deployed in a timely manner and that your network functions properly.

Taking a Look at the OSI Networking Model

Imagine how chaotic your morning commute would be if everyone decided for themselves which direction to travel on the freeway, what the different colors on traffic lights meant, and which road lanes were for parking. Fortunately, traffic laws do exist to regulate those things so everyone knows what to expect. Just like road traffic networks, computer networks also have rules to control traffic, and this section provides a brief introduction to those rules.

Layers and functionalities

Data flows through the layers of the network stack in an orderly fashion. Here's a brief explanation of the layers:

- ✓ **Physical layer:** Defines the electrical and physical specifications for devices

- ✔ **Data link layer:** Provides the means to transfer data between network entities
- ✔ **Network layer:** Performs network routing functions
- ✔ **Transport layer:** Controls the reliability of a given link
- ✔ **Session layer:** Controls the connections between computers
- ✔ **Presentation layer:** Establishes a context between Application Layer entities
- ✔ **Application layer:** Interacts with software applications that communicate

Figure 1-4 shows how a packet of information flows through the seven layers as it travel from one computer to another on the network.

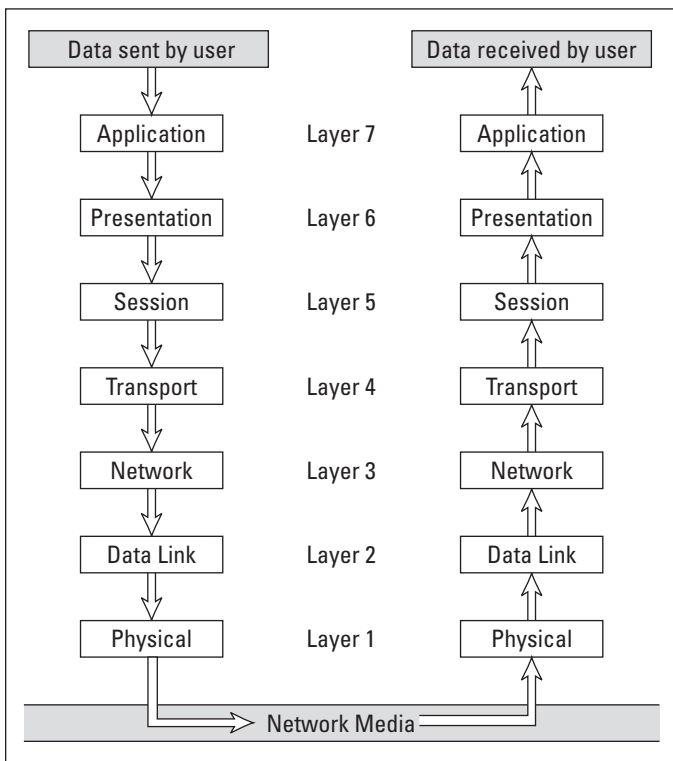


Figure 1-4: How data travels through the seven layers.

Following a packet through the layers

Packets follow a well-defined path through the layers. Here is the process:

- 1. The data begins its journey when an end-user application sends data to another network computer.**

The data enters the network through an Application layer interface.

- 2. The data then works its way down through the protocol stack.**

Along the way, the protocol at each layer manipulates the data by adding header information, converting the data into different formats, combining packets to form larger packets, and so on.

- 3. When the data reaches the Physical layer protocol, it's actually placed on the network media (in other words, the cable) and sent to the receiving computer.**

- 4. When the receiving computer receives the data, the data works its way up through the protocol stack.**

- 5. The protocol at each layer reverses the processing that was done by the corresponding layer on the sending computer.**

Headers are removed, data is converted back to its original format, packets that were split into smaller packets are recombined into larger messages, and so on.

- 6. When the packet reaches the Application layer protocol, it's delivered to an application that can process the data.**

Chapter 2

Understanding Network Protocols and Components

.....

In This Chapter

- ▶ Deciphering the Ethernet
 - ▶ Understanding network components
-

Networks use various protocols and components in order to function. This chapter briefly discusses those items.

The Ethernet Protocol

Ethernet has been around in various forms since the early 1970s when it was developed by Robert Metcalf at Xerox PARC. Since then, it has become the standard for network interconnection in today's LANs. Ethernet both defines standards for physical connection at Layer 1, as well as a common addressing and forwarding component at Layer 2 of the OSI model. (More on the OSI model in Chapter 1.)

The current incarnation of Ethernet is defined by the IEEE standard known as 802.3. Various flavors of Ethernet operate at different speeds and use different types of media. However, all the versions of Ethernet are compatible with each other, so you can mix and match them on the same network by using devices — such as switches and routers — to link network segments that use different types of media.



The actual transmission speed of Ethernet is measured in millions or billions of bits per second, or Mbps. Ethernet comes in several different speed versions. Keep in mind, however, that network transmission speed refers to the maximum

speed that can be achieved over the network under ideal conditions. In reality, the actual throughput of an Ethernet network rarely reaches this maximum speed.

Ethernet operates at the first two layers of the OSI model — the Physical and the Data Link layers. However, Ethernet divides the Data Link layer into two separate layers known as the *Logical Link Control (LLC) layer* and the *Medium Access Control (MAC) layer*. The two are generally considered sublayers rather than full layers. The LLC provides multiplexing and flow control, while the MAC provides addressing and channel access control.

Evolution of Ethernet

The demands of networking are constantly changing. Because of these changing demands, network designs, technologies, and protocols must also evolve. Take a look at some of the recent Ethernet enhancements:

- ✓ **Data Center Bridging:** Enhancements to LANs for use in data centers
- ✓ **Network devices:** Switches, hubs, and routers designed for higher performance
- ✓ **Storage Area Networks:** High-performance networks that enable storage devices to communicate

Data Center Bridging

Data Center Bridging (DCB) refers to enhancements to Ethernet local area networks for use in data center environments. Ethernet is designed to be a best-effort network that may drop packets when the network or devices are busy. Transport reliability was traditionally the responsibility of the Layer 4 transport protocols, such as the Transmission Control Protocol (TCP), with the trade-off being higher complexity, greater processing overhead and the resulting impact on performance and throughput.

Currently there are ongoing efforts to add extensions to the existing Ethernet protocol suite to provide reliability without incurring the penalties of TCP. A desire also exists for higher

granularity in control of bandwidth allocation and to ensure that bandwidth is used more effectively to make Ethernet a more viable transport for storage and server cluster traffic.

To meet these goals, new standards are being developed that either extend the existing set of Ethernet protocols or emulate the connectivity offered by Ethernet protocols. They're being developed respectively by two separate standards bodies: Institute of Electrical and Electronics Engineers (IEEE), and IETF. More information on DCB can be found in Chapter 3.

Network devices

Various networking devices, such as switches and routers, are the building blocks of any network. Network traffic often passes through multiple different networking devices as it travels along the network from one end point to another.

Switches and hubs

At first glance, hubs and switches may seem to be very similar. Originally, hubs were just about the only option for connecting devices to the network. Today, however, hubs are seldom used because of their limited functionality. Hubs operate at the physical layer (layer 1) of the OSI model and lack the intelligence to identify specific hosts on the network. Because hubs can't identify the target destination, every packet is forwarded to every endpoint with the end result being additional network traffic and poor performance.

Switches operate at the data link layer (layer 2) which means that a switch can identify the target destination and send the packet to that destination. By sending the packet to the correct destination network traffic is reduced, performance is enhanced, and security is improved because endpoints only see the packets that are intended for them.

Routers

Routers are advanced networking components that can divide a single network into two logically separate networks. They can

- ✓ Expand on the intelligence of switches by acting at the Network Layer (Layer 3) of the OSI model
- ✓ Examine the IP address of the packets that pass through it

- ✔ Determine the origin and target network of a message because IP addresses have both a network and a host address
- ✔ Divide a single network into two logically separate networks
- ✔ Form a logical boundary for the network, which prevents Ethernet broadcasts from crossing
- ✔ Operate based on protocols that are independent of the specific networking technology, which allows routers to easily interconnect various network technologies, both local and wide area

Storage area networks

Storage area networks (SANs) are high-performance networks that provide access to storage devices, such as disk arrays, tape libraries, and optical jukeboxes — devices that provide online access to any of a number of optical discs, such as CDR or DVD-R discs.

SANs use a variety of different protocols and technologies. The most commonly used protocols are SCSI, Fibre Channel (FC), Fibre Channel over Ethernet (FCoE) and iSCSI. Depending on which protocol is deployed, SANs can operate at both Layer 2 (Fibre Channel and FCoE) as well as Layer 3 (iSCSI). More information on these protocols can be found in Chapter 3.



Fibre Channel doesn't exactly follow the OSI model. FC uses a 5-layer model that's very similar to the OSI layers and is commonly described as a "layer 2 protocol." The 5 layers of the FC model include

- ✔ **FC4:** Protocol mapping layer, in which application protocols, such as SCSI or IP, are encapsulated into a PDU for delivery to FC2
- ✔ **FC3:** Common Services layer, which is a thin layer that could eventually implement functions, such as encryption or RAID redundancy algorithms
- ✔ **FC2:** Network Layer, which consists of the core of Fibre Channel and defines the main protocols

- ✓ **FC1:** Data Link Layer, which implements line coding of signals
- ✓ **FC0:** Physical Layer, which includes cabling, connectors, and so on

Standards versus Proprietary Protocols

Network vendors often struggle with the balance between openness and competitive advantage. A new and innovative technology that other companies can't copy may allow a vendor to claim a competitive advantage or even to collect licensing fees from competing vendors.

Unfortunately, such proprietary systems can lock you into a single vendor and allow them to charge you higher prices. In addition, if that vendor disappears or abandons that proprietary technology, there may be no one who can support or enhance the technology.

One alternative to proprietary technologies is an open-systems approach in which standards bodies, such as the Institute of Electrical and Electronic Engineers (IEEE) or ISO, define technologies. Ethernet, Transmission Control Protocol (TCP), Internet Protocol (IP), and Spanning Tree Protocol (STP) are all examples of technologies that became standards.

16 Getting Started with Networking, Wireless, and Security For Dummies

Chapter 3

Looking into Next Generation Data Centers (NGDC)

In This Chapter

- ▶ Checking out data centers and their complexities
 - ▶ Identifying new opportunities
 - ▶ Understanding new protocols and standards
 - ▶ Flattening the network
 - ▶ Virtualizing the network
-

Data centers came from mainframe computing. One large computer held all applications and communications. Now, many storage units are evolving from separate disk arrays to virtual pools of storage. In this chapter, you see how networking is the last major piece of the data center to virtualize.

Understanding Today's Data Centers

Data centers house the computing resources used to support the business-critical applications — along with accompanying computing resources, such as main frames, servers, and server farms. By consolidating all the critical computer resources under a controlled, centralized management, data centers enable enterprises to operate around the clock or according to their business needs. Often, redundant back-up data centers are deployed in order to reduce the risk of massive data loss in the event of a natural disaster or large failure.



A data center must have a strong network infrastructure to support the accessibility of critical applications to the entire enterprise. Each part of the data center infrastructure plays a specific role:

- ✓ **Aggregation:** Where the data center connects to the backbone corporate network
- ✓ **Front end:** Servers with which users interact
- ✓ **Back end:** Where data is stored, typically in relational database systems
- ✓ **Application:** Servers running code that links the front-end information to the back-end data
- ✓ **Storage:** The actual storage device that houses the data

Identifying Opportunities for Networking in the NGDC

In today's data centers, multiple layers of Internet Protocol (IP) traffic exist that travel from the NIC (Network Interface Card) to the switch to the router . . . and converge at the core. These multiple layers result in many "hops" for a single piece of data to travel, which induces latency and/or loss of packets. It can also mean retransmitting of the same data, and your IP network can incur multiple levels of chatter.

As you add applications, servers, and storage to facilitate more data storage and cloud environments, you have to provide more bandwidth for all these pieces to work together. This results in large capital expenditures on equipment and complex management for networking, along with the need for people who understand the new systems and can maintain them.

Bringing in New Protocols and Standards

Evolving technology often brings new and better ways to accomplish existing tasks. For example, Storage Area Networks (SANs) are evolving to offer higher performance

than they did in the past, and a part of this evolution comes in the form of new protocols and standards. With network speeds of 10G, 40G, and even 100G, existing protocols simply can't keep up.



Fibre Channel over Ethernet (FCoE) will be the standard that enables transport of Fibre Channel (FC) packets in SANs. FCoE has the same lossless standards as FC but can be encapsulated in Ethernet, increasing distance and lowering costs. DCB (Data Center Bridging) is a new standard that introduces the lossless Ethernet that will be required for FCoE and also benefit existing iSCSI installations. (More on these topics in Chapter 2.)

Transparent Interconnection of Lots of Links (TRILL) is a new standard that enables cloud computing as it increases switch and path utilization, provides underlying support for virtual servers and mobile devices, and enables faster convergence support.

Flattening Your Network

Flattening the network takes several layers of IP traffic and reduces them to a very few layers. This flattening improves latency, reduces hops, results in fewer dropped packets, and improves performance.

Flattening also means that there is less equipment to purchase, manage, power, and cool.



The idea is to remove as many intermediate devices as possible. Doing this decreases the time it takes for a packet to get from one point to another.

Among the ways to accomplish this flattening is to buy new technology, such as virtualization of the hardware and combined Top of Rack (ToR)/End of Row (EoR)/core-to-edge architectures. ToR refers to gathering all connections within the rack, consolidating them, and then using uplinks to the core to minimize cabling needs and latency issues. The PowerConnect 7000 series with 10G uplinks is a great choice for this approach.

With even bigger networks, you may want to connect all the ToR switches into a medium-sized chassis before the core.

Stepping into Network Virtualization

Virtualizing your network can offer some real advantages to your organization. Network virtualization is a natural next step in the process of server virtualization.

In the past, each additional switch in your network meant that you had an additional switch management console, additional static network paths, and no bandwidth sharing. If one of the switches failed, users were knocked off the network until the switch was replaced for the traffic rerouted. By using virtualization, these additional switches can share management, and bandwidth, and path redundancies.

Virtualization also means that fewer people are needed to manage the network because there are fewer physical devices to manage. Network virtualization enables you to better utilize your network as traffic and demand grow.

Chapter 4

Identifying Network Endpoints

In This Chapter

- ▶ Introducing servers
- ▶ Working with network interface cards
- ▶ Figuring out bus adapters
- ▶ Understanding arrays

The building blocks of networks are network hardware devices, such as servers, adapter cards, cables, hubs, switches, routers, and so on. This chapter provides an overview of these building blocks.

Serving as the Network Centers: Servers

Server computers are the lifeblood of any network. Servers provide the shared resources that network users crave, such as file storage, databases, e-mail, Web services, and so on. Depending on your needs, you may have a single physical server which serves multiple purposes or you may use dedicated computers which each provide a single server function.



An *NAS device* is a self-contained file server that's preconfigured and ready to run. All you have to do to set it up is take it out of the box, plug it in, and turn it on. NAS devices like the Dell PowerVault units are easy to set up and configure, easy to maintain, and less expensive than traditional file servers.

Accessing the Network with NICs

Every computer on a network, both clients and servers, requires a Network Interface Card (NIC) in order to access the network. A NIC is often a separate adapter card that slides into one of the computer's motherboard expansion slots. However, many newer computers have a built-in NIC, also called LAN on Motherboard (LOM), so a separate card isn't needed. Servers often have more than one NIC so they can connect to more than one network, such as your LAN and the Internet.

Every Ethernet NIC has its own *Media Access Control address* (MAC address) that uniquely identifies the NIC. This MAC address makes it possible to address each computer directly because no two NICs should ever have the same MAC address. Each NIC manufacturer has its own set of MAC addresses that are assigned by the Institute of Electrical and Electronics Engineers (IEEE).

The NIC is both an OSI layer 1 (physical layer) and layer 2 (data link layer) device because it provides physical access to a networking medium and provides a low-level addressing system through the use of MAC addresses.

Using a Host Bus Adapter

With the advent of new types of network connections, such as SCSI, Fibre Channel, and eSATA, a device called a *host bus adapter* (HBA) has come into use. These devices, which are also called host controllers or host adapters, can also include devices for connecting to IDE, Ethernet, FireWire, USB and other systems. The advent of iSCSI and Fibre Channel over Ethernet has brought about Ethernet HBAs, which include TCP Offload Engines. Most often, though, the term host bus adapter is used to refer to a Fibre Channel interface card. Each HBA has a unique World Wide Name (WWN), which is similar to an Ethernet MAC address. There are two types of WWNs on a HBA:

- ✓ A node WWN (WWNN), which is shared by all ports on a host bus adapter
- ✓ A port WWN (WWPN), which is unique to each port

HBA's are available in a number of different speeds: 1Gbit/s, 2Gbit/s, 4Gbit/s, 8Gbit/s, 10Gbit/s and 16Gbit/s.



A converged network adapter (CNA) is the combination of a NIC and HBA into one physical card with Ethernet ports for connecting to the network. CNA's are primarily used in Fibre Channel over Ethernet environments.

What Is an Array?

A *disk array* is a disk storage system that contains multiple disk drives, such as the Dell EqualLogic iSCSI storage array. It's differentiated from a disk enclosure, in that an array has cache memory, processor, and advanced functionality, like RAID and virtualization.

Typically a disk array provides increased availability, resiliency, and maintainability by using additional, redundant components (controllers, power supplies, fans, and so on), often up to the point when all single points of failure (SPOFs) are eliminated from the design. Additionally those components are often hot-swappable.

Chapter 5

Security 101

.....

In This Chapter

- ▶ Physically securing your network equipment
 - ▶ Figuring out user account security
 - ▶ Hardening your network and deciphering antivirus programs
 - ▶ Controlling traffic with firewalls
-

Before you had a network, computer security was easy. You simply locked your door when you left work for the day. You could rest easy, secure in the knowledge that the bad guys would have to break down the door to get to your computer. The network changes all that. Now, anyone with access to any computer on the network can break into the network and steal *your* files. Not only do you have to lock your door, but also you have to make sure that other people lock their doors, too.

Physical Security: Locking Your Doors

The first level of security in any computer network is physical security. Physical security is important for workstations but vital for servers. Any hacker worth his salt can quickly defeat all but the most paranoid security measures if he can gain physical access to a server.

Securing User Accounts

Next to physical security, the careful use of user accounts is the most important type of security for your network. Properly

configured user accounts can prevent unauthorized users from accessing the network, even if they gain physical access to the network. Users who write their usernames and passwords on sticky notes on their monitors should be publically flogged — or at least educated about why this is such a bad idea!

Hardening Your Network

You should also take steps to protect your network from intruders by configuring the other security features of the network's servers and routers. The following sections describe the basics of hardening your network.

Using a firewall

A *firewall* is a security-conscious router that sits between your network and the outside world and prevents Internet users from wandering into your LAN and messing around. Firewalls are the first line of defense for any network that's connected to the Internet. For more information about firewalls, refer to “Firewalls — Network Traffic Cops” later in this chapter.

Disabling unnecessary services

A typical network operating system can support dozens of different types of network services: file and printer sharing, Web and mail servers, and many others. In many cases, these features are installed on servers that don't need or use them. When a server runs a network service that it doesn't really need, the service not only robs CPU cycles from other services that are needed but also poses an unnecessary security threat.



When you first install a network operating system on a server, enable only those network services that you know the server requires. You can always enable services later if the needs of the server change.

Patching your servers

Hackers regularly find security holes in network operating systems. After those holes are discovered, the operating

system vendors figure out how to plug the hole and release a software patch for the security fix. The trouble is that most network administrators don't stay up to date with these software patches. As a result, many networks are vulnerable because they have well-known holes in their security armor that should've been fixed but weren't.

Even though patches are a bit of a nuisance, they're well worth the effort for the protection that they afford. Fortunately, newer versions of the popular network operating systems have features that automatically check for updates and let you know when a patch should be applied.

Using network appliances

You can harden your network by using a network appliance. Traditional Web antivirus gateways often lack scalability and performance for HTTP, HTTPS, and FTP scanning, leaving desktops to defend themselves. By using a network security appliance you get a proven record of enterprise robustness for effective virus scanning, plus complete visibility and control of enterprise Web communications.

Securing Your Users

Security techniques, such as physical security, user account security, server security, and locking down your servers, are child's play compared to the most difficult job of network security: securing your network's users.

The key to securing your network users is to create a written network security policy and stick to it. Have a meeting with everyone to go over the security policy to make sure that everyone understands the rules. Also, make sure to have consequences when violations occur.

Understanding Antivirus Programs

The best way to protect your network from virus infection is to use an antivirus program. These programs have a catalog

of several thousand known viruses that they can detect and remove. In addition, they can spot the types of changes that viruses typically make to your computer's files, thus decreasing the likelihood that some previously unknown virus will go undetected.



The people who make antivirus programs have their fingers on the pulse of the virus world and frequently release updates to their software to combat the latest viruses. Because virus writers are constantly developing new viruses, your antivirus software is next to worthless unless you keep it up to date by downloading the latest updates.

If you're looking to deploy antivirus protection on your network, here are several approaches:

- ✔ **Install antivirus software on each network user's computer.** This technique is the most effective if you can count on all your users to keep their antivirus software up to date. That's highly unlikely, so you may want to adopt a more reliable approach to virus protection.
- ✔ **Place antivirus client software on each client computer in your network.** An antivirus server automatically updates the clients regularly.
- ✔ **Use server-based antivirus software to protect your network servers from viruses.** For example, you can install antivirus software on your mail server to scan all incoming mail for viruses and remove them before your network users ever see them.
- ✔ **Limit Internet access:** Some firewall appliances include antivirus enforcement checks that don't allow your users to access the Internet unless their antivirus software is up to date. This type of firewall provides the best antivirus protection available.
- ✔ **Place appliances accordingly:** Place appliances responsible for catching Web-based threats, malware, and antivirus at the point where traffic crosses from your LAN to the Internet.

Firewalls — Network Traffic Cops

A *firewall* is a security-conscious piece of hardware or software that sits between the Internet and your network with a single-minded task: preventing *them* from getting to *us*. The firewall acts as a security guard between the Internet and your local area network (LAN). All network traffic into and out of the LAN must pass through the firewall, which prevents unauthorized access to the network.



Some type of firewall is a must-have if your network has a connection to the Internet. Without it, sooner or later a hacker will discover your unprotected network and tell his friends about it. Within a few hours, your network will be toast.

You can set up a firewall in two basic ways:

- ✓ **Purchase a firewall appliance.** This way is basically a self-contained router with built-in firewall features. Most firewall appliances include a web-based interface that enables you to connect to the firewall from any computer on your network by using a browser. You can then customize the firewall settings to suit your needs.
- ✓ **Set up a server computer to function as a firewall computer.** The server can run just about any network operating system, but many dedicated firewall systems run Linux.

Whether you use a firewall appliance or a firewall computer, the firewall must be located between your network and the Internet. As a result, all traffic from the LAN to the Internet and vice versa must travel through the firewall.

Looking at the Basic Types of Firewalls

Firewalls employ certain basic techniques to keep unwelcome visitors out of your network. The following sections describe the most common firewall techniques.

Packet filtering

A *packet-filtering* firewall examines each packet that crosses the firewall and tests the packet according to a set of rules that you set up. If the packet passes the test, it's allowed to pass. If the packet doesn't pass, it's rejected.



Packet filters are the least expensive type of firewall. As a result, packet-filtering firewalls are very common. Packet filters work by inspecting the source and destination IP and port addresses contained in each Transmission Control Protocol/Internet Protocol (TCP/IP) packet. However, packet filtering has a number of flaws that knowledgeable hackers can exploit. As a result, packet filtering by itself doesn't make for a fully effective firewall.



Consider adding an appliance like Dell's Power Connect J-SRX Series Service Gateway when connecting, securing and managing workforce centers with security in mind.

Stateful packet inspection (SPI)

Stateful packet inspection (SPI) is a step up in intelligence from simple packet filtering (see the preceding section). A firewall with stateful packet inspection looks at packets in groups rather than individually. It keeps track of which packets have passed through the firewall and can detect patterns that indicate unauthorized access. In some cases, the firewall may hold on to packets as they arrive until the firewall gathers enough information to make a decision about whether the packets should be authorized or rejected.

Deep packet inspection

Another more comprehensive way to inspect the packets is *deep packet inspection* (DPI). In DPI, the actual content of the packets is inspected so viruses, spam, and other harmful content can be blocked.



DPI also allows practices, such as data mining, eavesdropping, and content censorship, which make its use a controversial subject.

Chapter 6

Setting Up and Securing a Wireless Network

In This Chapter

- ▶ Looking at wireless network basics
 - ▶ Working with a wireless access point
 - ▶ Enabling the security features of your wireless network
-

This chapter introduces you to the ins and outs of setting up a wireless network. With wireless networking, you don't need cables to connect your mobile devices. Instead, wireless networks use radio waves to send and receive network signals. As a result, a mobile device can connect to a wireless network at any location in your office.

Wireless networks are especially useful for laptops, smartphones, and tablets. After all, the main benefit of these devices is you can carry them around with you wherever you go. For example, at work, you can use your laptop at your desk, in the conference room, in the break room, or even out in the parking lot. With wireless networking, your portable devices can be connected to the network no matter where you take it.

Diving into Wireless Networking

A *wireless network* is a network that uses radio signals rather than direct cable connections to exchange information. An example of this kind of network includes devices from Dell's PowerConnect W-Series, powered by Aruba.

A computer with a wireless network connection is like a cell phone. Just as you don't have to be connected to a phone line to use a cell phone, you don't have to be connected to a network cable to use a wireless mobile device.

In order to set up and use a basic wireless network, you need to understand the following key concepts and terms:

- ✔ **A wireless network is often referred to as a *wireless local area network (WLAN)*.** Some people prefer to switch the acronym around to *local area wireless network (LAWN)*. The term *Wi-Fi* is often used to describe wireless networks.
- ✔ **WLANs share bandwidth between multiple devices.** A wired network typically has a single device with dedicated bandwidth.
- ✔ **A wireless network has a name, known as a *service set identifier (SSID)*.** Each of the mobile devices that belong to a single wireless network must have the same SSID.
- ✔ **Wireless networks can transmit over any of several channels.** In order for portable devices to talk to each other, they must be configured to transmit on the same channel.
- ✔ **The simplest type of wireless network consists of two or more computers with wireless network adapters.** This type of network is called an *ad-hoc mode network*.
- ✔ **A more complex type of network is an *infrastructure mode network*.** All this really means is that a group of portable devices can be connected not only to each other but also to an existing cabled network via a device called a *wireless access point* or WAP. (Discover more about ad-hoc and infrastructure networks later in this chapter.)

Wireless Access Points

Unlike cabled networks, wireless networks don't need a hub or switch. If all you want to do is network a group of devices, you just purchase a wireless adapter for each one, put them all within 300 feet of each other, and *voilà!* — instant network.

But what if you already have an existing cabled network? For example, suppose that you work at an office with 15 computers all cabled up nicely, and you just want to add a couple of

wireless laptops to the network. Or suppose that you have a conference room where executives want to meet, use their mobile devices, and not worry about crawling under the tables looking for network connections.

That's where a wireless *access point (AP)* comes in. A wireless AP actually performs two functions:

- ✓ It acts as a central connection point for all your mobile devices that have wireless network adapters. In effect, the wireless AP performs essentially the same function as a hub or switch performs for a wired network.
- ✓ It links your wireless network to your existing wired network so your wired devices, such as desktops, servers, and printers, and your wireless mobile devices get along like one big happy family.

Infrastructure mode

When you set up a wireless network with an access point, you're creating an *infrastructure mode* network. It's called infrastructure mode because the access point provides a permanent infrastructure for the network. The access points are installed at fixed physical locations, so the network has relatively stable boundaries. Whenever a mobile device wanders into the range of one of the access points, it has come into the sphere of the network and can connect.



An access point and all the wireless devices that are connected to it are referred to as a *Basic Service Set (BSS)*. Each BSS is identified by an SSID. When you configure an access point, you specify the SSID that you want to use. The SSID is often a generic name such as *wireless*, or it can be a name that you create. Some access points use the MAC address of the wireless AP as the SSID.

Roaming

You can use two or more wireless access points to create a large wireless network in which mobile device users can roam from area to area and still be connected to the wireless network. As the user moves out of the range of one access point, another access point automatically picks up the user and takes over without interrupting the user's network service.

To set up two or more access points for roaming, you must carefully place the wireless APs so all areas of the office or building that are being networked are in range of at least one of the APs. Then, just make sure that all the mobile devices and access points use the same SSID and channel.

Two or more access points joined for the purposes of roaming, along with all the wireless devices connected to any of the access points, form what's called an *Extended Service Set (ESS)*. The access points in the ESS are usually connected to a wired network.

Wireless bridging

Another use for wireless access points is to bridge separate subnets that can't easily be connected by cable. For example, suppose that you have two office buildings that are only about 50 feet apart. To run cable from one building to the other, you'd have to bury conduit — a potentially expensive job. Because the buildings are so close, though, you can probably connect them with a pair of wireless access points that function as a *wireless bridge* between the two networks. Connect one of the access points to the first network and the other access point to the second network. Then, configure both access points to use the same SSID and channel.

Ad-hoc networks

A wireless access point isn't necessary to set up a wireless network. Any time two or more wireless devices come within range of each other, they can link up to form an *ad-hoc network*. For example, if you and a few of your friends all have portable devices with 802.11n wireless network adapters, you can meet anywhere and form an ad-hoc network. All devices within range of each other in an ad-hoc network are called an *Independent Basic Service Set (IBSS)*.

Securing Your Wireless Network

Before you dive headfirst into the deep end of the wireless networking pool, you should first consider the inherent security risks in setting up a wireless network. With a cabled network, the best security tool that you have is the lock on the front

door of your office. Unless someone can physically get to one of the mobile devices on your network, she can't get into your network. (Well, I'm sort of ignoring your wide-open broadband Internet connection for the sake of argument.)

If you go wireless, an intruder doesn't have to get into your office to hack into your network. She can do it from the office next door, the lobby, or the parking garage beneath your office. In short, when you introduce wireless devices into your network, you usher in a whole new set of security issues to deal with.



When you first install a wireless access point, change its administrative password and then secure the SSID that identifies the network. A client must know the access point's SSID in order to join the wireless network. If you prevent unauthorized clients from discovering the SSID, you prevent them from accessing your network.

Using MAC address filtering

MAC address filtering allows you to specify a list of MAC addresses for the devices that are allowed to access the network. If a device with a different MAC address tries to join the network via the access point, the access point denies access.

MAC address filtering is a great idea for wireless networks with a fixed number of clients. For example, if you set up a wireless network at your office so a few workers can connect their mobile devices, you can specify the MAC addresses of those devices in the MAC filtering table. Then, other devices won't be able to access the network via the access point.



Unfortunately, it isn't difficult to configure a device to lie about its MAC address. Thus, after a potential intruder determines that MAC filtering is being used, he can just sniff packets to determine an authorized MAC address and then configure his device to use that address. (This is called *MAC spoofing*.) So you shouldn't rely on MAC address filtering as your only means of security.

Additionally, establishing and maintaining authorized MAC addresses (white lists) may be time-consuming and require manual changes when new users and devices need wireless access or become unauthorized. With a large number of

people bringing personally-owned smartphones and laptops into the WLAN, MAC address filtering may not be the best option.

Role-based access

Role-based access is a security model where a policy is assigned to a user dependent on the user's credentials. Users validate their credentials residing on an authentication server by using an authentication method, such as 802.1X. After authenticated and authorized for network access, the user can be assigned a policy defining where the user may go, what may be accessed, time of day of access, as well as other parameters.

These user policies may place a user in a particular SSID, and separate firewalls may enforce these policies based on the SSID and attached VLAN (*Virtual Local Area Network*). However, some vendors provide integral policy enforcement firewalls allowing a unique policy for each user and device without the overhead of SSID and VLAN pairs.

Using encryption

One big problem with wireless networking is simply that it's wireless. This means that anyone nearby with a wireless adapter in their mobile device has the potential to snoop on your wireless connection. Fortunately, there's a solution — encryption.

By using encryption the wireless signals are protected from snooping because no one else will be able to read their contents. If you don't use encryption on your wireless network, it's almost like leaving your credit cards on a park bench — you can hope everyone is honest, but it's not a good bet.

The first wireless networks used an encryption method called *Wired Equivalent Privacy (WEP)*, which really didn't do much other than make users think that they were secure. Modern wireless networking equipment uses Wi-Fi Protected Access (WPA or WPA2), which actually does provide reasonable protection. WPA2 uses an AES (Advanced Encryption Standard)-based algorithm, CCMP, which is considered fully secure.

Chapter 7

Ten (Okay, Nine) Networking Strategies to Consider

In This Chapter

- ▶ Analyzing the needs
 - ▶ Determining a fix
 - ▶ Applying a solution
-

Choosing the correct networking equipment can be confusing, especially if you aren't sure which questions to ask. In this chapter, you see some questions that help you focus your efforts as you choose the networking options that best serve your organization's needs.

What Type of Network Topology Should I Consider?

You need to create a new network or expand an existing one, and the answer to the question depends on many factors, but one of the most important is whether you need to work with existing infrastructure. Most modern networks use a star topology for the wired sections of the network, but if you are expanding an existing network, which uses bus topology, it may be cheaper to stay with what you have. Don't forget, though, that future expansion needs may dictate that you switch to UTP cabling or even wireless networking.

What's the Best Way to Connect Workstations to My Existing Network?

Say your company is considering buying a neighboring building for a call center and you want to connect workstations to your existing network. Well, carrier pigeons are probably out of the picture (they're too slow and you don't want to be responsible for cleaning up after them). So a better option may be to use a pair of routers to extend the network between the two sites. That way, you only need a single wired or wireless connection between the buildings.

What Types of Products Should I Consider?

Your company is spending a lot of money on bandwidth, and you've been tasked with improving the performance of legitimate business processes while reducing those costs. If you suspect that a lot of bandwidth is being wasted on things like music downloads, you could announce a new company policy that anyone seen wearing earphones in their cube will be paraded naked around the building, but you may not want to contemplate how some employees would look in the buff.



Instead, consider adding a firewall appliance, which offers deep packet inspection so you can control which types of traffic are allowed into your network.

What Can I Do about Viruses and Poor Performance?

Even though you've installed anti-virus software on all your PCs, the occasional virus outbreaks still occur, and some users complain about poor performance.

Instead of relying on AV software, consider a network security appliance. By blocking the bad stuff before it ever enters

your network, you'll stop most problems in their tracks. And you won't have to listen to that know-it-all in accounting who insists that his snot-nosed, middle school nephew could make your PCs run faster.

What Should I Look for to Secure My Wireless Network?

You have a very large building and your workers move around a lot with their laptops. A reasonable approach is to make sure that all wireless equipment supports the WPA2 encryption standard.

How Can I Detect and Contain Rogue APs and Other Wireless Threats?

You suspect that some users have added unauthorized wireless access points (AP). The ability to detect unauthorized access points and accurately separate them from valid neighboring networks is critical. A state-of-the-art wireless LAN system not only can detect and classify rogues accurately but also stop rogue access points automatically and accurately point out the source of the threat as well.

An integrated policy enforcement firewall enables you to control the access privileges of each and every user, even if they move around in the network and use a multiple devices, such as laptops, PDAs, and WiFi phones. An integrated policy enforcement firewall provides the highest level of security and trust necessary to deploy mission-critical applications wirelessly.

What Considerations Help Me Plan for Future Expansion?



Be sure that the system you buy today supports your future needs. Ask how many users a single system can accommodate, and be sure it can be centrally managed.



A good wireless system eliminates the need for pre-deployment site surveys by configuring itself dynamically to provide best coverage and performance and to fill coverage gaps in the event of an AP failure. In addition, location-enabled wireless systems accurately locate unauthorized access points and malicious users. Later, location functionality may be used to implement location-enabled applications.

What Do I Need to Ask?

You need to add new network capacity without disrupting day-to-day operations. So what should you ask?

WLAN architectures that don't require configuration changes and upgrades to the existing wired LAN infrastructure are superior to those that require hardware and software upgrades to achieve an "integrated wired and wireless solution." The cost of upgrading software, configuring VLANs everywhere, and enabling mobility and management are too high not to be taken into account when purchasing a WLAN system just because they don't show up on the purchase order.

How Can I Identify the Actual Applications Consuming Bandwidth?

If you need to figure out what's eating up all your bandwidth and you want to monitor their behavior and quality, and set policy-based QoS controls to enable quality standards or manage bandwidth consumption, this is the section for you!

With certain types of network appliances, you can identify all the applications on the network and monitor response times and utilization at the application level. In addition, you can optimize application performance by using granular quality of service (QoS) controls to regulate traffic, ensuring business applications perform optimally.

Case Study A

Aruba's Approach: Taking the Campus Wireless

The Dell PowerConnect W-Series wireless platform is based on Aruba Networks technology. The following case study shows how Aruba Networks Liberty University enhances and expands its campus network by transitioning from the traditional wired Ethernet and older wireless technology to a more modern and capable 802.11n wireless system.

Analyzing the Situation

Liberty University is executing a progressive mobility strategy by transitioning voice and broadcast television services to its wireless LAN. The university's program was planned with careful attention to the return on investment, while offering its users an exceptional information technology (IT) experience.

IT has become the catalyst for growth at Liberty University, enhancing learning and life experience on campus while enabling a distance learning initiative that supports over 41,000 students across the world. IT also influences the university's affiliation with other organizations: Liberty University provides network services to several sister organizations, requiring that any new technology be extensible to remote facilities.

In the most recent upgrade, Liberty University focused on enhancing the wired and wireless networks as well as network security. The networking gear had historically been provided by a single vendor. The 802.11 b/g wireless APs provided some coverage across campus, but problems persisted and the network couldn't be sufficiently retrofitted to support the university's new requirements. With the upgrade, Liberty University

sought to move up to a best-in-class mobility infrastructure that supported all data and media services, including broadcast television (IPTV).

Understanding the Requirements

Liberty University was interested in several network enhancements:

- ✔ **Network availability:** The institution wanted campus-wide network access for increasingly popular handheld computing devices and smart phones. The campus' existing wireless infrastructure was simply not available everywhere these devices were used, and most of the devices didn't support wired Ethernet.
- ✔ **Wireless capacity:** With wireless fast becoming the primary form of network access, expectations grew that the wireless network should support voice, interactive curriculum, and broadcast quality video applications. With the legacy wireless network, two architectural issues made this task impossible:
 - The radio frequency (RF) or wireless capabilities of the legacy solution were designed for large sparsely populated areas, and the radios couldn't support Adaptive RF Management.
 - The legacy wireless network supported only 802.11 b/g, while the new applications needed the 300Mbps per channel and 5GHz operation of 802.11n Wi-Fi.
- ✔ **Network guest access:** The university campus regularly hosts large events at the Vines Center, a facility that seats over 8,000 people. Liberty University lacked a secure and easily administered guest access solution through which it could provide network access.
- ✔ **Identity and access management:** The existing system offered limited identity and access management, lacking features standard on other more sophisticated solutions.
- ✔ **Wireless management:** Monitoring, configuring, and optimizing Liberty University's large and growing network was becoming increasingly challenging by using the university's legacy management tools.

✓ **Security:** Security was always a priority, but over time, the network grew sufficiently, and the university needed a new security paradigm. The new approach needed to focus on delivering a consistent user experience to the growing cadre of mobile users. The network was originally designed to protect static users, but the increase in mobile users called for the use of role-based security to protect users regardless of where they worked or roamed.

Making a Business Case for Wireless

Prior to selecting a solution, the IT staff conducted a detailed review of campus infrastructure. The results highlighted the diminishing relevance and growing expense of Ethernet and justified the case for transitioning to pervasive 802.11n Wi-Fi coverage.

Liberty University determined that by using only Wi-Fi in its residence halls it would be able to significantly reduce operating costs associated with powering and maintaining closet infrastructure.

WLAN Deployment Details

The University has 212 buildings, over 52,000 local and distance education students, and more than 3,500 full time employees. Prior to selecting a vendor to upgrade its Wi-Fi network, the university staged both a 36-AP Aruba network at a remote site, and networks from its incumbent supplier at two residence halls.

At the conclusion of the evaluation period, Aruba's solution was deemed technically superior, and the university subsequently purchased a campus-wide Aruba 802.11n Wi-Fi network managed by Aruba's AirWave management.

Today, the network consists of approximately 800 AP-125 802.11n APs, three Aruba 3000 series Mobility Controllers, and three Aruba 6000 Mobility Controllers equipped with M3 controller modules and deployed in a redundant configuration.

Each M3 controller module supports a 10 Gbps connection to a distribution switch, providing ample backhaul capacity for 802.11n clients accessing bandwidth-intensive applications like broadcast video.

For guest access Liberty University provisions a separate guest SSID that's restricted by Aruba's Policy Enforcement Firewall (PEF) to accessing only the DMZ on the Internet edge switch.

Broadcast Television over 802.11n

Liberty University determined that maintaining a wired network just to support video would be excessively expensive, both with respect to the initial build out and on-going maintenance. Because consumers of video content were increasingly mobile, a determination was made that, in addition to being expensive, Ethernet was simply not a suitable video delivery mechanism.

The University launched a campus-wide program to encode and distribute multichannel, IP-based television (IPTV) campus-wide via wireless LAN to fixed televisions and computers, as well as roaming devices. The solution couples Aruba's adaptive high-speed 802.11n Wi-Fi with HaiVision network video compression and conversion technology. New Aruba multicast optimization technology efficiently transmits video streams while traffic prioritization techniques allow the wireless IPTV service to scale.

The result is an optimized video delivery infrastructure that leverages 802.11n everywhere possible; significantly enhances user mobility while reducing capital and operating expenses.

Liberty University's 802.11n wireless network now supports 15 live IPTV channels over the wireless network.

Case Study B

F5's Approach: Achieving Application Availability

In this case study, you see how Sysmex America, a customer of F5 Networks, was able to surpass 99.999 percent e-mail uptime on their Microsoft Exchange Server 2010 system and ensure business continuity and seamless communication among customers, partners, and employees.

Sysmex America develops clinical testing devices for the healthcare industry. Sysmex must keep e-mail systems highly available with cost-effective, centralized management and a flexible platform that it can later expand to support other key systems.

Understanding the Challenge

Sysmex America, based in Mundelein, Illinois, is the United States headquarters of Sysmex Corporation, a Japanese manufacturer of clinical testing devices and solutions for the healthcare industry. Sysmex relies on e-mail to communicate with its customers (hospitals, clinics, test centers, and blood centers), suppliers, and key business partners. Employees rely on e-mail to communicate information among research and development, production, marketing, and sales teams. Sysmex simply can't afford downtime — e-mail is the critical communications link that keeps business moving forward.

Sysmex relied on Microsoft Exchange Server 2003 for e-mail until the IT department determined that it could benefit from an upgrade. Microsoft Exchange Server 2010 held great promise for improving employee access to unified e-mail, calendars, and voice mail — from virtually any platform, browser, or device. Additionally, the newer Exchange Server version takes a unified approach to high availability and disaster recovery, helping IT departments achieve higher reliability.

High availability was the key driver behind the project, but Sysmex also wanted to ensure that its Microsoft Exchange Server 2010 deployment would be easy to manage, monitor, and protect so IT staff could spend less time troubleshooting and more time on strategic efforts.

Defining a Solution

Sysmex engaged Chicago-based Project Leadership Associates, a technology consulting firm, to review its design and implementation plan for Microsoft Exchange Server 2010. “Project Leadership said, ‘the plan looks great, but Microsoft recommends Application Delivery Controllers for load balancing,’” recalls Arthur Braune, Manager of Information Technology at Sysmex.

Sysmex evaluated a range of application delivery and load balancing technologies, including F5 BIG-IP Local Traffic Manager (LTM) from F5 Networks. BIG-IP LTM is an intelligent, flexible solution that helps companies improve application performance, ensure availability, enhance security, and simplify management.

The BIG-IP LTM offered the following important features that were important to Sysmex:

- ✔ Comprehensive load balancing
- ✔ TCP optimization
- ✔ HTTP server offload (request multiplexing)
- ✔ IPv6 Gateway
- ✔ Application connection persistence

- ✔ Customized health monitoring
- ✔ In-band server health monitoring
- ✔ Hardware accelerated SSL offload
- ✔ Advanced application switching
- ✔ Stateful failover
- ✔ Performance dashboard

To ensure a smooth rollout, Sysmex took advantage of F5's Application Ready Solution for Microsoft Exchange Server 2010. The solution provides step-by-step guidance on deploying F5 devices with Microsoft Exchange Server 2010, including detailed configuration tips for optimized performance.

With Exchange Server 2010, user access to e-mail is managed by the Client Access servers, so the IT department can intelligently manage all of the company's Exchange Server traffic on its BIG-IP devices. Sysmex is using the Database Availability Group feature of Exchange Server 2010, which uses continuous replication to keep database copies up to date and to automate recovery from failures at the disk, server, or data center level.

Sysmex also takes advantage of the F5 solution to offload CPU intensive Secure Sockets Layer (SSL) traffic from its Exchange Client Access servers to BIG-IP LTM. This frees up server resources, resulting in increased responsiveness and higher availability for users.

Realizing the Benefits

Sysmex America took advantage of F5 technology to enhance the availability of Microsoft Exchange Server 2010. Sysmex was very pleased with its F5 investment because of the performance that investment delivers and because the company can extend the solution to support other areas of its infrastructure.

This section covers some of the benefits Sysmex realized after implementing the F5 solutions.

Zero downtime for e-mail

Since deploying Exchange Server 2010 with BIG-IP LTM, Sysmex has experienced no unplanned e-mail downtime. “We have achieved five nines [99.999 percent] of availability,” says Braune. “The F5 solution helps us make the Exchange Server Client Access servers highly available. Users aren’t frustrated with outages, and we can keep business on track with our customers.”

The F5 solution also helps the IT department perform server maintenance without disrupting users. “The last time we did maintenance on the Exchange Server, no one even knew we had taken a server down,” adds Braune.

Superior support

F5’s Application Ready Solutions include deployment guides and provide best practices for smooth implementations and optimized performance of applications, platforms, and enterprise service-oriented architectures from some of the world’s largest software vendors. Sysmex followed the F5 deployment guide explicitly. The quality of F5 documentation, combined with the expertise of its support staff, made this deployment a very good experience. With F5, Sysmex gets what it looks for in a vendor — the ability to respond quickly and make things work.



The F5 deployment guide for Exchange Server 2010 was developed by F5 engineers working closely with Microsoft, and it delivers the best possible guidance to support a highly available and scalable deployment.

“Thanks to F5, we have complete confidence that we are following Microsoft best practices in implementing a critical business system. To an IT department, that in itself is an immeasurable benefit,” explains Braune.

Cost-effective platform for growth

The F5 solution for Exchange Server 2010 is feature-rich and highly expandable. It provides IT departments with flexible options for systems management, tools for simplified troubleshooting, and step-by-step configuration guidance. Everything that’s needed is on one BIG-IP device.

Sysmex considered a competitor's solution, but to get all that it has with F5, the company would have had to invest in twice as many appliances, incurring unnecessary cost and adding complexity. The F5 solution provides a strategic point of control and better visibility into Sysmex's systems. This simplifies security and troubleshooting. The company wouldn't have been able to create such a reliable Exchange Server environment without F5.

BIG-IP LTM special benefits

The F5 BIG-IP LTM solution offers some special benefits for the enterprise.

Plan for growth and avoid downtime

With BIG-IP LTM, you get advanced load balancing and application health monitoring capabilities. BIG-IP LTM enables you to seamlessly add physical or virtual servers and redirect traffic away from issues with specific servers or network components.

Accelerate your applications up to 3x

BIG-IP LTM reduces traffic volumes and minimizes the effect of client connection bottlenecks as well as WAN, LAN, and Internet latency to improve application performance up to three times.

Secure your applications and data

From network and protocol-level security to application attack filtering, BIG-IP LTM deploys a suite of security services to protect your applications. Add-on modules for secure access, spam reduction, and application protection can provide additional advanced security options.

Reduce servers, bandwidth, and management costs

Advanced TCP connection management, TCP optimization, and server offloading enable you to optimize the utilization of your existing infrastructure — tripling server capacity and reducing bandwidth costs by up to 80 percent. BIG-IP LTM helps you simplify system management by consolidating security, acceleration, and availability in one easy-to-manage platform. By using fewer servers, less bandwidth, less power,

and less cooling, while reducing the time spent time managing your infrastructure; you can significantly reduce your operational costs.

Take control over application delivery

The F5 TMOS platform gives you complete control of the connections, packets, and payloads for applications. Using F5's event-driven iRules, you can customize how you intercept, inspect, transform, and direct inbound and outbound application traffic. The F5 iControl API makes it easy to integrate with third-party management systems.

Case Study C

Dell's Approach: Reducing Costs and Streamlining Administration

In this case study, you see how Dell helped a large school division reduce IT costs and streamline administration while expanding the robust technology services offered to students, faculty, and staff.

Facing the Challenges

Technology plays a vital role in providing a rich academic experience for the 14,000 students of the one of the largest school divisions in Manitoba, Canada. Don Reece, IT director, said their vision is for technology to be ubiquitous and seamlessly integrated into the curriculum. The division employs a wide range of solutions to promote teaching innovation, to keep students engaged, and to enhance the efficiency of school administration. Reece says, “We strive to make the platform transparent, the tools intuitive, and the support attentive to educators and students.”

Technology touches nearly every aspect of school life. Students write on “smart boards” in class using compact netbooks with touch screens; faculty uses classroom computers and ceiling-mounted projectors to deliver colorful audiovisual presentations; teachers and administrators use a unified communications solution for phoning, videoconferencing, and paging to streamline communications while controlling costs.

To provide these and other technology services, the school division's IT group relies on a robust infrastructure built with help from Dell. The IT group created a virtualized environment using Dell servers and Dell storage area networks (SANs). The IT group also implemented Dell Fibre Channel switches to take advantage of the fiber-optic network the division had built between schools. The network connected the primary data center to the domain controller servers residing in each of the schools. By implementing the Dell switches, the school became the fastest educational network in Canada, running at 1 Gb speeds between buildings.

Recently, the IT group embarked on a new project to refresh its infrastructure. They wanted to increase server consolidation to streamline administration, better protect data, and cut costs. The group decided to migrate the domain controllers located at the schools to the virtualized environment in their primary data center. Reece added, "To deliver outstanding application performance for data, voice, and high-definition video out to the schools, we needed to increase network bandwidth."

Accelerating the Network

After evaluating several possible solutions, the division's IT group decided again to select Dell switches. "We tested switches from all of the major vendors. In addition to enabling high-bandwidth connectivity, the switches had to be reliable, provide the quality of service we need, offer an easy-to-use management interface, and be backed by outstanding support," says Reece. "Dell switches clearly met all of our criteria."

Based on open standards, the Dell switches give the IT group the flexibility for change while helping to control costs. The school division's unified communications solution currently uses a combination of Dell servers, Dell switches, and Polycom phone equipment. Because the Dell switches are nonproprietary, the type of SIP [Session Initiation Protocol] device or phone they use can be changed at the end without running into technical problems or having to pay the licensing fees required with competing switch solutions. In the near term, open-standard hardware enables the division to manage equipment without requiring in-depth expertise.

With network design assistance from Dell, the Division's IT group replaced every switch in the network, installing over 400 new PowerConnect switches in all. With the Dell switches in place, the school division now achieves 20 Gb speeds between sites.

At the primary data center, the IT group deployed Dell servers, which host a full range of virtualized applications as well as virtualized versions of the domain controllers previously located at individual schools. The Dell servers provide the raw processing performance and memory capacity required for running numerous virtual machines on each physical host. At the same time, the energy-efficient design and open-standard platform are helping the school control costs.

Deploying new servers and switches enabled the IT group to significantly consolidate the server infrastructure. All 34 domain controllers previously located in individual schools now run in the primary data center's virtualized environment. The Dell servers can easily handle the increased virtualization load, and the network can now deliver the bandwidth needed to provide services from the data center to the individual schools.

"We have moved from 34 servers to just 4 — and one of those four is a backup server," says Reece. "That's a 90 percent reduction from seven years ago. By virtualizing with Dell servers, we have reduced hardware acquisition, power, maintenance, and cooling costs by approximately 75 percent while increasing application performance."

Increasing Network Performance

By refreshing switches, the school division has improved its network performance by 20 times. The IT group has gained the capacity for new services and improved service quality for its various communications solutions. The division now has the bandwidth to support unified communications and provide redundancy without having to light up more strands of fiber. The unified communications solution is helping save more than \$200,000 per year (Canadian currency) while enabling double the number of phones. The quality of services to students, faculty, and staff has increased while

actually decreasing the costs by using the latest Dell PowerConnect switches to handle data, voice, and high-definition video.

While reducing the administrative budget, the IT group is expanding its services by improving the performance, reliability, and lifespan of infrastructure. The school can spend a larger percentage of its budget on new projects and services. For example, the IT group intends to enhance the school division's paging and videoconferencing capabilities.

The IT group is also working to help students and faculty develop ways to get connected and stay connected wherever they are. "More students and faculty are carrying their own handheld devices, netbooks, and tablets. Now we can provide a more robust, secure network so they can stay connected while at school," says Reece. "We are also improving virtual private network access so students and faculty can securely extend the academic experience to their home or any other place with WiFi. With our new infrastructure, they can more easily access school information and their own documents anytime, anywhere."

The Dell account team provided key assistance in refreshing the IT infrastructure. The Dell team invited technicians to Dell headquarters to executive briefings and helped meet with some of the engineers who design the switches. The team learned more about the technology and provided Dell with client feedback for future updates. In selecting servers and storage, the IT group also relied on input from Dell.

Case Study D

NetScout's Approach: Managing Performance

In this case study, you see how NetScout's nGenius Service Assurance Solution for J-Flow provides the management solution an organization needs to ensure an easily accessible, high-performing, and always-available network.

Understanding J-Flow

It's been said that you can't manage what you can't measure. That statement is certainly true if you want to manage an enterprise network for optimal performance. J-Flow provides the measurements you need for network management.

Like other flow-based statistics in routing devices, J-Flow in Juniper networking infrastructure products offers a means to collect IP traffic flow statistics from specific routing devices and locations throughout a global enterprise network. J-Flow leverages sampled data from enabled devices and sends those records at a user defined sample rate to a user-configurable universal collector.

J-Flow requires that the flow of IP packets is tracked as it passes through an enabled interface. Each flow is identified by seven criteria:

- ✓ Source IP address
- ✓ Destination IP address
- ✓ Source port number (TCP/UDP)
- ✓ Destination port number (TCP/UDP)
- ✓ Layer 3 protocol type (IP/ICMP)
- ✓ Type of Service (ToS) bit or Differentiated Services code point (DSCP)
- ✓ Input interface

J-Flow uses any combination of these criteria to keep each flow unique and distinguish one from another.

Getting to Know nGenius Service Assurance Solution for J-Flow

J-Flow is a source of application conversation and performance information that you can use in a variety of ways to meet the challenges associated with performance and problem management. Sometimes, though, you may want a little help using that J-Flow data more effectively.

The nGenius Service Assurance Solution can leverage J-Flow conversation data to provide you with service delivery management information as well as to serve as an early warning system. Using the nGenius Service Assurance Solution, you have the ability to go from the evidentiary information seamlessly to the powerful diagnosis information that pertains to the network incident in question, ultimately accelerating problem resolution.

The nGenius Service Assurance Solution provides you with a unified view of applications, hosts, conversations, and utilization data for easy network-wide analysis and reporting. Real-time views and historical reporting displays the information in an easy-to-use unified interface, alleviating the need to manually format and analyze the raw data.

The nGenius Service Assurance Solution provides detailed visibility into application performance, infrastructure optimization, and planning, as well as protecting business services delivery.

Analyzing J-Flow

Looking at a ton of raw data isn't a very efficient process and can lead to errors in addition to less than optimal solutions. The nGenius Service Assurance Solution provides a unique combination of capabilities that are necessary to cost-effectively optimize the performance of networks that drive today's networked business functions.

Multiple approaches using elements of the nGenius Service Assurance Solution are available — for either J-Flow only focused management strategies or for adding J-Flow to a broader, packet-based service delivery management program. The nGenius Service Assurance Solution provides a unique and powerful functionality revealing the details and trends collected by the nGenius Collectors and analyzed through the Common Data Model (CDM) technology. NetScout's CDM architecture provides the flexibility to show

- ✓ Application flow data from J-Flow records for all applications, not just the “Top N” for a complete analysis of service delivery
- ✓ Automated alarming for early warning problem detection
- ✓ Network and application utilization along with hosts and conversations trends
- ✓ Utilization statistics from MIB II data
- ✓ Network statistics and errors in QoS levels
- ✓ MPLS and VPN visibility

This information is displayed simultaneously on screens and reports, giving a more detailed and complete view of network and application performance.

The resulting rich traffic information supports challenging service delivery management tasks that include real-time monitoring, in-depth troubleshooting, and historical or ad hoc reporting for capacity planning and traffic engineering. Combining J-Flow and MIB II data with the nGenius Service Assurance Solution for Flows analysis capabilities enhances your ability to

- ✔ Plan and execute more informed infrastructure investments with evidence from your own environment
- ✔ View the impact of newly introduced applications throughout the infrastructure
- ✔ Evaluate better and less expensive alternatives to costly upgrades by tracking and trending who and what consumes bandwidth
- ✔ Spot potential bottlenecks with intelligent baselines, anomaly detection, and alarming capabilities
- ✔ Accelerate problem resolution
- ✔ View real-time J-Flow data and alarms alongside historical reports for every J-Flow enabled interface on the network
- ✔ View users and activity contributing to increased WAN utilization to predict and prevent congestion issues
- ✔ Receive alerts and generate ad hoc reports for collaboration with other IT groups and third party providers
- ✔ Optimize infrastructure and planning
- ✔ View quality of service (QoS) traffic to validate adherence to corporate policy for priority delivery of services
- ✔ Analyze trend projections in bandwidth consumption for every application across LAN and WAN infrastructure
- ✔ Validate QoS configuration changes within the network

Identifying Complex Applications from J-Flow Conversations

J-Flow supports IP and its well-known TCP and UDP-based applications, for example Lotus Notes, HTTP, or Telnet. These applications are identified by their well-known TCP or UDP

ports and are recognized by most J-Flow collectors, including nGenius Collectors. However, a number of applications are more complex in nature, such as SAP or Exchange, which can be transported on multiple ports. Other collection tools are unable to differentiate ports, and as a result, may label them as TCP Other or UDP Other.

NetScout leverages the powerful capabilities in CDM technology to recognize these complex and customized applications. An application, such as SAP that uses a range of ports or a customized application that isn't defined by a "well-known" port, can be monitored by using CDM technology for recognizing, labeling, tracking, and aggregating the activity of that application as it traverses the global network environment.



These critical complex and customized applications can now be monitored and reported, providing an important system-wide view of all activity thus empowering you to make more informed decisions about capacity planning, troubleshooting, and service delivery management.

Benefits to the nGenius Service Assurance Solution for J-Flow

The nGenius Service Assurance Solution leverages J-Flow as a data source, offering metrics and details vital for service delivery management. J-Flow, as part of the nGenius Service Assurance solutions from NetScout, offers numerous benefits including helping to

- ✓ Optimize infrastructure and capacity planning with historical trending and analysis in customizable ad hoc and customizable reports
- ✓ Protect service delivery with detailed application recognition applied to J-Flow-collected records
- ✓ Simplify operations and improve efficiencies with unified service delivery management solution

Case Study E

Riverbed Technology, Inc.'s Approach: Making the Network Deliver

This case study shows how Riverbed Technology, Inc helped CSX accelerate WAN traffic and improve network speed to better handle their vast system of rail-based transportation services. Riverbed Technology is an IT performance company for networks, applications, and storage. Riverbed provides comprehensive WAN optimization solutions.

The customer, CSX Corporation, is one of North America's largest providers of rail-based transportation services. The company ships a variety of industrial and consumer materials and operates approximately 21,000 miles of rail across the eastern United States.

Defining the Problem

In 2000, CSX consolidated its network and moved its remote servers to a centralized data center. The company's infrastructure was now more streamlined, but end-users experienced slow response times across the network.

CSX's challenges were multifold. The first challenge was early in 2000. CSX went through a centralization of servers; 98 percent of all the servers' services — file, print, database, messaging — were centralized into one data center. Then it deployed Citrix thin-clients, but it had a mainframe, too. The customer found that web browsing, Microsoft Outlook, and

even opening up Excel spreadsheets across the WAN were impacting the performance of the Citrix applications and even the mainframe applications.

CSX's next challenge was accelerating hundreds of applications, including Oracle applications, Java-based web applications, and SharePoint 2006 for its 650 branches; 450 of which have PCs and printers while the rest have thin-client terminals and printers. Many remote users would wait 5 to 10 minutes to access data. In order to speed up application performance for users in the branches while keeping a centralized IT infrastructure, CSX decided to evaluate WAN acceleration technologies.

Examining the Options

The CSX IT team worked with a system integrator to evaluate WAN optimization solutions. As a first step, the system integrator took detailed field measurements of their end-users' wait times.

After collecting the field measurement information, they recreated the field test scenarios in CSX's Jacksonville network lab, where they tested the WAN optimization solutions. CSX looked at several options.

The customer spent about three months evaluating technologies and just kept coming back to Riverbed. It wanted full buy-in from the Wide Area Networking group as well as the applications folks and the Enterprise Architects, so it kept bringing teams in and showing them the different products. Everyone agreed that the Riverbed product was the best performer, and had the most flexible deployment capability of being in-path, out-of-path, and virtual in-path with the WCCP protocol.

Another key reason CSX chose Riverbed was the ability to support SSL traffic. CSX has over 390 applications that are SSL-encrypted through a single sign-on portal. When it was time to deploy the Steelhead appliances in CSX's enterprise network, Riverbed offered the most scalable solution with the least disruption to their network.

Creating a Solution

To accelerate WAN traffic and improve network speed, CSX chose the Riverbed Steelhead Appliance solution. By offering a quick deployment process, Riverbed technology enabled CSX to see immediate improvements and achieve ROI in just eight months.

CSX decided to use an out-of-path design at the data center, and used a virtual-in-path WCCP design in remote offices. This deployment allowed telnet and mainframe traffic to pass through the router without being touched.

Deploying Riverbed throughout CSX was completed in a phased approach. First CSX focused on 13 regional sites, with each office having between 200 to 250 remote computers connected across the 768K circuit. The Steelhead appliances deployed in over 100 locations that had from 10 to 200 users. The company shipped the appliances out to the remote sites, and had the communication operations staff install the Steelheads; CSX walked them through configuring the WCCP commands. Average deployment only took about 30 minutes of hands-on deployment time, though. Now CSX can deploy remote sites in the middle of the day, without impacting the business.

Gauging the Benefits



Riverbed solutions have improved WAN performance across its network; CSX is helping employees access data and applications faster than ever before while keeping the costs of equipment upgrades low. This provided improved application performance and significant ROI.

By deploying Riverbed Steelhead appliances to over 100 sites, CSX has been able to improve application performance. Monthly, the company removes about 15 terabytes of duplicate data from the wide area network. It's cleaning that traffic off the WAN and allowing real business traffic through — not the duplicate bytes — 91 percent of the data has been reduced. With CIFS traffic, CSX now gets an average of 85 percent reduction per month, and HTTP traffic is reduced by over 50 percent.

By deploying Riverbed, CSX has been able to improve performance for its key enterprise applications without having to upgrade their bandwidth. CSX was able to finance the Riverbed deployment with funds that were earmarked for PC upgrades. The customer's business approach invested in the wide area network optimization solution instead of upgrading all its PCs at the remote sites. In effect, end-users now had a PC that was running five times faster than before, and support complaints dropped. CSX was able to delay the PC refresh project and extend the life of a PC for another year in the field; it basically paid for the Riverbed deployment within the first eight to nine months.

CSX has also noticed a qualitative benefit. After deploying over 100 Steelhead appliances, the CSX team noticed that Internet browsing traffic decreased from 25 percent to 5 percent, and now end-users were able to download Excel spreadsheets from their file shares and utilize the network more effectively. Not only did the client improve performance of applications, but it was truly a productivity boom. The end-users immediately noticed the performance gains resulting from the deployment of Riverbed appliances at their offices.

Drawing a Conclusion

CSX Corporation was looking to address application performance for 390 applications across their consolidated IT infrastructure. This transportation company wanted to provide LAN-like application performance to its 100 remote offices. With Riverbed Steelhead appliances, CSX has been able to increase employee productivity and forgo a PC refresh for remote end-users. The payback period for the Riverbed investment was 8 to 9 months.

Networking: Dell and Force10 at
www.dell.com/networking