OUR 15TH YEAR

PACKET®

CISCO SYSTEMS USERS MAGAZINE

THIRD QUARTER 2003

Wireless LANs 30

Going Mobile in a Big Way

Cisco Systems

cisco.com/go/packet

PACKET®

CISCO SYSTEMS USERS MAGAZINE THIRD QUARTER 2003 VOLUME 15 NO. 3

30


39

▼ ON THE COVER

Wireless Takes Off

The market for wireless is heating up for businesses and consumers alike. *Packet*® looks at the status of the wireless landscape and examines the issues that the latest technology trends raise for user deployments

▼ FEATURES

Reining the RF

The enormous task of managing and securing large enterprise wireless LANs (WLANs) has been a challenging one—that is until now. With the Cisco Structured Wireless-Aware Network, enterprises have a standards-based solution for managing and securing their WLANs that is scalable and cost-effective

Wireless Gems

New and planned capabilities in Cisco's WLAN product portfolio and software focus on improving your ability to centrally configure, upgrade, secure, and manage hundreds or thousands of wireless access points

Wireless Worlds Apart

Understanding the differences between business-class and home wireless solutions can help you choose the right one for your needs. Now, with Linksys products for the home and Aironet® solutions for businesses, Cisco offers WLAN solutions to meet the disparate requirements of both markets

Tips from the Trenches

Early adopters of wireless networks share their experiences with—and unexpected benefits of—WLAN technology

Mobile Data Services: Beyond the Hype

What are the real challenges to enabling users with anytime, anywhere, any-device access to the network? This article explores key considerations when rolling out a mobile data service in your enterprise

SPECIAL REPORT

PACKET

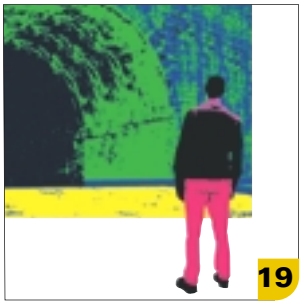

IP Tunneling Primer

Storage Networking Solution

Healthcare Technology

PACKET ONLINE EXCLUSIVE

Wireless Bridging: Applications and Deployment Considerations
Wireless bridges inexpensively interconnect LANs across enterprises and provide an economical backhaul option in service provider networks. What considerations should you make before deploying a high-speed, outdoor wireless network? Find out at cisco.com/go/packet.

PACKET MAGAZINE

DAVID BALL
EDITOR-IN-CHIEF

JERE KING
PUBLISHER

JENNIFER REDOVIAN
MANAGING EDITOR

SUSAN BORTON
SENIOR EDITOR

JOANIE WEXLER
CONTRIBUTING EDITOR

R.J. SMITH
SUNSET CUSTOM PUBLISHING
PRODUCTION MANAGER

MICHELE GERVAIS, MARK RYAN,
NORMA TENNIS
SUNSET CUSTOM PUBLISHING
PRODUCTION

JEFF BRAND
ART DIRECTOR

EMILY BURCH
DESIGNER

ELLEN SOKOLOFF
DIAGRAM ILLUSTRATOR

BILL LITTELL
PRINT PRODUCTION MANAGER

CECELIA GLOVER TAYLOR
CIRCULATION DIRECTOR

KEVIN TWOMEY
COVER DESIGN

SPECIAL THANKS TO THE FOLLOWING CONTRIBUTORS: PEGGY CASEY, JANICE KING, RHONDA RAIDER

ADVERTISING INFORMATION:
Kristen Bergman, 408-525-2542
kbergman@cisco.com
View *Packet* magazine at cisco.com/go/packet.

PUBLISHER INFORMATION:
Packet magazine (ISSN 1535-2439) is published quarterly by Cisco Systems and distributed free of charge to users of Cisco products. Application to mail at Periodicals Rates pending at San Jose, California, and additional mailing offices.

POSTMASTER: Please send direct address corrections and other correspondence to packet@external.cisco.com or to *Packet* in care of:

Packet Magazine
PO Box 2080
Skokie, Illinois 60076-9324
USA
Phone: 847-647-2293

FROM THE EDITOR

The James Bond Factor

WIRELESS LANS (WLANS) IN THE CONSUMER market have taken off with tremendous gusto; I wouldn't be surprised if the access point beats out the Xbox as the number one gift this holiday season. Why the sudden popularity? It could be because there are more multicomputer homes than ever before, and consumers are smartening up to the benefits of sharing one broadband connection to the Internet? Most people aren't interested in stringing Category 5 cable throughout their house, and are no doubt attracted to the elegant solution that WLANs provide—mobility being an added benefit. Or, it could be the simple fact that prices for WLAN gear have reached the sweet spot where most consumers are ready to jump in?

Most likely it is a combination of reasons. But there is another reason—one that most journalists and analysts tend to neglect. That is something I like to call the *James Bond Factor*.

Let's face it. Wireless is cool. Browsing the Web poolside. Checking your e-mail or Instant Messaging friends from your local coffee shop. Logging on to your corporate network from the airport. This is "Bond" gadgetry at its coolest, and it is only going to get cooler.

Wireless products are sprouting up in the market that will allow you to attach more and more "nontraditional" devices to the network. In fact, a slick gadget from Linksys (see page 42) actually lets you play music files on your computer or from the Internet over your home stereo system—without wires. How cool is that?

But no matter how much we techno geeks love our toys, it is the cool technologies that solve major hurdles at work that make most Netheads feel like Agent 007. Until recently, the task of managing and securing large enterprise WLANs was a challenging one—fraught with labor-intensive, manual processes. Today, if James Bond had to build a secure, manageable, scalable WLAN, "Q" would outfit him with the Cisco Structured Wireless-Aware Network (see page 36): a standards-based management solution that leverages your Cisco infrastructure to rein in the cost of operating large WLANs by providing the management features and tight controls you have come to expect from traditional wired networks.

One of those features that everyone expects and depends on is network availability. Our six-page special report on high availability networking (page 53) will help enterprises define, measure, and improve availability in their own networks. Service providers, look for an article on improving the availability of your networks in the next issue of *Packet*®.

Also, coming next issue is part two of the extremely popular "Deployment Diary" written by New York University's Jimmy Kyriannis (part one appeared in *Packet* Second Quarter 2003). Jimmy's tetrahedron core design will be fully deployed, so look for part two of his exploits in our Fourth Quarter issue.

While you're waiting, there is plenty to learn in this issue, including a primer on IP tunneling (page 19), migration strategies for the latest spanning tree protocols (page 23), IP telephony troubleshooting (page 7), another installment of the fanatically popular Reader Tips (page 16), and much more.

So, ask Money Penny to forward your calls, pour yourself a vodka martini—shaken, not stirred—and enjoy the read.

David A. Ball

BALL. . . DAVID BALL
Editor-in-Chief
daball@cisco.com

Mail

Same IPv6 Conversion, Different Result

While reading "IPv6 Time Is Now" [First Quarter 2003], I found a mistake or incongruence in the article.

In the configuration of the link-local address, the article states that you can automatically configure the address on any interface using the local-link prefix FE80::/10 (1111 1110 10). In the configuration of the SITE-local address, the article states that the site-local prefix is FEC0::/10 (1111 1110 11).

I think the conversion of the binary 1111 1110 10 into hexadecimal is actually FE20, not FE80. Similarly, the binary 1111 1110 11 in hexadecimal should be FE30, not FEC0.

I would also appreciate any information about the IEEE EUI-64 format mentioned in the article.

—Marcet Torres B, CCNA student, Santa Marta, Colombia

The following is a response from Packet® *article author Raj Gulani —Editors*

First, to make sure we are using the same conversion rule, here is the binary-to-hexadecimal conversion:

Binary	Hexadecimal
00000	**0**
0001	1
0010	**2**
0011	**3**
0100	4
0101	5
0110	6
0111	7
1000	**8**
1001	9
1010	A
1011	B
1100	**C**
1101	D
1110	**E**
1111	**F**

When converting from binary to hexadecimal, you split the binary number into groups of four and translate them into hexadecimal digits.

An IPv6 address prefix, in the format **ipv6-prefix/prefix-length,** represents bit-wise contiguous blocks of the entire address space. The **/prefix-length** variable is a decimal value that indicates how many of the high-order contiguous bits of the address comprise the prefix (the network portion of the address).

Therefore, for Link Local Prefix ,FE8O::/10, 10 is the prefix length using the above binary to hexadecimal conversion rule.

We split them into groups of four:

1111	**1110**	**10(00)**	**0000**
F	**E**	**8**	**0**

I believe you converted **1111 1110 0010 0000,** which would be FE20. Similarly, this would be the case for site local addresses.

For information about the EUI-64 format, refer to the following documents.

IPv6 for Cisco IOS® Software: cisco.com/packet/153_2a1

RFC 2373, IPv6 Addressing Architecture: ietf.org/rfc/rfc2373.txt

Guidelines for 64-Bit Global Identifier (EUI-64) Registration Authority: standards.ieee.org/regauth/oui/tutorials/EUI64.html

Sci-Fi Art, Mixed Reviews

I have just received the First Quarter 2003 issue of *Packet*. Is there any chance of getting an electronic version of the cover image? Any Star Trek fan will know why that image is so great! The higher the resolution, the better!

—Ivar Holck, Comnet, Bergen, Norway

You can download a PDF of the cover image at cisco.com/packet/153_2a2 —Editors

I have been reading *Packet* for four years, and I appreciate the great job you do of publishing a technical magazine that appeals to both novices and advanced networking professionals. But because I care so much about my favorite magazine, I have to say something about the recent design [First Quarter 2003]. It is awful. The article "A Good Defense" wins the worst design award for 2003. The cover of the issue, especially the title, is very dull. Yellow text on yellow background? Star Wars-like patterns inside badly done bubbles? I know you can do better.

What happened to the images of the cute Cisco people reminiscent of the 1920s and 30s? Last year's wireless issue [Second Quarter 2002] was my favorite. Don't get me wrong. I still think the issue has great content, but please, please, hire some good graphic designers and keep your readers happy.

—Tomasz Oczapowski, Pacific Publications, New South Wales, Australia

Read *Packet* on the Move

I subscribe to the print version of *Packet* and would like to download a PDF of the magazine. Having the ability to view the magazine in this format is a great way to read it on the move.

—Name withheld

To download a PDF version of any Packet *article, visit cisco.com/go/packet and scroll down to the bottom of the article. You will find options to download the article alone or the entire issue.*

SEND YOUR COMMENTS TO *PACKET*

We welcome your comments and questions. Reach us through e-mail at **packet-editors@cisco.com**. Be sure to include your name, company affiliation, and e-mail address. Letters may be edited for clarity and length. **Note**: The *Packet* editorial staff cannot provide help-desk services.

User Connection

New Web Portal Helps Certified Customers Share Knowledge

A NEW COMMUNITY WEB portal for Cisco certified network professionals provides an interactive forum for customers with valid Cisco certifications (such as CCIE®, CCNA®, CCDA®, CCNP®, CCDP®, CCIP®, or CCSP) to exchange information with Cisco subject matter experts, employees, partners, resellers, and other customers, as well as share and gather networking best practices.

Currently in the pilot phase, the Cisco Certifications Community offers easy access to news, discussion groups, video presentations, white papers, articles, and other resources. In addition to covering major technologies such as IP voice, storage, wireless, and security, the Cisco Certifications Community provides updates on the Cisco certification program, including recent changes, recertification requirements, and upcoming courses and exams. Other useful features include tips on passing CCIE lab exams and "Ask the Experts" and "Tech Talk" video-on-demand segments.

"Like courses and exams, communities also provide lifelong learning opportunities for our certified professionals," says Don Field, senior manager in the Internet Learning Solutions Group at Cisco.

"Besides offering practical advice on networking to Cisco certified individuals, the certification community portal will allow Cisco certified individuals to share knowledge—perhaps one of the best ways to learn," adds Field.

KNOWLEDGE SHARING: The new Cisco Certifications Community offers certified individuals practical advice and the opportunity to share learning.

To view and participate, visit cisco.com/go/certcommunity. Access to the portal is available at no charge to Cisco certified individuals who are registered users of Cisco.com. ▲▲

Number of CCIE Experts Swells to Half Million Worldwide

CISCO CAREER CERTIFICATIONS, one of the IT industry's fastest-growing certification programs, crossed the half million mark this year.

The number of Cisco certifications issued worldwide has doubled in less than two years. The Cisco Career Certifications program, which includes the highly regarded CCIE® expert-level certification and the popular associate-level CCNA®, has positioned Cisco as the networking industry's foremost source of IP expertise.

In the past year, Cisco certifications have garnered industry accolades from trade publications such as CertCities.com for "Best Overall Certification Program." Cisco was also recently recognized by the US National Security Agency and the Committee on National Security Systems as the only industry vendor to meet the 4011 training standard for information assurance professionals in federal departments and agencies.

"The continuing growth and adoption of Cisco career certifications, particularly in security and IP telephony, is evidence of increasing confidence in the competence of certified professionals," says Tom Kelly, vice president, Internet Learning Solutions Group at Cisco.

Industry leaders and analysts have long attributed Cisco's market-share gains to the cadre of Cisco certified individuals in the marketplace, which includes 9500 CCIE worldwide.

The widely respected Cisco Career Certifications program brings valuable, measurable rewards to network professionals, their managers, and the organizations that employ them.

Cisco offers three levels of general certification—from Associate to Professional to Expert—available in various paths (or tracks) and designations. A variety of Cisco Qualified Specialist focused certifications are available to gain knowledge and skills in specific technologies, solutions, or job roles.

For more information on certifications, visit cisco.com/go/certifications. ▲▲

Best Practices for IP Communications Services

An ever-increasing number of companies worldwide are recognizing that IP communications can fundamentally change and significantly improve the way business is conducted. Indeed, IP communications has the potential to be a strategic benefit to virtually every business that takes the time to properly plan, thoroughly evaluate, and effectively implement its applications, technology, and ongoing support requirements. And as more companies and organizations adopt converged IP networks and IP communications solutions, they need to know that these solutions will be reliable and scalable, and that they can integrate with the existing network environment. As a result, a need has arisen for effective support to aid in the successful installation and operation of these complex solutions.

Cisco and its partners develop and support some of today's most complex IP communications solutions. Cisco has formulated a wealth of innovative best practices to help customers and partners create tailored services portfolios that can accelerate the successful deployment and implementation of IP communications solutions, as well as ensure consistency and quality, and ultimately, increase customer satisfaction.

A straightforward, comprehensive guide to the standard services that are required with a Cisco IP communications solution, the *Cisco IP Communications Services Blueprint* is based on customers' needs, partners with highly qualified and well-trained services providers, and best-in-class tools and technology. The blueprint describes the critical services deliverables required to plan, design, implement, and operate a Cisco IP communications solution. The blueprint, which also includes reference pricing, is available at cisco.com/packet/153_3b1.

Steps to Success Web Site

For Cisco partners that provide IP communications solutions to their customers, the Steps to Success Web site offers vital IP communication services life-cycle methods, procedures, tools, and proven approaches to services that can help companies deploy and operate complex IP communications solutions. Visit the Web site at cisco.com/go/stepstosuccess (requires Cisco.com partner login).

IP Telephony Readiness Tool

The IP Telephony Readiness Assessment Web-based Tool has been updated. Offered by Cisco Advanced Services, this tool helps organizations confirm whether they have met common network design requirements for a high-availability Cisco IP telephony implementation, and whether the planned IP telephony solution meets common identified design and configuration criteria.

Now this popular tool, which has been widely used to perform high-level assessments of customers' readiness to deploy IP telephony, has been refreshed with the latest technical content.

Tips for Adopting IP Communications

If your company is in the planning stages of IP communications implementation, consider these four important guidelines.

1. Form a Multidisciplinary IP Communications Planning/Evaluation/Implementation Team

Create a small, hands-on IP communications program team that includes representatives of the IT, finance, and purchasing organizations, as well as strategic business units. A team composed of qualified representatives from these organizations will provide a good balance of business, economic and technical goals, objectives, benefits and issues.

2. Align IP Communications-Enabling Capabilities and Costs with Business Goals and Objectives

Identify and quantify the near and longer-term benefits that can be delivered by effective IP telephony implementation and support. These should include productivity gains that can be achieved through improved applications access and sharing for mobile and remote staff members—plus expense reduction, containment, and avoidance opportunities. Weigh these benefits against the immediate and ongoing IP telephony expenses in order to determine the return on investment.

3. Conduct a Thorough Network Readiness Audit

Determine the readiness of your network infrastructure to efficiently and effectively accommodate IP communications. Identify and quantify the costs associated with the required upgrades to the network infrastructure. Recognize that the upgrades will also facilitate a broad range of other necessary applications (beyond just IP telephony), so allocate the costs appropriately.

4. Evaluate Technology and Support Models

Focus on the ability of the channel to provide meaningful support in the planning, implementation, and ongoing operational phases. Thoroughly examine the experience and capabilities in the areas of planning, diagnostics, responsiveness, and problem solving. The technology and applications are important, but world-class service and support are essential.

For more information on Cisco IP communications services, visit cisco.com/go/ipcservices.

The Readiness Assessment tool takes you through a series of questions that results in a summary report with red, yellow, or green readiness status in several technical areas. The report also provides quick visual pointers to key areas that might require further study or evaluation and includes an option for a more detailed printable report.

Available to Cisco customers and partners at cisco.com/go/iptassessment (Cisco.com login required), the tool also has a link to the IP telephony support available from the Cisco Advanced Services team.

For the majority of companies, migrating to IP communications is no longer a question of "if," but rather a matter of "when." While internal and external factors will continue to influence the timing of migration to IP communications, waiting too long to deploy IP-based communications can cause some companies to be competitively disadvantaged in a relatively short time. ▲▲

Cisco Worldwide Events

August 11–15	HP World 2003	Atlanta, Georgia, USA
August 25–29	IBM NSTC	Miami Beach, Florida, USA
October 12–18	ITU Telecom World 2003	Geneva, Switzerland
October 30–November 1	Networkers China 2003	Beijing, China
November 24–26	Networkers South Africa 2003	Johannesburg, South Africa
December 9–12	Networkers Brazil 2003	Sao Paulo, Brazil

cisco.com/warp/public/688/events.html

Cisco Learning Credit Program Streamlines Training Process

Deploying new networking technologies requires a skilled and knowledgeable workforce. But when purchasing networking solutions and equipment, decision makers often address training as an afterthought to their technology purchase. Now, through the Cisco Learning Credit Program, managers can define their networking and training needs concurrently and prepay for high-quality training from more than 150 authorized Cisco Learning Partners.

The Cisco Learning Credit Program, currently available in the US and Canada only, enables you to simultaneously purchase necessary training and Cisco hardware, software, or services. And the program streamlines the administrative process by allowing you to manage a single budget and purchase order, as well as reduce expense reimbursements.

Identify Training Needs

Learning partners can help you more effectively identify your network training needs no matter how complex, and develop a plan to meet them.

The ability to prepay for training, using learning credits, makes it easier for you to add training to your overall Cisco purchase.

How the Program Works

Authorized Cisco Learning Partners work with you to develop a tailored training program to determine how many Cisco learning credits you will need. After designing a training program and obtaining the appropriate number of credits, you can redeem the credits at participating worldwide Cisco Learning Partner locations for Cisco instructor-led training courses, online courses, live labs, and remote labs. Customers can choose from more than 190 courses.

Cisco Learning Credits are managed in an online database that tracks transactions, validates redemptions, and sends monthly statements. The Web-based Learning Credit Management tool, available 24 hours a day, lets you view credit balances, review account transactions, generate reports and monitor courses taken by individuals and departments.

Pricing and Availability

You can purchase Cisco Learning Credits along with a hardware purchase or as a standalone item for an integrated network solution. Each Learning Credit pays for US$100 of training from a participating Cisco Learning Partner. Valid for one year, credits are available for purchase in packs of 100, 500, or 1500, and can be redeemed according to how many you need.

Cisco Learning Partners are the only authorized providers of Cisco training content. They must meet stringent guidelines for authorization and their instructors must pass rigorous exams to become certified Cisco instructors.

For more on the Cisco Learning Credit Program, visit cisco.com/go/learningcredits.

The Learning Credits program will expand to countries outside the US later this year. ▲▲

My NetPro

Now you can personalize your *Cisco Networking Professionals Connection* Web page. To customize your page based on interest in forums and topics, posts, subscriptions, bookmarks, and preferences, visit cisco.com/packet/153_3g1.

JAMES CARRIER

PACKET WINS A MAGGIE

In April *Packet*® was honored with a Maggie Award from the Western Publications Association (WPA) at the 52nd annual awards ceremony in Los Angeles, California. *Packet* took home the trophy for "Best Quarterly/Trade Publication." Other nominees in the category were *Dell Insight,* Oracle's *Profit Magazine, Launchpad,* and *Tribal College Magazine.*

Tech Tips

New Cisco IOS® Software packaging simplifies image selection. Cisco IOS Packaging consolidates the number of Cisco IOS Software packages in each software release by using consistent package names across hardware platforms. Designed for Cisco 1700, 2600XM, 2691, 3725, and 3745 access routers, the first phase of Cisco IOS Packaging is available now in Cisco IOS Software Release 12.3. cisco.com/packet/153_4e1

Attend a mentoring session in advanced networking technology. Cisco Network Deployment Mentoring courses are conducted by Cisco networking experts with the latest information on integration and deployment practices. Participants will receive mentoring and guidance on using the Cisco deployment best practices. Course attendees receive training in network design, implementation, operation, and integration. Courses include intensive hands-on lab sessions. cisco.com/packet/153_4e2

Find new, improved troubleshooting information for Cisco hardware platforms. The Cisco TAC Web site's Hardware Troubleshooting index page provides new and improved documentation that can help you troubleshoot potential hardware issues and identify what might be causing a hardware failure. cisco.com/packet/153_4e3

Get automatic end-of-sale notifications about specific products. The Cisco Product Alert tool (formerly the Field Notice tool) can now send automatic end-of-sale e-mail notifications for Cisco products. To set up a profile in the tool, name your profile, select the products you're interested in, and enter your e-mail address. The tool sends you e-mail updates about reliability, safety, network security, and end-of-sale issues for the Cisco products you've specified. (available to registered users only). If you already have a profile, you'll receive end-of-sale announcements automatically for the products you selected. cisco.com/packet/153_4e4

IP Telephony Detective

Developing a Troubleshooting Methodology for Cisco IP Telephony

BY PAUL GIRALT

In a perfect world, networks would diagnose problems for you, tell you the root cause, and resolve problems automatically. In the real world, problems arise and it is up to you to troubleshoot and bring them to resolution. IP telephony networks are particularly challenging to troubleshoot because they involve a wide range of components that work as a system to provide voice and data services. When troubleshooting a problem, you should employ a solid methodology for dividing the problems into smaller, manageable components and investigate each one until you find the root cause.

Troubleshooting problems can be broken down into two stages: *data gathering* and *data analysis* (although your analysis might lead you to collect additional data). Generally, these two stages encompass the following:

1. Gather data about the problem:
   a. Identify and isolate the problem.
   b. Use topology information to isolate the problem.
   c. Gather information from end users.
   d. Determine the problem's timeframe.
2. Analyze the collected data:
   a. Use deductive reasoning to narrow the list of possible causes.
   b. Verify IP network integrity.
   c. Determine the proper troubleshooting tool(s), and use them to find the root cause.

Sometimes, several service-affecting problems occur simultaneously. In fact, this is not uncommon because multiple problems often manifest themselves as symptoms of the same root cause. When multiple problems occur simultaneously, focus on the problem that has the greatest impact on users. For example, if some users are reporting dropped calls and others are reporting occasional echo, the two problems are probably unrelated. Troubleshoot the dropped-call problem first, because keeping calls connected is more critical than removing the occasional echo on an active call.

Stage 1: Gather Data About the Problem

Half the battle in troubleshooting a problem is determining which piece of the puzzle is the source of the problem. With so many different components to an IP telephony network, the first step is to isolate the problem and, if multiple problems are being reported, determine which of them might be related to each other.

Identify and isolate the problem. You must determine which parts of the problem are symptoms and which are the root cause of the problem. For example, if a user complains of an IP phone resetting itself, it's logical to assume that something is wrong with the phone. However, the problem might lie with Cisco CallManager or one of the routers and switches that make up the underlying data network. So, although the symptom is a phone reset, the root cause could be a WAN outage or CallManager failure.

Always remember to look at the big picture when searching for the root cause and not let the symptoms of the problem steer you in the wrong direction. To help you visualize the big picture, detailed topology information is essential.

Use topology information to isolate the problem. One of the first lines of defense in the troubleshooting process is possessing current topology information. One of the most important pieces of topology information is a detailed network diagram, which should include network addressing information and the names of all the devices. It should also clearly show how the devices are interconnected and the port numbers used for the interconnections.

In addition to a network diagram, you should use some method to store information such as IP address assignments, device names, and password information. For a small network, you can use something as simple as a spreadsheet or even a plain text file. For larger deployments, a database or network management application, such as CiscoWorks, is recommended.

You also need documentation of your dial plan. Some deployments—especially those heavily using toll bypass—have very complex dial plans. Knowing where a call is supposed to go just by knowing the phone number and from where it is dialed helps you quickly understand a problem.

When your topology information is complete, it should include all the following information:

- Interconnection information for all devices, including device names and port numbers. If any patch panels exist between devices, the port numbers should be listed.
- IP addressing for all network devices (routers, switches, and so on)
- IP addressing for all telephony and application servers and voice gateways (including data application servers)
- IP addressing for endpoints; that is, scopes of a Dynamic Host Control Protocol (DHCP) pool
- WAN and PSTN service provider names and circuit IDs for each circuit
- Spanning-tree topology, including root bridges for all virtual LANs (VLANs) and which ports should be forwarding and blocking
- Dial plan information
- Software version information for all devices

If you are troubleshooting a network you didn't design, topology is one of the first pieces of information you should obtain, if it is available. If it is not available, obtain

the information from someone who is familiar with the network, and then make a quick sketch. A general topological understanding of the network, or at least the piece of the network in question, helps when you are trying to differentiate the problem from its symptoms.

What is worse than not having topology information? Having *incorrect* topology information can lead to countless hours heading down the wrong path. If you are going to keep topology information, make sure it is updated often.

Use all the topology information you have to narrow down which pieces of the network might be involved in the problem you are troubleshooting. To further isolate the problem, interview the end users who reported the problem to gather additional information.

Gather information from end users. The more detail about the problem you can gather before you begin troubleshooting, the easier it is to find a resolution—and that means less frustration for you. Here is some information you should collect from users:

- Details about exactly what the user experienced when the problem occurred
- Phone numbers for all parties involved in the problematic call or calls (you can use this as search criteria if you need to look through traces)
- Actions performed by the user when the problem occurred, including which buttons were pressed and in what order
- End-user observations, including text messages displayed on the phone or recorded announcements
- Information about the user's device. For example, if the user experienced a problem while using a Cisco IP Phone 7960, get the phone's MAC address and IP address, along with registration information and any other statistics available from the phone.

The information provided by end users might not be enough to even begin troubleshooting. Sometimes the proper diagnostic tools are not enabled when the problem occurs, forcing you to ask the users to inform you the next time the problem occurs. Be sure to turn on tracing or debugs before making the request so that when the problem occurs again, you will capture the data. Also, point out to users the importance of letting you know immediately when a problem occurs, as many of the diagnostic trace files overwrite themselves within several hours or days (depending on the amount of traffic on your system).

> Always remember to keep the IP network in mind and look at every layer in the OSI model, starting from Layer 1. Check your physical layer connectivity. . . . Then make sure you have Layer 2 connectivity by checking for errors on ports, ensuring that Layer 2 switches are functioning properly, and so forth.

Determine the problem's timeframe. In addition to *what* the problem is, you should try to determine *when* the problem occurred. Determining the earliest occurrence can help correlate the problem with other changes that might have been made to the system or other events that occurred around the same time, such as scheduled maintenance windows or known network outages.

If relying on end users to provide the "when" information, ask them to note the time on their IP phone when the problem occurred. The phone's time should be synchronized with the clock on the Cisco CallManager to which the phone is registered. As long as you have your CallManager servers and network devices time synchronized, getting a phone-based time from the user makes finding the proper trace files very easy.

Time synchronization throughout your network is vital to troubleshooting problems, especially when many devices are involved. For example, a call could traverse several CallManager servers, interactive voice response (IVR) servers, voice-mail systems, and gateways—each of which has its own clock to write timestamps to the log files. The best way to synchronize the devices in your network is to use the Network Time Protocol. To find out more about this topic, see "How to Configure Time Synchronization for Cisco CallManager" at cisco.com/packet/153_4a1.

Although it is important to use information about when the problem started happening, it is equally important to not assume that the problem was a direct result of an event. For example, if a user reports a problem the day after an upgrade was performed on CallManager, the upgrade might have caused the problem—but you should not automatically assume that this is the root cause.

Stage 2: Analyze the Collected Data

Now that you have collected data from a variety of sources, you must analyze it to find the root cause and workaround for your problem.

Use deductive reasoning to narrow the list of possible causes. The next part of your fact-finding mission is to identify the various components that might be involved and to eliminate as many of them as possible. The more you can isolate the problem, the easier it is to find the root cause.

Although not all of the following factors apply to every problem, where applicable, you must check the following pieces involved in the call. You can use your topology information to help obtain this information.

- Cisco CallManager nodes involved in the signaling
- Network devices that signaling, voice traffic, or both traverse
- Gateways or phones involved in the call

Continued on page 12

PAUL GIRALT is an escalation engineer with the Cisco Technical Assistance Center (TAC) in North Carolina and trains TAC teams around the globe. He has been troubleshooting complex IP telephony networks since Cisco CallManager release 3.0, and has resolved problems for some of the largest Cisco IP telephony deployments. Co-author of the Cisco Press book *Troubleshooting Cisco IP Telephony* (ISBN: 1-58705-075-7), Giralt holds a bachelor's degree in computer engineering from the University of Miami, Florida. He can be reached at pgiralt@cisco.com.

IP Telephony Troubleshooting, Continued from page 10

Methodology for Troubleshooting Cisco IP Telephony

Troubleshooting Cisco IP Telephony: A Cisco AVVID Solution from Cisco Press lays out the methodology you need to identify and resolve complex problems in an IP telephony network. Written by Paul Giralt, Addis Hallmark, and Anne Smith, the book provides comprehensive coverage of all parts of a Cisco IP telephony solution, including Cisco CallManager, IP phones, gateways, analog devices, database and directory replication, call routing, voice mail, applications, network infrastructure, and more.

You'll learn how to read trace files and when to turn on tracing, IOS voice debugging, and how to troubleshoot voice quality issues. The book also shows you how to break down problems to find the root cause.

Descriptions of each part of a Cisco IP telephony solution are provided, to help you in understanding the functionality of each part, how each part interacts with other parts in the solution, and what steps to take and tools to use to identify and resolve the cause of a problem.

Troubleshooting Cisco IP Telephony (ISBN: 1-58705-075-7) is available at online booksellers such as Amazon.com, local bookstores, computer and electronics superstores, and other venues where computer technology books are sold. For information on corporate bulk or education purchases and sales outside the US, visit ciscopress.com/sales.

- Other devices involved, such as conference bridges or transcoders

Concentrate your energy on the smallest subset of devices possible and try to find a commonality to multiple users' problems. For example, if all the users on a particular floor are having the same problem, concentrate on the problem one user on that floor is having and determine which device in the network the users share that might be causing the problem. If you fix the problem for that one user, in most cases you will fix it for all the affected users.

Verify IP network integrity. Remember that your IP telephony network is only as good as your IP network. A degraded network or a network outage can cause a wide range of problems, from slight voice quality degradation to a total inability to make or receive calls on one or more phones.

Network health is especially important when troubleshooting voice quality problems because most voice quality problems stem from packet delay and/or loss due to network congestion or degradation.

Always keep the IP network in mind and look at every layer in the OSI model, starting from Layer 1. Check your physical layer connectivity (cables, patch panels, fiber connectors, and so on). Then make sure you have Layer 2 connectivity by checking for errors on ports, ensuring that Layer 2 switches are functioning properly, and so forth.

Determine the proper troubleshooting tool. After you narrow down the appropriate component(s) potentially causing the problem and have detailed information from the user(s) experiencing the problem, you must select the proper troubleshooting tool or tools. Use the tracing and debugging facilities available in Cisco CallManager and other devices to determine exactly what is happening.

Because CallManager is central to almost all problems, you will typically want to look at a Cisco CallManager (CCM) trace file on the CallManager(s) involved in the problem call. This step is the most demanding on your troubleshooting skills, because you analyze the detailed information provided in the various tools and use it to search for additional clues using other tools.

Some tools you can use to troubleshoot problems in an IP telephony network include the following:

- CallManager trace files: CCM, Signal Distribution Layer (SDL), computer telephony integration (CTI), and so on
- Application server trace files (MIVR, JTAPI)
- IOS® Gateway **debug** and **show** commands
- Q.931 Translator
- Dick Tracy Utility (for troubleshooting WS-X6608 and WS-X6624 gateways)
- Call detail records
- Microsoft Performance (PerfMon) or CallManager Serviceability's Real-Time Monitoring Tool
- Network Packet captures (Sniffer Pro, Ethereal, for example)

Remote Access to All Devices

To make troubleshooting your network as easy as possible, ensure you have remote access to all the devices on your network. Use either Secure Shell (SSH) or Telnet to access routers and switches. Enable Windows Terminal Services or Virtual Networking Computing (VNC) on your Cisco CallManager servers and application servers to access them remotely.

If there is a chance that you will have to troubleshoot a problem when you are not on your corporate network, consider adding a virtual private network (VPN) concentrator to your network, which will enable you to connect to the network regardless of where you are.

◆ ◆ ◆

When troubleshooting IP telephony problems, remember to first put on your detective hat and gather enough information to isolate the problem to a few components of the system. Then dig deeper into each component by dividing the problem into more manageable pieces. Finally, apply your expertise to each of these smaller pieces until you find the resolution to your problem. ▲▲

Hidden Gems

Do you know about these five IOS CLI commands?

BY MARK BASINSKI

Among the many IOS® command-line interface (CLI) commands, there are several "hidden gems"—commands that have been implemented relatively recently, or aren't well documented, or, for one reason or another, are not as well known as they should be.

In this article, let's take a look at the following five CLI hidden gems:

- Do command
- Privilege command enhancement
- Show output line numbering
- Show redirect
- Console memory reservation

Do Command

Introduced in Cisco IOS Software Release 12.2(5)T, the **do** command allows you to run EXEC-level commands in global configuration mode or other configuration modes or submodes. You can use this command during configuration, such as an interface, if you need to check other parts of the configuration (for example, the Internet address).

Syntax:
do <EXEC-level command>

EXAMPLE DO COMMAND

```
3600-4(config)#int e2/0

3600-4(config-if)#do show int e2/1

Ethernet2/1 is up, line protocol is
down

  Hardware is AmdP2, address is
  00b0.643d.f321 (bia 00b0.643d.f321)

  Internet address is 200.40.89.1/24
```

Privilege Command Enhancement

Introduced in Cisco IOS Software Release 12.2(11)T, the enhanced **privilege** command allows privilege levels to be set more efficiently (in the past, each level needed to be set independently). Use the **privilege** command whenever a privilege level is set.

Syntax:
privilege <mode> all level x <submode>

In the example shown below, a privilege level of 5 is assigned. Only users with level 5 or higher privilege will be allowed to configure ARP.

EXAMPLE ENHANCED PRIVILEGE

```
3600-4(config)#privilege config all
level 5 arp
```

Show Output Line Numbering

Introduced in Cisco IOS Software Release 12.2(4), this command displays line numbers to a **show** output. It is especially useful when you call the Cisco Technical Assistance Center (TAC) and need to identify a particular portion of a very large configuration.

Syntax:
show run linenum

EXAMPLE LINE NUMBERING

```
3600-4#show run linenum

Building configuration...

Current configuration : 209254 bytes

   1 : !

   2 : ! Last configuration change at
       22:02:42 PDT Mon May 3 1993

   3 : ! NVRAM config last updated at
       21:49:44 PST Sat Mar 13 1993

   4 : !

   5 : version 12.2

   6 : service timestamps debug
       datetime msec localtime

   7 : service timestamps log
       datetime msec localtime
.....
3600-4#show run linenum | include rtr

  88 : rtr responder

  90 : rtr 1

  92 : rtr schedule 1 start-time
       12:54:00 Mar 19

  93 : rtr 2

  95 : rtr schedule 2 start-time
       09:45:00 Mar 19

  96 : rtr 3

  98 : rtr schedule 3 start-time
       11:02:00 Mar 19

  99 : rtr 4  101 : rtr schedule 4
       start-time 11:02:00 Mar 19
```

MARK BASINSKI, CCIE® No. 4422, is a product manager in the Internet Technologies Group at Cisco, focusing on IOS configuration, programmatic interface, and SNMP infrastructure. At Cisco since 1997, Basinski has been a senior TAC support engineer and team lead for network management, and a technical marketing engineer in the Enterprise Management Business Unit with a focus on CiscoWorks products and IP telephony management. He can be reached at mbasinsk@cisco.com.

Show Redirect

Introduced in Cisco IOS Software Release 12.2(8)T, this command allows the **show** output to be directed to other devices, such as slot, ftp, tftp, or disk. You can use the **show redirect** command to analyze **show** output offline (see example).

Syntax:
show <command> | redirect <URL>

EXAMPLE SHOW REDIRECT

```
Gsr-10#show version | redirect
slot0:ver.txt
Gsr-10#show run | redirect
tftp://171.69.1.129/running.txt
```

Console Memory Reservation

Introduced in Cisco IOS Software Release 12.0(11)S, this command reserves memory for console access. Without reserving console memory, when the router is very low in memory or memory is highly fragmented, console access isn't allowed. Simple servicing of the router is denied. With this command, a small block of memory is reserved to ensure console access.

Syntax:
memory reserved console <nKB>

show memory console reserve

Note that the general rule is to reserve memory three times that of NVRAM. NVRAM memory can be seen with the **-dir nvram:** command. ▲▲

EXAMPLE CONSOLE MEMORY

```
Router#dir nvram:
Directory of nvram:/
  123 -rw- 1962 <no date>
  startup-config
  124 ---- 46 <no date> private-config
  1 -rw- 0 <no date> ifIndex-table
  2 -rw- 12 <no date> persistent-data
  3 -rw- 12 <no date> persistent-data
129016 bytes total (123884 bytes free)
Router#memory reserved console 387
(three times of 129K)
Router#show memory console reserve
Memory reserved for console is 262144
bytes
```

One Stop for Network Management Tools and Resources

Need the latest list of management Information Bases (MIBs) supported by Cisco devices? What about Simple Network Management Protocol (SNMP) traps? Or descriptions of command-line interfaces (CLIs) for all Cisco network management products?

These are among the most requested and downloaded items on the Cisco Technical Assistance Center (TAC) Web site. Now you can find this information and other key resources to help manage Cisco devices on your network in one centralized location with the online Cisco Network Management Toolkit. Included in the toolkit are:

- A comprehensive listing of the latest MIBs supported by Cisco products. Arranged by product category and available for individual download, there are MIB files for Simple Network Management Protocol (SNMP) version 1 and SNMPv2, including SNMPv1 conversions of SNMPv2 MIBs. (This list replaces the former *Cisco MIB User Quick Reference*.)
- A directory of application notes for using Cisco MIBs
- A directory of helpful MIB-related scripts and files
- Currently available object identifiers (OIDs)
- Currently available traps
- Product documentation describing CLIs for all Cisco network management products
- An archive of MIBs, OIDs, and schemas for Cisco IOS® Software 10.0 and earlier releases
- Cisco NetFlow data sheets, white papers, and other information. A feature of Cisco IOS Software, NetFlow collects and measures traffic data as it enters specific routers or switch interfaces.

Software Support Pages

The Network Management Toolkit also includes a compendium of technical-support documentation created by the Cisco TAC: configuration and design guides, troubleshooting documents, Web-based tools, and information on product features and versions.

The technical-support pages cover a range of Cisco network management software products, including CiscoWorks, Resource Manager Essentials, Element Management Framework, Cisco WAN Manager, CiscoWorks Wireless LAN Solution Engine, FlowAnalyzer, Network Registrar, Cisco Cable Manager, and many more.

Check out the Cisco Network Management Toolkit at cisco.com/packet/153_4f1. Access to full toolkit resources requires Cisco.com login.

Software Center Updates

The Cisco Network Management Toolkit resides in the Software Center on Cisco.com, under "Network Management Software." The Software Center provides the latest updates and versions and releases of all Cisco software products, from IOS to network management to Aironet® wireless. Registered Cisco.com users can find the Software Center at cisco.com/packet/153_4f2. ▲▲

Ask your peers and Cisco experts questions or share your own knowledge about SNMP traps, MIBs, and other related management topics at the *Cisco Networking Professionals Connection* Network Management forum: cisco.com/discuss/infrastructure.

Virtual Private Networks At a Glance

Sponsored by Cisco Press

Why Should I Care About VPNs?

A virtual private network (VPN) is a set of solutions and technologies designed to make secure (encrypted) site-to-site and remote-access connections over public networks. VPN connections can be set up quickly over existing infrastructures and provide an excellent alternative to dedicated private networks such as Frame Relay and ATM. Benefits of VPNs include:

Cost savings: VPNs use cost-effective public IP networks to connect remote office users to the main corporate site, eliminating expensive dedicated WAN links.
Security: VPNs provide a high level of security using advanced encryption and authentication protocols.
Scalability: VPNs can be easily set up over the existing Internet infrastructure, allowing corporations to add large amounts of capacity without adding significant infrastructure.
Compatibility with broadband technology: VPNs allow mobile workers, telecommuters, and "day extenders" to connect to their corporate network via high-speed, broadband connectivity, such as DSL and cable.
Ease of access: Network access can be provided from anywhere in the world by local Internet access points of presence (POPs).

VPNs offer almost the same level of information security as traditional private networks, and can be simpler to set up, less expensive to operate, and easier to administer.

What Problems Need to be Solved?

The two primary technical considerations in setting up VPNs are **tunneling** and **encryption**.

Tunneling is the process of encapsulating the protocol header and trailer of one network protocol into the protocol header and trailer of another. Prior to the packet being sent across the network, it is encapsulated with new header information that allows an intermediary network to recognize and deliver it. When the transmission ends, the tunneling header is stripped off, and the original packet is delivered to the destination. While tunneling allows data to be carried across a third-party network, it does not protect data against unauthorized inspection or viewing. To ensure tunneled transmissions are not intercepted, traffic over a VPN is typically **encrypted**. It is important to realize, however, that encrypted data can still be intercepted, and attempts can be made to decrypt captured data.

Deployment Modes

Site-to-site VPNs link company headquarters, remote locations, branch offices, and e-business partners to an internal network over one shared infrastructure. Site-to-site VPNs can be intranets or extranets.
Remote-access VPNs allow corporate users and mobile workers to access a corporate intranet securely by using their cable, DSL, or the local numbers of an ISP to dial in and connect to the network. Leveraging local ISP dialup infrastructures enables companies to lower communications costs and increase productivity because of the robust technology that supports the Internet and other public networks.

VPN Architecture

There are several methods (both Layer 2 and Layer 3) and technologies available for establishing a VPN. VPNs can be established and managed on the customer premises or over the network by a service provider. You can also combine several methods simultaneously to meet specific needs.

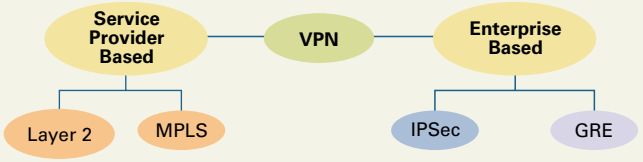

Generic Routing Encapsulation

Generic Routing Encapsulation (GRE) is a workaround method for routing packets over an IP network that would otherwise not be routable. It can also be used for routing multicast packets over noncompatible networks. GRE can also be used to route non-IP protocols (for example, AppleTalk, IPX) over IP networks. The figure below shows the concept of encapsulating a packet. When the GRE packet reaches the destination, the GRE header and trailer are stripped off and the original protocol works as normal.

Putting It All Together

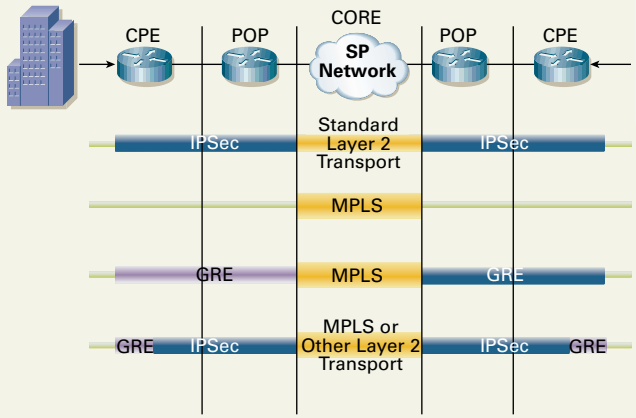

IPSec tunnels over a service provider network. This is a trusted virtual point-to-point connection.

MPLS VPN. Packets are assigned labels at the ingress of the service provider network and routed according to forwarding instructions.

Simple GRE over a Layer 2 transport. Could be used for transporting AppleTalk over an IP network.

This GRE/IPSec scenario could be implemented by a company to transport AppleTalk over a secure link, via a service provider utilizing MPLS or other transport.

IP Security

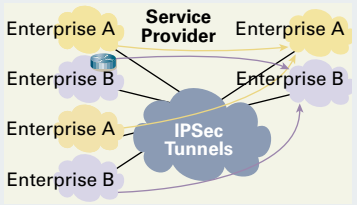

IP Security (IPSec) is an Internet standard for establishing and managing data privacy between network entities over unprotected networks, such as the Internet. IPSec services are provided at the network layer, allowing simple and effective encryption of IP traffic. Prior to IPSec, data had to be encrypted on an application basis, or hardware encryptors had to be placed at all intermediate points in the network. With IPSec, encryption takes place at the network edge, which is faster and cheaper.

Multiprotocol Label Switching

Multiprotocol Label Switching (MPLS) uses a method of forwarding packets that is based on labels. The labels may correspond to IP destination networks, as in traditional IP forwarding, but they can also represent other parameters, source addresses, quality of service (QoS), or other protocols. MPLS implements label swapping between different modules within the network. The two main components in an MPLS network are the **control plane** and **data plane**. The control plane takes care of routing exchanges and the exchanging labels between adjacent devices. The data plane forwards packets based on labels and is independent of routing or label exchange protocols. MPLS is designed for use on virtually any media and Layer 2 encapsulation. For frame-based encapsulations, MPLS inserts a 32-bit label between the Layer 2 and Layer 3 ("frame-mode" MPLS). MPLS allows service providers to offer their customers similar services as those found in Frame Relay or ATM networks, with the conveniences of an IP network. With MPLS networks, VPNs operate as logical "ships in the night" across a common routed backbone. The VPN appears as a "private" routed WAN to the enterprise.

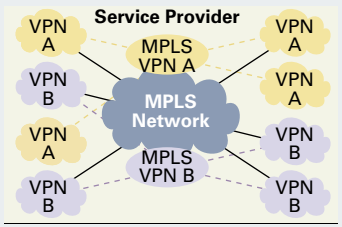

Cisco Systems

PACKET

Reader

Packet® *reader tips continue to stream in from all parts of the globe, and for that we thank you. This issue includes networking wisdom from Australia, Canada, India, the Netherlands, Pakistan, Poland, Serbia and Montenegro, and the United States. While every effort has been made to verify the following reader tips,* Packet *magazine and Cisco Systems cannot guarantee their accuracy or completeness, or be held responsible for their use. —Editors*

Troubleshooting

TIP *BRI ISDN Interference*

I work in a large BRI ISDN installation, and when we have rains or heavy winds some random connections to the schools are lost. This simple command on our core router clears the buffers of line-induced interference and redials the site.

```
#clear interface BRI <number>
```

We have had such success that I have set up a Web page so our secretarial staff can reset the BRIs themselves with a click of a link.

— Mike Sosnkowski, Fullerton School District, Fullerton, California, USA

TIP *Cable Tracing*

The following Cisco IOS® Software commands can make cable tracing a thing of the past. To find what device is connected to a specific port:

1. Determine the MAC address of the connected host (this only works if the switch port is directly attached to the end station):

```
Switch>show mac-address-table interface fa0/2
Non-static Address Table:
Destination Address  Address Type  VLAN  Destination Port
-------------------- ------------ ----  --------------------
0002.a567.abcd       Dynamic       150  FastEthernet0/2
```

2. Log into the segment's router, and locate the MAC in the ARP cache. If the station is inactive, first initiate a broadcast ping to populate the table.

```
Router>sh arp | include 0002.a567.abcd
Internet  10.1.1.16     158   0002.a567.abcd  ARPA  Vlan150
```

3. If the IP is dynamically assigned using Dynamic Host Control Protocol (DHCP), locate the entry in the server's state table to determine the hostname. Alternatively, in a Microsoft Windows environment, you can use the ping utility to find the name:

```
C:\>ping -a 10.1.1.16
Pinging jdoe [10.1.1.16] with 32 bytes of data:
Reply from 10.1.1.16: bytes=32 time=2ms TTL=127
```

This process can save significant time, because it can be done from any administrator console. To identify a given system's port, just reverse the steps. Instead of supplying a specific port to the **show mac-address-table** command, enter the MAC address.

— Ryan LaTorre, UNIS LUMIN Inc., Oakville, Ontario, Canada

TIP *Cisco Bandwidth on Demand Feature*

I would like to share my personal experience with the Cisco Bandwidth on Demand (BOD) feature, which is triggered only by outgoing traffic levels on ISDN BRI channels. To configure BOD, you must define the load-threshold value under the following interface configuration command:

```
dialer load-threshold (load)
```

Load is the numerical value from 1 to 255, where 1 is the minimum load and 255 is maximum load.The load value is directly related to the bandwidth of the link, which is defined under the **bandwidth** command.

To enable the BOD feature properly, you must change the bandwidth of the interface from the default setting of 1544K to 128K when you are using ISDN BRI. The commands for RouterA are as follows:

```
RouterA(config)#int bri 0
RouterA(config-if)#bandwidth 64000
RouterA(config-if)#dialer load-threshold (load)
```

The same commands are used for int bri 1.

— Haseeb Nasir Ali, Transnet Communications, Karachi, Pakistan

TIP *Cisco Express Forwarding Using Per-Packet Switching*

We have more than four E1 links between the same source and destination pair. Earlier, we had problems with Open Shortest Path First (OSPF) running over these E1 links. The problem was that one of the links was very heavily used—to nearly 100 percent utilization—whereas the others we loaded to 78 to 90 percent.

To solve the problem, we turned on Cisco Express Forwarding globally using the **ip cef** command. Then, on each interface, we explicitly started per-packet switching using the **ip load-sharing per-packet** command.

We thought the CPU load would shoot up to 90 percent, but it was pretty cool at about 25 percent on Cisco 3640 routers at peak loads.

To verify Cisco Express Forwarding is working, use the **sh ip cef summary** or **sh ip cef co** command.

— Dharmesh Shah, Reliance Industrial Infrastructure Ltd, Mumbai, India

Editor's Note: There are consequences using per-packet load sharing, including the potential for out-of-order packets. For more information on load sharing, visit cisco.com/packet/153_4c1.

TIP *Fiber Patch Cable Testing*

Here is an easy way to test if your multi- or single-mode fiber patch cables are damaged: Plug a working Gigabit Interface Converter (GBIC) into a switch or a router. Simultaneously connect both receive connectors to the GBIC, and if the cable is OK, you should get a link up. Do the same with the transmit connectors. Of course, this won't tell you if the cable is degraded, but it is a good way to troubleshoot fiber connection problems in a multipatch link with multiple fiber cables and patch panels involved.

— Georg Pauwen, Network Services, The Netherlands

Tips

Submit a Tip

Help your fellow IT professionals and show off to your peers by submitting your most ingenious technical tip to packet-editors@cisco.com. Who knows, you may see your name in the next issue of *Packet*.

TIP *Flapping Cable Modem*

If you have problems with a flapping cable modem (apart from cable debugs), use the **ping docsis <MAC> verbose** command:

```
uBR-VXR#ping docsis 0020.40db.4e68 verbose
```

Queuing 5 MAC-layer station maintenance intervals, timeout is 25 msec:

```
Reply from 0020.40db.4e68: 2 ms, tadj=-1, padj=0.25, fadj=0
Reply from 0020.40db.4e68: 2 ms, tadj=0, padj=0.25, fadj=0
Reply from 0020.40db.4e68: 2 ms, tadj=0, padj=0.25, fadj=0
Reply from 0020.40db.4e68: 2 ms, tadj=0, padj=0.25, fadj=0
Reply from 0020.40db.4e68: 2 ms, tadj=-1, padj=0.25, fadj=0
```

This will quickly tell you if the cable modem is, for example, adjusting its power level rapidly.

— Lukasz C. Jokeil, Klonex Telecom, Opole, Poland

TIP *Router Load*

Piping process output using the following command comes in handy when troubleshooting router load:

```
(config)# show process cpu | excl 0.00%__0.00%__0.00%
```

However, this will list only CPU active/consuming processes. You can save precious time by also mapping an alias using the following:

```
(config)# alias exec ps show proc cpu | excl 0.00%__0.00%__0.00%
```

— Igor S. Jovanovic, ISC d.o.o., Belgrade, Serbia and Montenegro

TIP *SSH Problems*

If you experience problems connecting your router through Secure Shell (SSH), it might be due to the hostname and router's Domain Name System (DNS). If you generate RSA keys before you configure the hostname and domain, SSH won't work. Remove and regenerate keys.

From config mode, remove keys with this command:

```
crypto key zeroize rsa
```

Generate new keys:

```
crypto key generate rsa
```

Verify keys:

```
sh cry key mypubkey rsa
```

Between other lines, the output should now contain: Key name: *router_hostname.your_domain*. This will enable your SSH connection.

— Vujanovic Sasa, NIL, Belgrade, Serbia and Montenegro

TIP *WAN Link Congestion*

An effective method of troubleshooting a congested WAN link is to use Network-Based Application Recognition (NBAR) applied to the serial interface in question to see what the application usage is in bits per second. The router will need to have Cisco Express Forwarding enabled. Changing the default load interval on the interface to one-minute averages provides more granularity.

```
interface Serial0/0
 ip nbar protocol-discovery
 load-interval 60
```

Define an alias called traffic to pull off the results you need quickly for applications that have been active over the last minute:

```
alias exec traffic show ip nbar protocol-discovery Stats
Byte-R | ex 0
0
```

Here is an example of a site with mostly Microsoft Exchange activity and a single, compressed voice-over-IP call in progress:

```
router#traffic
Serial0/0
          Input              Output
 Protocol   1 minute bit rate (bps)   1 minute bit rate (bps)
 ------------ --------------------
-----------
   exchange    216000             16000
   rtp        16000              16000
   smtp        2000              0
   unknown      1000              22000
   Total       235000             54000
```

— Gary Hills, Legato Systems, Burlington, Canada

*Editor's Note: You can also configure flow caching on the inbound interface with **ip route-cache flow**, and look at the output with **show ip cache flow**. Information will be based on IP protocol/port on a per-flow basis, instead of Layer 7 protocol.*

Maintenance

TIP *Deleting Individual Interfaces*

When a router is fully configured but you want to delete any interface configuration alone, use the **default interface** command:

```
router(config)#default interface serial0
```

This deletes all the configurations under that interface and saves you from using the *no* form of the **reverse** command to disable the functionality.

— N. Sivarajan Thiruvadi, Homebush west, NSW, Australia

TIP *Upgrading IOS*

You can upgrade your IOS image by configuring a router as a Trivial File Transfer Protocol (TFTP) server, and download the IOS image directly from that router. You can configure a router as a TFTP server with the following command:

```
(config)# tftp-server flash:<flash name>
```

— Vivek Mohta and Rajneesh Das, Sify Ltd., Kolkata, India.

All About Tunnels

BY ERIC MATKOVICH

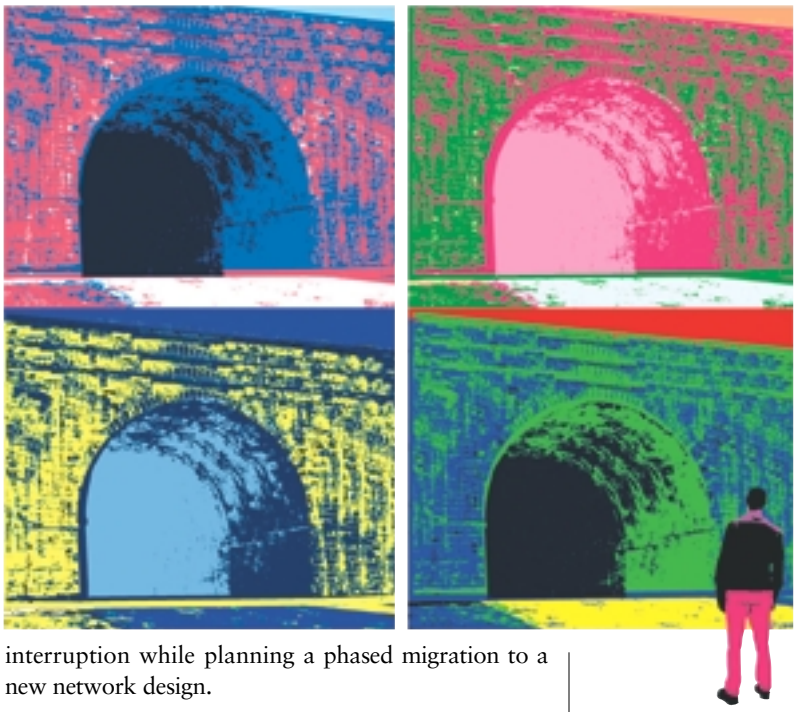

JOHN RITTER

NETWORKS MUST CONTINUALLY ADAPT to new business and capacity requirements. As such, they frequently gain support for additional protocols that deliver valuable new features and efficiencies.

Tunneling, which is supported in router and switch software, involves encapsulating a packet supporting one protocol within another packet that runs the same or a different protocol. We'll explore the reasons why you might want to do this later in this article. However, first you need to understand what the key components of tunneling are. These include the *tunneled protocol*, the *tunneling* or *transport protocol*, and in some cases, a *multiplexing identifier*.

To better understand the components, think of tunneling as driving a car. The tunneled protocol is the passengers in the car and the tunneling protocol is the car itself. What distinguishes our car from other cars are characteristics such as color, year, and model. This is analogous to the multiplexing field and is how you logically separate one set of traffic from another. Different tunneling technologies use different forms of multiplexing, but the concept is the same.

The tunneled protocol (the one to be encapsulated) gains one or more additional headers that can be used to identify different tunnels between a pair of devices and ultimately deliver the payload to a remote peer.

You can apply tunneling to protocols operating at the same layer of the Open Systems Interconnection (OSI) model or at different layers. You can derive a wide range of applications from various tunneling protocols, such as connecting isolated network segments, nondisruptive network renumbering, Layer 2 transport, security, and controlling routing behavior.

Disparate Network Segments

A common problem solved by tunneling is to bridge disparate network segments while implementing a new network addressing scheme to achieve optimal route aggregation. Or, in the case of acquisitions or relocations, an administrator might encounter address overlap, different network protocols, or segmentation policies that need to be preserved. The network administrator can use tunneling to prevent service interruption while planning a phased migration to a new network design.

Latest Protocols

New, emerging tunneling and transport protocols are integral to realizing the benefits of a converged network. Not limited to Layer 3 applications, tunneling can now support Layer 2 transport and will eventually support some Layer 1 circuit emulation services.

Let's take a closer look at some of the new tunneling forms and their primary applications.

Generic Routing Encapsulation

Generic Routing Encapsulation (GRE) takes packets or frames from one network system and places them inside frames from another network system in a peer-to-peer configuration. Enterprises often use GRE to transport legacy Layer 3 protocols over an IP backbone.

A common scenario might be an acquisition in which one enterprise using IP has recently acquired another organization that supports servers or end stations running the Novell IPX protocol. One of the first tasks the network administrator might face is to

reduce the recurring cost of duplicate WAN links between sites in common locations.

Before planning the more comprehensive migration of the new organization at the application layer, the administrator might build a GRE overlay network and replace the long-distance WAN links to terminate at regional locations into the primary enterprise's existing core. IPX applications on the servers that must communicate with remote IPX-enabled devices could use GRE to tunnel IPX across the existing IP core.

In this scenario, GRE has solved a problem by reducing recurring costs, while not inhibiting the organization's productivity.

GRE consists of a packet header with components that allow it to identify data for processing when it arrives at the associated peer. These components include an IP Version 4 (IPv4) tunneling or delivery header; a GRE header with optional fields that include tunnel key, checksum, and sequencing fields; and the payload (or tunneled Layer 3 packet).

Most GRE applications are used to solve problems in the enterprise. As with most technologies, you should weigh the advantages and disadvantages of using GRE. For example, GRE is a point-to-point tunneling protocol, so it would be operationally overwhelming to deploy GRE tunneling in a full-mesh configuration, because the number of tunnels would rapidly spiral out of control. Therefore, GRE is ideal in limited, tactical deployments.

More scalable tunneling tools have recently been developed to handle wide-ranging network transitions—such as merging new IP Version 6 (IPv6) network components onto existing IPv4 networks. Similarly, in service provider environments, multicast virtual private networks (VPNs) displace the need to use full-mesh GRE tunneling to send multicast traffic among Multiprotocol Label Switching (MPLS) VPN sites.

L2TP and L2TPv3

Layer 2 Tunneling Protocol (L2TP) solves some of the remote-access requirements of emerging service providers and enterprise outsourcing. Its primary function is to divorce the Point-to-Point (PPP) endpoint from the Layer 2 circuit termination point. By doing so the subscriber can use local access servers, thus bypassing long-distance toll charges and tunneling PPP frames across the packet-switched network (PSN) for remote termination on an L2TP network server (LNS).

L2TP's typical application is in the "wholesaling" environment. For example, an incumbent local-exchange carrier (ILEC) has built out a dial, ISDN, or xDSL infrastructure in the local area and has excess capacity. The ILEC might want to sell off a portion of its capacity to a nationwide Internet service provider

OPERATIONAL EFFICIENCY: Service providers use tunneling to converge traffic from customer sites with disparate Layer 2 access connections across a single-protocol backbone.

TYPICAL USES OF TUNNELING

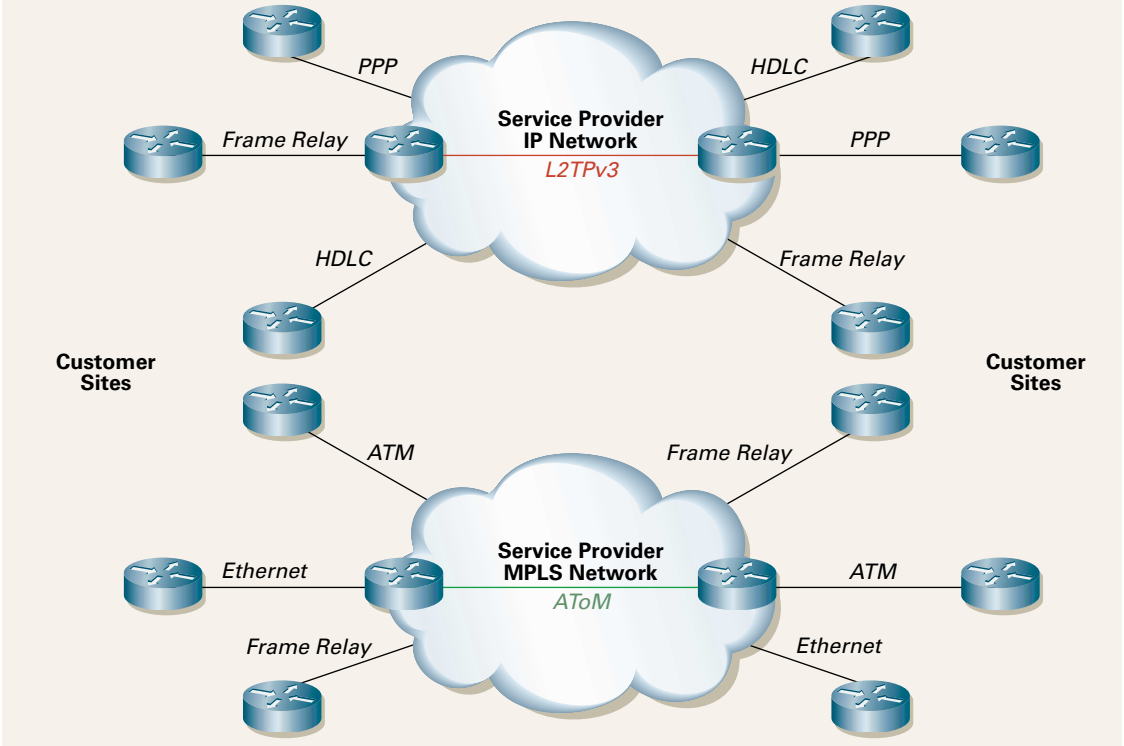

(ISP). The advantage to the ISP is that it doesn't have to build out new infrastructure to service limited user requirements, and this helps the ILEC by generating revenue on what would normally be idle equipment.

L2TP's main components include a reliable control channel that is responsible for session setup, negotiation, and teardown, and a forwarding plane that adds negotiated session IDs and forwards traffic. The Layer 2 circuits terminate in a device called an L2TP Access Concentrator (LAC), and the PPP sessions terminate in an LNS, which authenticates the user and starts the PPP negotiation.

L2TP is primarily used by service providers to deploy VPNs directly to business customers and other ISPs in a wholesale dial scenario, and by enterprises to support remote users. L2TP is an industry standard that combines components of the Microsoft Point-to-Point Tunneling Protocol (PPTP) and the Cisco Layer 2 Forwarding (L2F) protocol, which are proprietary protocols.

L2TPv3, which is currently an Internet Engineering Task Force (IETF) draft standard, expands L2TP to include several new service models. It also introduces support for tunneling alternate Layer 2 Protocol Data Units (PDUs) rather than just PPP. The primary application is for service providers to consolidate multiple Layer 2 networks onto a single high-speed IP network infrastructure for operational savings or, optionally, to allow traditional Layer 3 service providers to offer conventional Layer 2 services without building out an entirely separate core. Layer 2 protocols that L2TPv3 can tunnel include ATM, Ethernet, IEEE 802.1Q, Frame Relay, High-Level Data Link Control (HDLC), and PPP (see figure).

IPSec VPNs

Using tunneling for security is a prime example of how you can use tunnels to deliver enhanced network features and benefits.

IP Security (IPSec) is a sophisticated and complex set of services. Simply put, though, it employs tunneling, in part, to address security vulnerabilities that building a VPN over the public Internet can open up.

IPSec has different methods for securing and encrypting data. Two common encryption methods are IP authentication header (AH) and IP encapsulating security payload (ESP), respectively. Each method can operate in transport or tunnel mode and encrypts all or part of the original IP packet, including the original IP headers. In transport mode, only the IP payload is encrypted and the original IP headers are left intact. In tunnel mode, the entire original IP datagram is encrypted and becomes the payload in a new IP packet.

The AH authenticates the integrity and source of IP packets but does not provide confidentiality through encryption. ESP allows encryption but does not protect the new IP header. So for strong authentication plus confidentiality, you can deploy AH and ESP in tandem in either transport mode or tunnel mode.

MPLS VPNs

MPLS VPNs are "per-customer tunnels" enabled by service providers operating MPLS cores. The primary application is for enterprises that want to outsource all or part of their WAN core to a service provider.

MPLS VPNs allow service providers to separate customer address space on the provider edge (PE) router. When customer traffic enters the service provider's core, a unique VPN label is appended to the incoming IP packet and then switched across a tunnel label-switch path (LSP). By isolating the customer traffic with a VPN label, core routers swap labels based on the outer tunnel label and do not require any knowledge of the inner VPN label or customer payload. LSPs can be determined based on quality-of-service (QoS) requirements, traffic engineering, or minimum bandwidth requirements, and are established through either Label Distribution Protocol (RFC 3036) or extensions to Resource Reservation Protocol, known as RSVP-TE (RFC 3209).

The inner MPLS VPN label, along with VPN routes, is distributed between service provider edge routers using the Multiprotocol Border Gateway Protocol (MP-BGP) extensions, and the label is appended to customer traffic prior to traversing the core network. Then the outer label switches traffic between the core routers until the traffic reaches the destination edge router. The tunnel LSP and the VPN label are removed just before sending the traffic on to the customer's network. Label imposition and disposition are completely transparent to the customer. This service requires that the service provider is participating in the customer routing at the customer edge-to-provider edge demarcation.

AToM

Any Transport over MPLS (AToM), which conforms to the IETF Draft Martini working document, "Layer 2 Transport over MPLS," is analogous to L2TPv3 for use in service provider MPLS networks. It transports Layer 2 frames


ERIC MATKOVICH

ERIC MATKOVICH is a technical marketing engineer in Cisco's Internet Technologies Division. He can be reached at ematkovi@cisco.com. Cisco technical leader Tony Hain, who serves as the technology director of the IPv6 Forum's North American Task Force Steering Committee, and Cisco technical marketing engineer Chiara Regale also contributed to this article.

across an MPLS network. It works similarly to an MPLS VPN in that it appends a label to customer traffic prior to sending it to the destination peer. However, it uses an extended or directed Label Distribution Protocol (LDP) session to negotiate a unique virtual circuit or pseudowire label to identify customer traffic of varying Layer 2 types. The customer uses a native Layer 2 interface (such as Frame Relay, ATM, or Ethernet) to connect to the service provider's network. When the service provider network receives a packet from the customer, it appends a label to the packet and switches it using a tunnel LSP to the destination edge router. The receiving PE router strips off the label and forwards the packet to the customer site in its native Layer 2 PDU format (see figure).

Like L2TPv3, AToM can be used as a network consolidation tool or to allow a service provider to offer new services.

IP in IP

One use of encapsulating IP in IP is when service providers offering an IP VPN service must avoid overlapping internal addressing schemes in multiple customer networks. Encapsulating IP packets generated by the customer within IP packets that conform to the service provider's own addressing scheme segregates the traffic of different customers while it traverses the WAN and keeps it private.

Another application of running IP in IP is to allow multiple versions of the protocol to coexist in different segments of a network. The migration to IPv6 by enterprises and service providers, for example, is expected to significantly increase in the next two years. Several forms of tunneling will aid in IPv4-to-IPv6 network migrations by decoupling the protocol dependencies between network segments.

The most common implementation of tunneling for migration to IPv6 networks is likely to involve wrapping IPv6 packets in IPv4 network address headers, because many network operators will leave the majority of their IPv4 networks in place and gradually add IPv6 segments. These tunnels are called *6to4* tunnels and allow an IPv6 router to automatically tunnel packets across an IPv4 network between two isolated IPv6 sites using a single IPv4 address.

Another type of tunnel, *Teredo*, encapsulates IPv6 in the User Datagram Protocol (UDP) over IPv4. It is primarily used to circumvent the need for IPv4 Network Address Translation (NAT), which can be unwieldy with certain applications.

For a comprehensive discussion of IPv4-to-IPv6 migration tools, including several types of tunneling, refer to the *Packet® Online* Exclusive, "New Tunneling Techniques," at cisco.com/packet/153_5a1.

"802.1Q in Q" for VLAN Extension Services

Some service providers want to offer transparent LAN services that preserve and extend customers' virtual LAN (VLAN) groupings and associated access privileges across a metropolitan-area network (MAN) and, possibly, a WAN. To do this, they can use *Cisco 802.1Q Tunneling* (also called Cisco *802.1Q-in-Q*).

Cisco 802.1Q Tunneling enables service providers to use a single VLAN to securely transport most or all of a single customer's VLANs across their MAN or WAN backbone. In this case, the software adds an extra 802.1Q tag to customer traffic in the switch at the edge of the service provider's network. This tag assigns a unique VLAN ID number to each customer to keep each customer's VLAN traffic segregated and private.

The tag also enables as many as 4094 customer VLANs to be backhauled across a single service provider VLAN through the use of a *tunnel port* that is assigned to each customer site. All of a single customer's VLANs that are configured in the tunnel port on the service provider's WAN edge switch are aggregated and backhauled over a single VLAN.

Service providers do not have to assign a unique VLAN ID number to each individual customer VLAN, which quickly consumes the 4094-ID VLAN space supported by Ethernet's 802.1Q technology. In this way, encapsulating multiple customer 802.1Q VLANs into a single service provider 802.1Q VLAN (thus the name, "Q in Q") affords service providers a scalable approach to offering Ethernet services. To transport not only customers' data traffic but also customers' Layer 2 control traffic (such as Spanning Tree, Cisco Discovery Protocol, and VLAN Trunking Protocol), service providers must configure on 802.1Q tunneling ports Cisco Layer 2 Protocol Tunneling, a separate feature that is available in the same Cisco IOS® Software release.

An Important Tool

Tunneling, in its many forms, has proven itself a valuable tool to enterprises and service providers alike. Advantages include:

- Ability for network segments with protocols in common to communicate with one another over a dissimilar backbone, providing a migration path from legacy to newer protocols
- Assurance of privacy and security in shared networks that support multiple enterprise customers
- Consolidation of multiple protocol types on a common backbone for reduced operations costs
- Ability for service providers to deploy offerings that extend customers' LAN-based capabilities in a scalable way

As networks evolve to solve new business problems, tunneling will likely continue to enable them to adapt in ways that are nondisruptive and protect existing investments. ▲▲

Planting New Spanning Trees

Implementing IEEE 802.1w and 802.1s

BY CHIARA REGALE

A ROBUST, RELIABLE NETWORK NEEDS to transfer traffic efficiently, provide redundancy and recover quickly from faults. In a Layer 2 network, where routing protocols are not available, Spanning Tree Protocol offers redundant connections and eliminates the danger of data traffic loops by building a loop-free logical forwarding topology from a meshed physical topology. The original Spanning Tree Protocol, IEEE 802.1D, typically recovers a link failure within 50 seconds [convergence time = (2 x Forward_Delay) + Max_Age]. While such an outage was acceptable when the protocol was designed, today's mission-critical applications (voice and video, for example) require much faster network convergence.

To speed up network convergence and address scalability limitations related to spanning tree and virtual LAN (VLAN) interaction, the IEEE committee developed two new standards: Rapid Spanning Tree Protocol (RSTP), defined in IEEE 802.1w, and Multiple Spanning Tree Protocol (MST), defined in IEEE 802.1s.

This article highlights key features of 802.1w and 802.1s, interoperability with legacy spanning tree protocols, and suggests some protocol migration guidelines.

Key Definitions

IEEE 802.1D—IEEE proposal that specifies a single, loop-free topology in the whole Layer 2 domain. Since 802.1D has no virtual LAN (VLAN) awareness, the forwarding topology is unique, independent from the VLANs active in the network. This approach is scalable from a computational overhead standpoint, but does not offer load-balancing across redundant links in cases where multiple VLANs are configured.

PVST+—Cisco Spanning Tree that builds a different logical topology for each VLAN. This approach allows load balancing by having different forwarding links per VLAN, but it is more CPU intensive—bridge protocol data units (BPDUs) are generated and processed for each VLAN.

IEEE 802.1w Rapid Spanning Tree Protocol

The IEEE recognized that the convergence characteristics of the original 802.1D Spanning Tree Protocol were inadequate for modern switched networks and applications, and designed a new protocol, 802.1w Rapid Spanning Tree Protocol (RSTP) to address the convergence problems of 802.1D. The characteristics of the IEEE 802.1w RSTP incorporated many of Cisco's value-added spanning tree extensions to the original 802.1D protocol such as Portfast, Uplinkfast, and Backbonefast. (For more information on these Cisco features, go to cisco.com/packet/153_5c1.) The IEEE 802.1w protocol provides rapid recovery from failure of a switch (bridge), switch port (bridge port), or an entire LAN by relying on an active bridge-to-bridge handshake mechanism instead of timers specified by the root bridge as in 802.1D. RSTP changes the way in which the topology is maintained by using spanning tree "hellos" as link-local keepalives. This behavior makes the original 802.1D fwd-delay and max-age timers largely redundant, and are now used primarily for backup purposes to the normal operation of the protocol.

Besides the new concepts outlined in the following section, RSTP introduces new BPDU handling and a new topology change mechanism. Each bridge generates BPDUs every "hello time," even if it doesn't receive any from the root bridge. BPDUs play the role of keep-alive messages between bridges. If a bridge fails to receive BPDUs from a neighbor, it assumes that it lost connectivity to that neighbor, which allows for faster failure detection and convergence.

In RSTP, a topology change occurs only if non-edge ports move to a forwarding state. A loss of connectivity—a port moving to blocking state, for instance—will not trigger a

ROB BRODMAN

CHIARA REGALE

CHIARA REGALE is a technical marketing engineer in the Metro Ethernet Group at Cisco. She focuses on design and validation of solutions in the Metro Ethernet arena. After earning a Master of Science Degree in Telecommunications Engineering from Politecnico of Turin, Italy, Regale joined Cisco's Catalyst® 6000 Spanning-Tree Protocol development team, where she designed and implemented features to improve Spanning-Tree performance. She can be reached at chiarar@cisco.com.

NEW 802.1w PORT ROLES

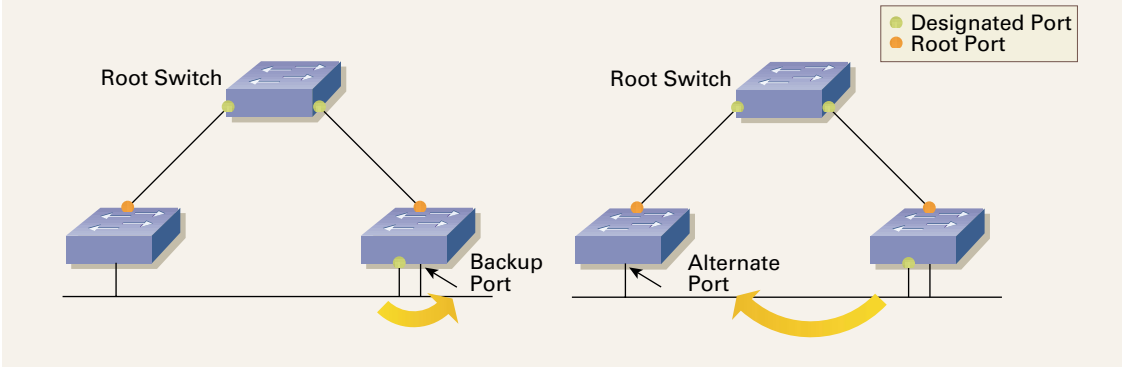

FIGURE 1

topology change as in 802.1D. Topology change notification (TCN) in 802.1w differs from 802.1D to reduce the flooding of data. In 802.1D, the TCN was unicasted to the root bridge and then multicasted to all bridges. The receipt of an 802.1D TCN causes a bridge to fast age all entries in the forwarding table irrespective of whether the bridges forwarding topology was affected. RSTP, by contrast, optimizes this operation by explicitly telling the bridge to flush all entries except those entries that were learned via the port on which the TCN was received. This change in TCN behavior significantly reduces the amount of MAC addresses flushed during a topology change.

Port Roles

RSTP differentiates clearly between port state (whether forwarding or blocking traffic) and port role (whether or not it plays an active role in the topology). Beside the root port and designated port definitions inherited from 802.1D, two new roles (see Figure 1) are specified:

- *Backup Port*—acts as a backup for the path provided by a designated port toward the leaves of the spanning tree. It can only exist where there are two or more connections to a shared LAN segment, or where two ports are connected in loopback by a point-to-point link
- *Alternate Port*—provides an alternate path toward the root bridge to the one provided by the current root port

These new port roles in RSTP enable a rapid transition of an alternate port to forwarding on failure of the root port. The process is explained in the example that follows.

Port States

The state of a port controls the operation of the forwarding and learning process. RSTP defines three states: *discarding*, *learning* and *forwarding*. A root or designated port plays an active role in the topology, while an alternate or backup port does not participate in the active topology. In a steady network, root and designated ports are forwarding, while alternate and backup ports are in discarding state.

Overview of Rapid Convergence

As stated earlier, RSTP aims to transition root ports and designated ports to forwarding, and alternate and backup ports into blocking as quickly as possible. To prevent forwarding loops, RSTP uses an explicit "handshake" between bridges to ensure that the port roles assigned across the network are consistent.

Figure 2 depicts the agreement/proposal handshake that takes place before transitioning a port into forwarding. When a link becomes active, both "p1" and "p2" are designated ports in discarding state.

RSTP AGREEMENT/PROPOSAL MECHANISM

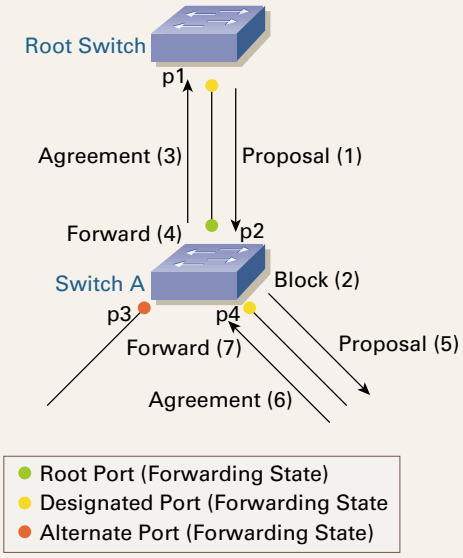

FIGURE 2

In such a scenario, "p1" sends a proposal BPDU toward Switch A. Upon receipt of the new BPDU, Switch A will ascertain that the root switch has a better cost to root. Because the BPDU contains higher root priority, Switch A starts the sync mechanism before moving its new root port "p2" into forwarding. A port is in sync with the root information if it is in blocking state or if it is an edge port (either at the edge of the bridged LAN or connected to end stations).

Port 3 ("p3") already fulfills the above requirements, since it is already blocking. Therefore, no actions are taken on this port. However, "p4" is a designated port, and has to be blocked. Once all the interfaces on Switch A are in sync, "p2" acknowledges the proposal previously received from the root and can safely transition into forwarding. Upon receipt of Switch A's acknowledgement, the root switch will transition "p1" to forwarding immediately. A similar wave of proposal/agreement messages propagates from "p4" toward the leaves of the network.

As this handshake does not rely on timers, it quickly propagates toward the edge of the network, and rapidly restores connectivity following a topology change. If an agreement does not echo a proposal message, the port reverts to 802.1D mode and transitions to forwarding through the legacy listening-learning sequence. It should be noted that 802.1w works only on point-to-point links. In case of a shared media, the 802.1w protocol will revert to 802.1D operation.

INTEROPERABILITY BETWEEN RSTP/MSTP AND PVST

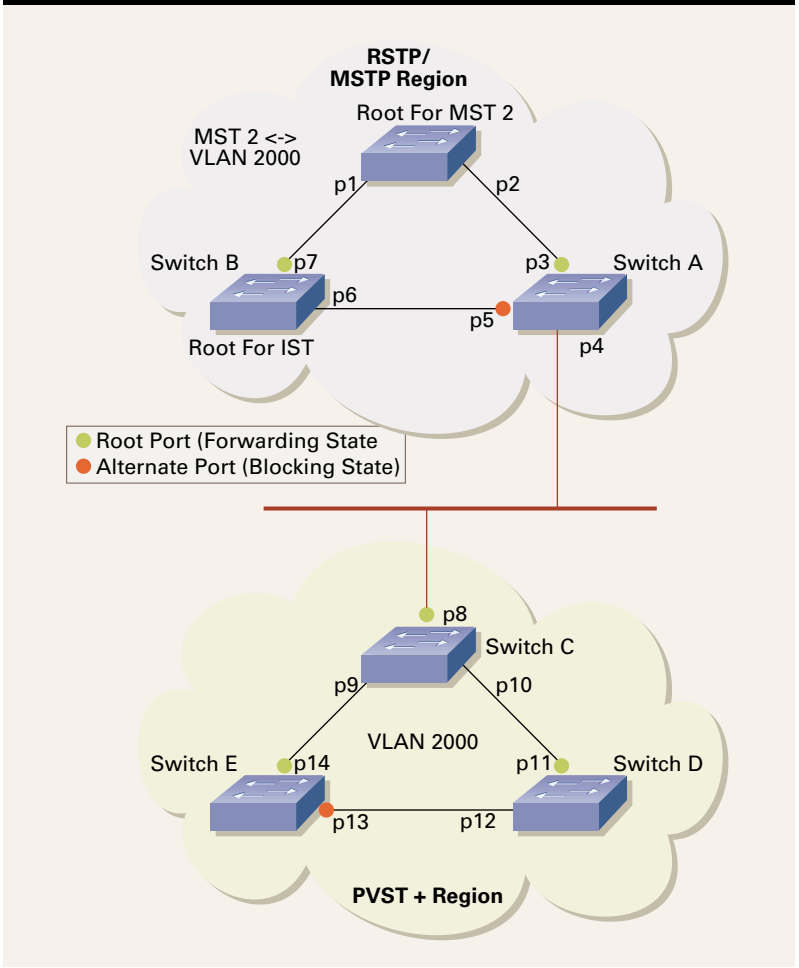

FIGURE 3

Multiple Spanning Tree Protocol

MST, inspired by Cisco MISTP [Multiple Instance Spanning Tree Protocol], improves RSTP scalability by aggregating a number of VLAN-based spanning trees into distinct instances, and by running only one (rapid) spanning tree per instance. To determine the VLAN-instance association, 802.1s introduces the concept of MST regions. Each switch running MST has a single MST configuration identified by an alphanumeric configuration name, a configuration revision number, and a 4096-element table that associates each of the 4096 VLANs potentially supported to a given instance. To be part of a common MST region, a group of switches must share the same configuration attributes. It is important to remember that switches that differ in one or more configuration attributes are considered to be in different regions.

To ensure consistent VLAN-instance mapping, the protocol has to identify the boundaries of the regions. Therefore, the region's characteristics are included in BPDUs. The switches need to know whether they are in the same region as a neighbor, hence a digest of the VLAN-instance mapping table is sent along with the revision number and the name. After a switch receives a BPDU, it extracts the digest and compares it with its own computed digest. To avoid a spanning tree loop, if two switches disagree on any of the parameters received within the BPDU, the port on which the BPDU was received is declared a boundary port.

IEEE 802.1s introduces the concept of IST (Internal Spanning Tree) and MST instances. The IST is an RSTP instance that extends the 802.1D single spanning tree inside the MST region. IST connects all the MST bridges and appears from a boundary port as one virtual bridge that encompasses the whole bridge's domain. The MST Instances (MSTI) are RSTP instances that only exist inside a region. They run RSTP by default, without any additional configuration. Unlike the IST, MSTIs neither interact nor send BPDUs outside a region. MST interoperates with legacy and PVST+ switches. Cisco implementation defines 16 instances: one IST (instance 0) and 15 MSTIs, while 802.1s allows for one IST and 63 MSTIs.

Migration strategies and sample configurations highlighting the forwarding topology for several switches can be found at *Packet Online* at cisco.com/packet/153_5c2.

Interoperability with Legacy Spanning Tree

RSTP and MSTP are both able to interoperate with legacy spanning tree protocols. However, 802.1w's fast convergence benefits are lost when interacting with legacy bridges.

To retain backward compatibility with 802.1D-based bridges, IEEE 802.1s bridges listen for 802.1D format BPDUs on their ports. If an 802.1D BPDU is received, the port uses standard 802.1D behavior to ensure compatibility. For instance, in Figure 3, as soon as "p4" on Switch A detects PVST+ BPDUs for at least twice the "hello time," it starts sending PVST+ BPDUs. It should be noted that if the PVST+ bridge is removed from the network, Switch A won't be able to detect the topology change and a manual intervention would be required to restart the protocol migration.

Figure 3 represents the forwarding topology for VLAN 2000, which is mapped to MST #2 in the RSTP/MSTP region. The root switch for both the IST and MST # 2 resides within the RSTP/MSTP region. MSTI BPDUs are not sent out the boundary port "p4", only IST BPDUs are. The MST region simulates a PVST+ neighbor by replicating the ISTP BPDU on all the VLANs active in the PVST+ domain. Consequently, the PVST+ domain receives the BPDUs sent on the IST and elects Switch B as root switch for VLAN 2000 (Note that Switch B is the root for the IST.)

If a topology change occurs in the PVST+ domain, the corresponding topology change notification (TCN) BPDU generated in the legacy cloud will be handled within the MST region by the IST, without affecting the MST forwarding topology. In order to avoid misconfigurations that may lead to loops, it is highly recommended to configure the root switch for the PVST+ instances (i.e., IST root) in the MST domain. ▲▲

FURTHER READING

For more information on spanning tree, RSTP, and MSTP, check out the following white papers:

- **Configuring Spanning Tree: cisco.com/packet/153_5c3**
- **Understanding Rapid Spanning Tree Protocol 802.1w: cisco.com/packet/153_5c4**
- **Understanding Multiple Spanning Tree Protocol 802.1s: cisco.com/packet/153_5c5**

Identification, Please

"Hardened" Cisco AAA appliance joins identity-based networking to strengthen security.

Recent technology strides by Cisco give organizations the tools to reinforce network security by authenticating users based not only on the MAC and IP addresses of their client devices but also on their personal identity. Verifying the identity of users before granting them corporate network access has become possible under an initiative called Cisco Identity-Based Networking Services (IBNS), an IEEE 802.1X-based technology framework. IBNS now spans a broad range of Cisco LAN infrastructure products to identify users attempting to gain access to the wired or wireless LAN (WLAN). IBNS is supported across the family of Cisco Catalyst® switches, in Cisco Aironet® 1100 and 1200 series access points, and in the Cisco Secure Access Control Server (ACS), which functions as a RADIUS.

Cisco Secure ACS 3.2 recently became available in a one-rack-unit "hardened" appliance form factor called the Cisco Secure ACS Solution Engine. "The appliance is another option to make it much more difficult to penetrate server operating systems and third-party applications that might have known security vulnerabilities," explains Michel Chahine, Cisco Secure ACS product manager in the Secure Managed Networks Business Unit at Cisco.

What's Driving Identity-Based Networking?

To maximize productivity, many organizations have adopted a policy to enable active physical ports through the network and implemented WLANs. These open environments, combined with the ease of obtaining an IP address via autoconfiguration with the Dynamic Host Control Protocol (DHCP), have rendered many corporate networks more susceptible to unauthorized access. While these technologies open up significant productivity opportunities to organizations by enabling greater mobility, they also present significant security considerations. In an unsecured environment, it is considerably easier for unauthorized outsiders to gain network access to eavesdrop on sessions, launch denial-of-service (DoS) attacks, or hijack an IP address and hack into network-based data resources.

These are only a few of the reasons that the security landscape is evolving to control network access based on a user's identity, says Ken Hook, a Cisco Catalyst 6500 Series product manager who helped develop IBNS. With this initiative, Cisco has extended the industry-standard 802.1X authentication framework across the Catalyst 6500, 4500, 3550, and 2950 series switches and Cisco Aironet series access points. Depending on whether the user is attempting to gain access via a wireless or wired connection, any of these devices can function as an *authenticator* in the 802.1X model, a role that involves passing messages and authentication challenges between clients and authentication servers. The Cisco Secure ACS 3.2 (both in the appliance and in a Microsoft Windows server-based form factor), performs the role of *authentication server* in the 802.1X framework, and can interface with many data stores for user identity lookups (for example, Microsoft Active Directory). The basic premise of Cisco IBNS is to *keep the outsiders out* through authentication and authorization and *keep the insiders honest* by introducing accountability into the network as a deterrent, says Ian Foo, a technical marketing engineer for Cisco's Enterprise Solutions Design team, the solutions architect for Cisco IBNS.

The IEEE standard protocol for port-based network access control, 802.1X acts as a link layer transport for messages carried within a higher-level authentication protocol—for example, EAP, the Extensible Authentication Protocol (RFC 2284)—using the appropriate frame format of the Layer 2 LAN being accessed, such as Ethernet (802.3) networks or WLANs (802.11). If desired, 802.1X could be used on other LAN media types. Cisco Secure ACS also supports EAP, providing centralized intelligence and control for the user authentication, authorization, and accounting (AAA) process used to support access control at the network perimeter.

Cisco Secure ACS generally resides in a data center and serves to identify and make intelligent policy-based decisions regarding users attempting to access the network. Identification of a client based on a username/password combination, a device's Layer 2 MAC address, or a Layer 3 IP address are all possible. In a typical configuration, the client device requesting network connectivity (known as a *supplicant* in 802.1X terminology) communicates via the intermediary authenticator device, usually the network access device residing at the network edge.

The authenticator passes the client's credentials—such as username and password or Public Key Infrastructure (PKI) certificates—to the Cisco Secure ACS software. The ACS, playing the role of authenti-

Continued on page 29

For more information on Cisco IBNS and the Cisco Secure ACS Solution Engine, see cisco.com/packet/153_5b2 and cisco.com/packet/153_5b3, respectively.

Identity-Based Networking, Continued from page 27

cation server within the 802.1X framework, then iterates through a series of communication exchanges with the client, via the authenticator. If the process is successful, and the client's identity and credentials are valid, the ACS operates with the understanding that the user is indeed the person represented and issues the necessary authorizations and policy instructions to the edge device to apply to that given client's session. (For a more detailed summary of the steps involved in this process, see "The Basics of 802.1X Authentication" at *Packet Online*, cisco.com/packet/153_5b1.) Depending on the authentication method used, the ACS can also verify that the portion of the network accessed is the one represented and allowed to that particular client.

To benefit from this functionality, all devices in the system must support the IEEE 802.1X protocol and an authentication algorithm, typically a version of EAP. The various EAP methods include EAP-Transport Layer Security (TLS), which leverages the security of PKI crypto; Protected EAP (PEAP); EAP-MD5; Cisco LEAP; EAP-Generic Token Card (GTC); and EAP-Microsoft Challenge-Handshake Authentication Protocol (MSCHAP). Using the appropriate EAP method, one-time passwords, hashed challenge-response mechanisms, or public key certificates with RSA encryption may be used to authenticate a client securely across the network. Cisco IBNS also supports PKIs for stronger authentication through the use of public key encryption with x.509 v3 certificates. This feature allows IBNS and ACS to integrate into the strongest enterprise security models that may be seeking to, or that might already be, leveraging PKI.

Why Use an Appliance?

The Cisco Secure ACS Solution Engine runs a 2.66-GHz processor in a single-rack-unit form factor, and provides the same basic features and functions as Cisco Secure ACS for Windows in a dedicated, application-specific appliance package. Its hardened security posture, plug-and-play setup, and reliability benefits are aimed at improving total cost of ownership with reduced setup and maintenance. Chahine recommends the Cisco Secure ACS Solution Engine as Cisco's latest resilient, reliable AAA server offering. In this appliance form factor, the operating system is configured to automatically reboot on system crash, and the serial console service is configured to automatically restart if it fails. And Cisco Secure ACS software implements the monitor that will restart Cisco Secure ACS services if they fail. The appliance also provides a remote administrator command line interface (CLI) that supports serial line and Telnet connections, and the Cisco Secure ACS service can be reimaged, reloaded, upgraded, and rebooted remotely from the CLI.

What's more, the appliance runs its own preinstalled operating system (OS) so that IT staffs don't have to worry about installing software and maintaining system versions as they change. The OS is a closed system with all services subject to security attacks stripped down or removed, and is one reason that the hardened appliance is more difficult for hackers to infiltrate. For example, NetBIOS—the widely deployed transport-layer program that allows applications on different LAN-attached computers to communicate—"is known for its vulnerabilities and has been removed from this platform," says Chahine. The appliance also closes down all TCP ports not in use by Cisco Secure ACS to prevent hackers from using them as a network entry point.

802.1X Extensions in Cisco IBNS

On top of the IEEE 802.1X standard framework, Cisco has layered on capabilities to address additional enterprise networking needs. For example, in addition to verifying a user's identity and then activating the port to which the user has connected, Catalyst switches will dynamically provision access to the appropriate virtual LAN (VLAN) assigned to that user. "Internally, this could mean that a user is automatically associated with the finance or engineering VLAN. Also, consultants or others on site can be auto-provisioned to the so-called 'guest VLAN,' or a more selectively restricted VLAN, which would provide them with Internet access and, depending on the nature of their relationship with the enterprise, access to select other corporate resources," explains Hook.

Additionally, the rise in the use of IP telephony has resulted in many organizations building *voice VLANs (VVLANs)*, primarily so that voice traffic is always mapped to the highest-priority 802.1p queue on the LAN to avoid latency and service degradation. Cisco Catalyst switches running the "802.1X with VVLAN" feature can autodetect the presence of Cisco IP phones and map them to the VVLAN. In addition, when a workstation is connected to the switch ports on the Cisco IP phone, they are required to authenticate to gain access to the network.

In the case of a daisy-chained phone and PC connection to a Catalyst switch, the switch will detect the PC and require authentication to gain access to the network. This capability enables both IP telephony and IBNS to operate simultaneously on the same Catalyst switch port, says Hook.

◆ ◆ ◆

By combining access control and user profiles, Cisco IBNS not only offers greater flexibility and mobility to users—resulting in productivity gains—but greatly enhances security for an organization's network connectivity, services, and applications. Additionally, notes Foo, the centralized control architectural nature of Cisco Secure ACS allows for simplified, yet redundant, management of client accounts, information, and policy to minimize operational overhead. ▲▲

WIRELESS TAKES OFF

BY JOANIE WEXLER

WIRELESS NETWORKING ACTIVITY IS TEARING through business and consumer markets like a tornado. Wireless LANs (WLANs), in particular, are finally spilling out of their vertical-market niches—where their payback has long been impressive and easily measured—and into mainstream business and consumer applications.

A confluence of factors drives the masses toward mobility.

Infonetics Research, for example, reports that in 2002, about 38 percent of large North American companies and about 20 percent of organizations in the UK and Germany had adopted WLANs. The research firm expects these penetration numbers to more than double over the next four years.

What's spurring the deployment uptake?

For one thing, businesses that have installed WLANs for use by knowledge workers are reporting as much as 90 minutes of time-savings per employee per day. And some enterprises with vertical applications are realizing they can leverage their existing infrastructures for additional uses. Hotels currently offer WLAN connectivity as an amenity to woo guests, for example, while using the same infrastructure for internal applications, such as automated guest check-in at the curb.

Retailers that traditionally use WLANs for inventory applications are Webcasting product demonstrations and other live events in stores using wireless connectivity. Some hospitals talk of piggybacking on the WLAN infrastructure used by mobile doctors and nurses to provide Internet access services to patients and visitors.

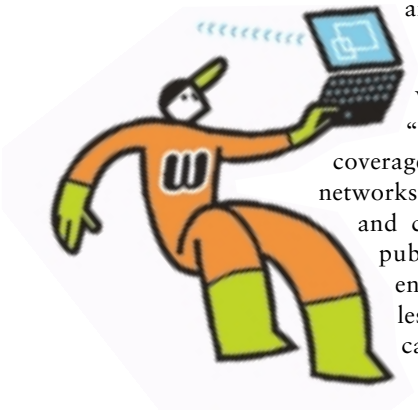

Meanwhile, cellular carriers around the world are deploying WLAN technology (also called "Wi-Fi") to inexpensively fill in some coverage gaps with the multimegabit-speed networks in public venues such as airports and convention centers. The resulting public-access WLAN services help enterprises further justify their wireless investments, because employees can use their WLAN client connections to conduct business in places other than just their corporate offices. Other market factors, too, are falling into place. A few cases in point:

- Baseline Wi-Fi standards have matured and have been certified for interoperability.
- Emerging centralized configuration and management capabilities address grand-scale implementations.
- Home users are pushing their employers toward wireless.
- About 40 percent of enterprise-class laptop computers now ship with WLAN capabilities built right in, according to the Wi-Fi Alliance, a multivendor industry consortium.
- Internetwork roaming technology is simplifying network access for end users and driving acceptance.

The list goes on. But these are among the reasons Synergy Research Group expects the worldwide enterprise Wi-Fi market to grow to US$2 billion by the end of 2003. The firm also anticipates that the market will pick up steam at a compound annual growth rate (CAGR) of 14 percent and will surpass the US$3 billion mark by year-end 2006.

Let's look at the status of the WLAN landscape and examine some issues that the latest technology trends raise for user deployments.

Standards Prompt Architectural Considerations

The 802.11b, a, and g standards have been approved, so enterprises can soon build their WLAN infrastructures based on a mix of technologies to optimize the coverage and capacity of their deployments.

The 802.11b standard is already well established in the marketplace. A variety of system vendors and chipmakers support it, and the technology has made substantial inroads in enterprise deployments, as evidenced by the Infonetics research.

ILLUSTRATION: LEO ESPINOSA
PHOTOGRAPH: © KEVIN TWOMEY

802.11a interoperability testing is now in progress by the Wi-Fi Alliance. Note that in Europe, extensions to 802.11a are described in an emerging draft standard called 802.11h.

In May 2003, the 802.11g Task Group ratified the 802.11g standard, and in June, the IEEE Standard Board issued its final approval. Wi-Fi Alliance testing is slated to begin later this year.

Because of the standards progress, businesses are now evaluating running 802.11a and 802.11g WLANs, which each run at speeds of 54 Mbit/s, alongside 11-Mbit/s 802.11b networks. The availability of all these WLAN types opens up new opportunities, as well as architectural and management considerations.

For example, 802.11g is backward compatible with 802.11b. This is extremely desirable from a migration perspective, because G and B clients and access points can be mixed and matched. Still, if you deploy only B and G access points in a given coverage area (and no A), you remain limited to three nonoverlapping channels. This could lead to interference.

Also, industry-standard G technology works such that once a B client associates with a G-based access point, G network throughput is impacted, even if the B client is not transmitting any packets. This is because special headers, which create network overhead, must be appended to G packets to enable B packets to detect and avoid G traffic.

Fortunately, tri-mode chipsets supporting B, A, and G recently began shipping, which means that decisions about which types of access points and clients to deploy will soon no longer be interdependent. In other words, organizations could run any mix of B, A, and G access points that among them have as many as 27 nonoverlapping channels—depending on imminent decisions about global cooperation on the use of the 5-GHz spectrum—for easily configuring networks that avoid interference. Tri-mode clients supporting B, A, and G could simply associate with the best network connection available to them in a given location.

Eventually, client devices will also support cellular protocols. As a result, if no WLAN connection was available, but a device was in coverage range of a 3G service, a connection could be secured using the lower-speed (but more ubiquitous) WAN service.

While these client capabilities are only just emerging, internetwork roaming is simplifying the user's ability to connect—and stay connected—without fiddling with settings or swapping out PC cards to access different types of networks. User acceptance is often driven by ease of use, and multinetwork products will go a long way toward fulfilling that requirement.

Developments with 802.11a/5-GHz Spectrum

The biggest global issue surrounding 802.11a networks is mismatched spectrum usage in different parts of the world. At the time this issue of *Packet*® went to press, the World Radiocommunication Conference (WRC) was taking place in Geneva, Switzerland. The WRC aims to coordinate spectrum-allocation and frequency-usage regulation on a global basis.

On this year's WRC agenda was the harmonization of the 5-GHz spectrum worldwide. One goal was to define a common set of global rules for running 802.11a WLANs in the 5-GHz frequency range. Let's take a brief look at the issues on the WRC agenda at press time.

Inconsistent Spectrum Use and Power Limits. Because the 5-GHz RF band is not used consistently in countries around the world, it is difficult for users and vendors to buy and manufacture 802.11a products that can be used wherever multinational customers might have facilities. For example, the US and Canada allow Wi-Fi to operate both in the 5.15-to-5.35-GHz band and the 5.725-to-5.825-GHz band. Europe allows WLANs to operate in these bands, too, but also in the 5.47-to-5.725-GHz band. Japan allows usage in a different spectrum.

Having common frequencies and spectrum-usage restrictions would allow customers to purchase one version of a product for use in different countries. Vendors would no longer have to build products to separate sets of specifications and could more easily achieve economies of scale.

Additional Spectrum for the US? The Wi-Fi Alliance petitioned the US Federal Communications Commission (FCC) in January 2002 to open the mid-band of 5.47 to 5.725 GHz for Wi-Fi use in the US to match Europe's 802.11a frequency allocation. The additional spectrum would provide an extra 255 MHz of capacity for 802.11a products in the US, enabling more WLAN channels and reducing the potential for interference. Initially, the US military was concerned about potential interference issues. However, industry and government arrived at a compromise to protect military radars and allow Wi-Fi use in the middle band. So use of this spectrum in the US is on the table at the WRC.

In addition, because of its own concerns about interference with military radar, Europe has imposed power limits—known as transmit power control, or TPC—and dynamic frequency selection (DFS) requirements on the use of that spectrum. DFS thresholds describe the "loudness" of a signal beyond which a WLAN must move to another channel so as not to interfere with military radar. The European Telecommunications Standards Institute (ETSI) requires that 802.11a LAN products built for use in Europe conform to these specifications.

The ETSI requirements are described in the emerging extension to the 802.11 standard called 802.11h, mentioned earlier. 802.11h is in the final stages of approval by the IEEE.

Rx for Security Concerns

Recent security advances such as Wi-Fi Protected Access (WPA) product compliance and interoperability certification mean that dynamic encryption keys, 802.1X authentication, and other important security capabilities have been added to Wi-Fi products for standards-based security.

WPA is a subset of the forthcoming IEEE 802.11i standard for stronger security

JOANIE WEXLER

JOANIE WEXLER is an independent analyst and editor in the computer networking industry and a frequent contributor to *Packet* magazine. Based in California's Silicon Valley, she has covered the networking and IT industries for 14 years, in her current freelance capacity and in several senior editorial roles on staff at International Data Group weekly publications. Wexler currently authors *Network World*'s twice-weekly "Wireless in the Enterprise" electronic newsletters. She can be reached at joanie@jwexler.com.

than was offered in 802.11's initial Wired Equivalent Privacy (WEP) algorithm, which relies on static encryption keys. While vendors took it upon themselves to plug the security holes with robust solutions of their own [such as the Cisco Wireless Security Suite], the additional availability of standards-based technology allows interoperability among different vendors' clients and access points.

In addition, many vendors are adding automated capabilities to detect security policy violations, to log and track user connections, and to authenticate users based on their personal identity, rather than only on a MAC or IP address (see "Identification, Please," page 27). This action is important, given that small devices can easily be lost or stolen, opening up potential access opportunities to outsiders.

Also, many vendors are adding air-monitoring capabilities to their centralized wireless management systems. The monitors detect unauthorized access points and alert administrators about them. Well-meaning employees often bring these "rogue" access points from home, and simply plug them into Ethernet switch ports for their own convenience, unwittingly introducing network vulnerabilities into their organizations.

In fact, home use is where Wi-Fi has gained its momentum, according to Gemma Paulo, a senior analyst, voice and data communications, at In-Stat/MDR in Scottsdale, Arizona. It all began when home users turned to WLANs for a quick fix to the issue of sharing broadband Internet connections among family members.

"Once accustomed to this convenience, home users began asking for wireless capabilities at work—or installing their own unauthorized access points themselves," Paulo notes.

As a result, some enterprises are installing WLANs as a defensive measure so that they can secure and control them. The issues with rogue access points are that 1) they cause interference if not modulated properly by an RF technician; and 2) they can become a potential point of entry for an intruder.

If only a few access points are deployed in a single office, store floor, or warehouse, walking around with a handheld scanner to detect rogue access points is not a particularly burdensome task. However, detecting them among hundreds or thousands of infrastructure devices becomes a much more cumbersome undertaking without a high level of automation.

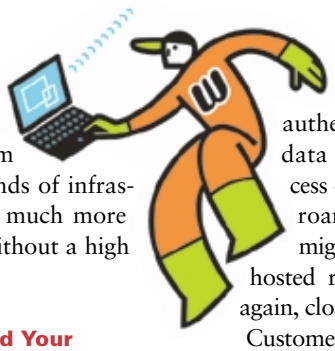

Getting Your Arms Around Your Air Space

In the past, centralized management tools for large-scale deployments were rudimentary but are now growing much more sophisticated. In addition to the air-monitoring capabilities described in the previous section, other capabilities that enable large-scale deployment and management include the following:

- Site-survey tools that automate the process of determining where best to install access points for large-scale coverage and adequate per-user capacity
- Mass configuration of access points and mass firmware upgrades
- Network self-healing; the ability of access points to automatically increase their transmit power in an attempt to fill in any resulting coverage gap if one access point should quit working

There are many architectural approaches to deploying and managing intelligent network features for WLANs. A current industry debate is over how to distribute intelligent network features and functions across access points, wiring-closet devices, and data center equipment.

The "right" choice will likely vary depending on the size of an organization and how geographically distributed it is. For example, highly distributed organizations might require some features that normally would be hosted by a centralized device, such as a RADIUS server, to be bundled into local access points in case the WAN connection to the central device is lost.

One general guideline is to have intelligence located on devices that make the most sense based on your applications. For example, quality of service (QoS) and packet buffering might reside on access points, closest to the user, because the additional latency incurred by packets having to travel deeper into the network to determine their priority could defeat the purpose of having QoS in the first place. Similarly, while it might be desirable to host initial user authentication services in the data center, the subsequent process of reauthentication as a user roams among access points might be more efficient when hosted right on the access points, again, closer to the user.

Customers will begin to get a feel for what works best for them as they gain experience with emerging products.

◆ ◆ ◆

Over the next 12 to 18 months, there will likely be a "settling" of approaches to WLAN architectures as enterprises gain experience with emerging products and evaluate the merits of new automated management capabilities.

The cellular carriers are likely to continue investing in Wi-Fi for supplementary coverage and, as laptops and other mobile devices ship with embedded Wi-Fi connections, users will find themselves using wireless connections almost in spite of themselves.

For the time being, 802.11b—the proven workhorse of the WLAN industry—will remain the primary network in place within enterprises so long as businesses continue to view Wi-Fi as an extension to their wired networks. And voice is poised to join those networks.

As more users hop on existing WLANs with additional applications, the need to add higher-speed infrastructures with more channels using 802.11a, g, and h technologies will emerge. What will enable enterprises to build these larger, more complex WLAN infrastructures are sophisticated configuration, management, and security capabilities that automate network-wide control of the air space. ▲▲

FURTHER READING

- **"Wireless in the Enterprise" newsletter archives (Network World Fusion): www.nwfusion.com/newsletters/wireless**
- **"Wi-Fi Fanfare," *Packet Online* Exclusive: cisco.com/packet/153_6a1**

REINING THE RF

Managing and Securing a Corporate Wireless Network

FOR MANY ENTERPRISES, MANAGING A wireless network has felt like riding a runaway horse without a bridle—until now. Cisco Systems reins in its 802.11 wireless LAN (WLAN) solutions with the Cisco Structured Wireless-Aware Network, a powerful framework for deploying, managing and operating scalable, business-class WLANs.

BY GAIL MEREDITH

Many organizations have been reluctant to deploy extensive WLANs because management tools were weak or nonexistent, and WLAN management was labor intensive. "The capital cost of a wireless network is low," says Pejman Roshan, product manager in the Wireless Networking Business Unit at Cisco. "But without the proper management framework, operational costs of large WLANs, in terms of man-hours required, can be high. What's more, there is a perception among network managers that wireless networking is unreliable because managing it requires high human interaction."

However, if you don't deploy a WLAN, your users probably will. It's as easy as plugging an access point into an Ethernet switch port. Unauthorized or "rogue" access points present a significant security risk, because default settings are insecure, and the access point is unknown. Anyone with an 802.11 client device could log into your network, even from outside the building.

"Rogue APs [access points] are far more common than most network managers realize," says Bruce McMurdo, technical marketing engineer in Enterprise Solutions Engineering at Cisco. While employees who install rogue access points are simply ignorant of the security breach, malicious intruders can take advantage of the leak or install their own device, then hide their exploits.

WLAN Management Requirements

"Enterprises are intrigued by wireless networking," says Roshan, "but all aspects must be ready. As with wired LANs, network managers need a mature solution to troubleshoot and manage a wireless LAN." This solution must provide fault, configuration, accounting, performance, and security (FCAPS) management functionality, and security controls that include the ability to detect and remove rogue access points.

With the Cisco Structured Wireless-Aware Network, the company delivers a comprehensive, standards-based management solution that leverages the customer's existing Cisco infrastructure. It allows enterprises to manage operational costs because it requires no additional network infrastructure, headcount, or specialized training in radio frequencies (RFs). This new functionality in the network is scalable and cost-effective because it resides on Cisco access points today, and Cisco routers and switches starting in 2004, eliminating the need for additional infrastructure—such as dedicated wireless switches—that sit in the data path. "There is no 'bump' in the wire," says Roshan.

The Cisco Structured Wireless-Aware Network solution includes the following components:

- Cisco IOS® Software features
- Cisco Aironet® Series wireless LAN access points
- Cisco Aironet Series wireless LAN client adapters
- CiscoWorks Wireless LAN Solution Engine (WLSE) version 2.x
- Cisco Wireless Security Suite
- Cisco Compatible client adapters and mobile devices from other vendors
- Cisco Secure Access Control Server (ACS)
- Cisco Catalyst® 3750, 4500, and 6500 series switches and Cisco 2600XM and 3700 Series routers (starting in 2004)

The only dedicated hardware in the solution is the CiscoWorks WLSE version 2.x. But unlike a dedicated wireless switch, the CiscoWorks WLSE is not an infrastructure device in the data path. Each CiscoWorks WLSE 2.x appliance can manage up to 2500 access points from the network operations center. It can support larger networks by adding management domains, and integrates with other CiscoWorks tools, such as CiscoWorks LAN Management Solution. Its extensible markup language (XML) interface can communicate with HP OpenView, and Tivoli management systems.

Assisted Site Surveys

Proper management of the WLAN begins with installation. Precise site surveys are vital to successful deployment but are often conducted

by consultants who use a mix of unsophisticated tools, experience, and intuition for varying results. The Cisco Structured Wireless-Aware Network automates and assists with site surveys, making it possible to conduct them by in-house IT professionals who may not be well versed in RF propagation and measurement. The Assisted Site Survey Tool, part of the CiscoWorks WLSE version 2.5, lets network managers complete site surveys in five simple steps:

1. Import floor plan of location or draw an approximate diagram within the tool.
2. Add initial access point locations to the diagram to estimate adequate RF coverage.
3. Install Cisco Aironet access points at locations indicated in the diagram.
4. Set access points to "AP Scan Mode;" all units assume the same channel and transmit at maximum power; they detect each other's presence and automatically adjust transmit power, frequency, and other settings.
5. Access point RF settings are fine-tuned using "Client Walkabout" mode; an individual with a wireless client device walks the facility area, including the perimeter; access points detect client RF measurements and automatically adjust their signals further to provide optimal RF coverage.

Wireless Domain Services

One feature that makes the Cisco Structured Wireless-Aware Network so scalable is a component of its distributed management framework called Wireless Domain Services (WDS). WDS is a collection of Cisco IOS Software features that enhance WLAN client mobility and simplify WLAN deployment and management. It is deployed in Cisco Aironet access points, and will be deployed in Cisco routers and Catalyst switches in 2004. The current WDS feature set includes fast secure roaming and IEEE 802.1X local authentication.

Fast secure roaming allows authenticated client devices to roam securely from one access point to another without perceptible delay in reassociation. Secure handoff services from the access points enable <150 ms roaming within a subnet. This feature supports latency-sensitive applications such as enterprise resource planning (ERP), Citrix-based solutions, and wireless voice over IP (VoIP), without dropping connections during roaming. (Client devices must support the Cisco Centralized Key Management protocol.)

Fast secure roaming works by designating a Cisco infrastructure device such as a router, switch, or access point as a WDS device. Users initially authenticate with the central RADIUS server, and the local WDS handles roaming and reauthentication by holding a "master key" securely transferred from the central RADIUS server (see figure).

The *IEEE 802.1X local authentication service* configures access points to act as local RADIUS servers for up to 50 local accounts when a central RADIUS server is unavailable. "The concept is similar to Cisco Survivable Remote Site Telephony [SRST] for IP phones," says Roshan. "Local Authentication Service enables users in branch offices or other remote sites to be authenticated locally so they can continue working if their WAN access is temporarily cut off."

"Until now, there hasn't been a very secure scheme for voice using 802.1X authentication," says Sri Sundaralingam, a technical marketing engineering in Cisco's Wireless Networking Business Unit. "Now, you can enable fast secure roaming with 802.1X authentication to dynamically generate encryption keys for voice. It's a huge break for VoIP handhelds."

Detecting Rogue Access Points

The most efficient way to monitor enterprise airwaves is to use the infrastructure itself, not dedicated monitoring devices or individuals walking around with scanners. By fourth quarter 2003, the Cisco Structured Wireless-Aware Network will be able to detect rogue access points using a robust method that employs both Aironet access points and Cisco and Cisco Compatible client devices. Since rogue access points are most likely to exist where legitimate access points do not, using access points alone to find rogues is not as effective as using them in conjunction with client devices.

Client-assisted scanning increases the RF coverage area over access points alone by 10 to 20 times. Clients sample applicable airwaves every few seconds on a channel-by-channel basis for both beacons consistent with ad-hoc rogues and RF signatures that indicate the presence of camouflaged rogues. The local WDS device collects data about these samples and uploads it to the CiscoWorks WLSE 2.5 appliance.

The CiscoWorks WLSE software compares sampled data with valid MAC addresses for known access points, and reports potential rogue devices (or other 2.4-GHz network interference) and its approximate location (calculated by triangulation) to the IT

LEO ESPINOSA

FAST SECURE ROAMING

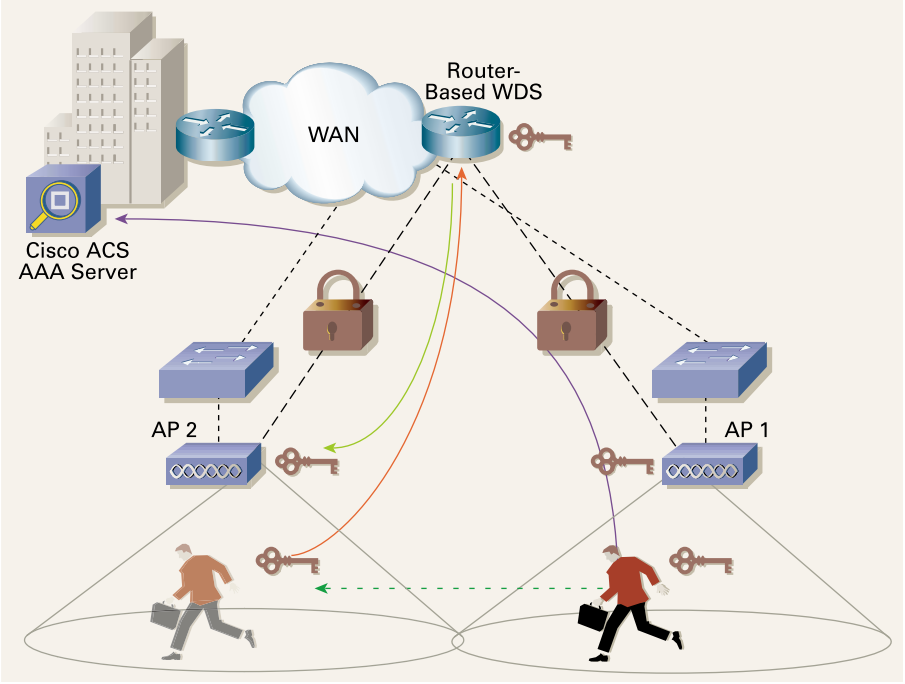

THE HANDOFF: Using WDS, authenticated client devices roam securely from one access point to another without perceptible delay using the access point's secure handoff services.

staff. This appliance provides a much-needed, useful tool that proactively alerts the IT staff of potential security breaches.

Cisco Wireless Security Suite

"Many enterprises have not deployed wireless LANs because they think they're insecure," says Brian Cox, technical marketing engineer in Enterprise Solutions Engineering at Cisco. "They heard about the weaknesses in WEP [Wired Equivalent Privacy] a few years back and haven't kept up to date."

Cisco addressed the issues of weak encryption and poor user authentication in the WEP standard two years ago (see "How to Build a Secure WLAN," First Quarter 2003 issue of *Packet*). Cisco technologies that solved these problems have been incorporated into the Wi-Fi Protected Access (WPA) specification and upcoming 802.11i standard. The Cisco Aironet solution complies with both standards. For more information about WPA, visit www.wi-fi.org/OpenSection/protected_access.asp.

Securing a WLAN takes more than WEP. It includes three components: the authentication framework, the authentication algorithm, and the data privacy or encryption algorithm. All three components are included in the Cisco Wireless Security Suite, which monitors security policies and includes the following features:

- Authentication framework based on the 802.1X standard
- Authentication algorithm based on Extensible Authentication Protocol (EAP)
- Data privacy via Temporal Key Integrity Protocol (TKIP), which includes Message Integrity Check (MIC), per-packet keying, and dynamic key rotation for broadcast and multicast traffic

Cisco Aironet solutions use 802.1X authentication for seamless integration with Cisco AVVID (Architecture for Voice, Video and Integrated Data) networks, which use the same standard for user and device authentication. It provides a layer between network devices such as Ethernet switches and access points, and authentication technologies such as EAP, Public Key Infrastructure (PKI), and digital certificates. This flexibility gives enterprises the freedom to choose technologies that best fit their security requirements and budget.

For WLANs, EAP is available in numerous forms including Cisco LEAP, EAP-Transport Layer Security (EAP-TLS) and types that operate over EAP-TLS, such as Protected EAP (PEAP). PEAP is a draft authentication type designed to allow hybrid authentication using server-side PKI. On the client side, PEAP can use any form of EAP authentication, such as token cards, one-time passwords, and the Message Digest 5 (MD5) algorithm.

The new Cisco Secure Access Control Server Solution Engine 3.2 , a component of Cisco Structured Wireless-Aware Network, supports authentication of several different EAP types, including Cisco LEAP, PEAP, and EAP-TLS. This high-performance, highly scalable, centralized user access control appliance provides a dedicated, locked down security system for centralized authentication of wireless devices (for more on the Cisco Secure Access Control Server Solution Engine 3.2, see "Identification, Please," page 27.)

What's next for WLAN security? "Cisco plans to extend its intrusion detection technologies into RF," says Roshan. "This will give IT managers insight into every aspect of their wired and wireless networks." ▲▲

FURTHER READING

Wireless Management

- ***White paper: The Need for a Cisco Structured Wireless-Aware Network:*** **cisco.com/packet/153_6c1**
- ***Cisco Wireless LAN Solution Engine version 2*:** **cisco.com/packet/153_6c2**

Wireless Security

- ***White paper: A Comprehensive Review of 802.11 Wireless LAN Security and Cisco Wireless Security Suite:*** **cisco.com/packet/153_6c3**
- ***Wi-Fi Protected Access:*** **www.wi-fi.org/OpenSection/protected_access.asp**

Wireless Design

- ***Cisco AVVID Wireless LAN Design:*** **cisco.com/packet/153_6c4**

WIRELESS GEMS

Multiple unencrypted VLANs, 802.11a bridge, and wireless IP phone join Cisco portfolio.

CISCO RECENTLY MADE SWEEPING enhancements to its wireless LAN (WLAN) product portfolio and software feature set with the Cisco Structured Wireless-Aware Network. Many of this framework's new and planned capabilities focus on improving an enterprise's ability to centrally configure, upgrade, secure, and manage hundreds or thousands of wireless access points. Such centralized control is a critical capability that requires a large degree of automation if businesses are to grow their wireless networks on a grand scale.

CISCO IP PHONE

Other Cisco Structured Wireless-Aware Network enhancements help enterprises conduct automated site surveys, "manage the air" by detecting and adjusting signal strength, and discover and alert network administrators of security-policy violations or unauthorized rogue access points.

Along with these management and security capabilities, Cisco has included a few additional wireless gems in its latest rollout that will play a telltale role in the advancement of WLANs. These include a capability called *multiple unencrypted virtual LANs (VLANs)*, an *802.11a WLAN bridge* for interconnecting LANs in multiple sites across a metropolitan area, and Cisco's first *802.11b-capable voice-over-IP (VoIP) phone*.

Partitioning the Air

Support for multiple unencrypted VLANs in Cisco Aironet® access points is especially useful for enabling multiple guest VLANs in WLAN hotspot deployments. "This feature enables ISPs to provide open WLAN access through a single infrastructure targeted to multiple customer groups," explains Sri Sundaralingam, technical marketing engineer in Cisco's Wireless Networking Business Unit.

"Today, users with WLAN adapter cards visiting an airport, convention center, or other large public venue are likely to detect multiple service providers' access points in that location. Alternatively, one access point infrastructure in the venue, owned by a single service provider, could accommodate services from several different providers or several different private groups by partitioning multiple VLANs simultaneously," says Sundaralingam.

So, for example, if a user had a Wi-Fi service subscription with a particular provider, that user would associate only to that service VLAN, and traffic would be kept to that provider's domain.

Bridging the Metro Campus

Wireless bridges have long been used to enable enterprises to interconnect LANs

ILLUSTRATION: LEO ESPINOSA
PHOTOGRAPH: ©KEVIN TWOMEY

in different buildings within a metropolitan or campus area. They eliminate the need to dig trenches, lay cabling, or pay recurring leased-line fees to a service provider. When using over-the-air connectivity with roof-mounted bridges, "there is no recurring operational expense to a service provider. You replace it with a single capital expense that can pay for itself in less than six months," according to Jon Leary, Cisco product line manager for the new IEEE 802.11a-based Cisco Aironet 1400 Series Wireless Bridge.

The new bridge runs at the 802.11a underlying physical-layer rate of 54 Mbit/s across distances of up to 7.5 miles. Actual throughput is usually about 28 Mbit/s, depending on distance and type of traffic.

"Larger packets are transferred over the air more efficiently than smaller packets," says Leary. The 5.8-GHz space avoids interference better than 802.11b and the 802.11g 2.4-GHz band, because it uses four non-overlapping channels (depending on local regulation in the country where it is installed) compared to three channels in the 2.4-GHz band. "But radio waves are not diffracted as well in the 5.8-GHz band relative to the 2.4-GHz band. This makes 802.11a a line-of-sight technology," Leary notes.

There are two versions of the Cisco 1400 Series Wireless Bridge. One ships with an integrated, high-gain (22.5 dBi) directional antenna. The other version allows users to choose one of three remote antennas for professional installations:

- A 9-dBi omnidirectional antenna for connecting multiple locations within two miles surrounding a central location
- A 9.5-dBi sector antenna for connecting multiple locations within a 60-degree arc from a hub location, for better interference rejection than the omnidirectional antenna provides
- A very focused, narrow-beam 28-dBi dish antenna that can transmit at distances of more than 20 miles

Both bridges can be configured in two basic architectures: point-to-point and point-to-multipoint. The architecture chosen will reflect the number of buildings to be connected. For connecting multiple buildings, users can designate one bridge as a hub, or a "root bridge" through which other remote bridges communicate with other sites. This feature means companies do not have to deploy a full mesh of point-to-point bridges to interconnect multiple sites.

According to Leary, customers will likely use sector or omnidirectional antennas as hub devices and more focused antennas as remote bridges or in point-to-point configurations, because those bridges must communicate with only a single other location.

The Cisco 1400 Series Wireless Bridge has been fully sealed and weatherized for outdoor use. The devices contain a high-power amplifier integrated with 802.11a radios to get to the 250-milliwatt transmit power that enables the longer ranges supported by the bridge.

Leary points out that while 802.11 wireless bridges have long existed, the method for 802.11-based bridging has not been standardized. "That is happening now. For example, with bridging the MAC timing needs to be extended to accommodate longer transmission ranges," he says. The 802.11j subcommittee is addressing this mechanism.

Phone Rings Up Wireless Users

The new Cisco Wireless IP Phone 7920 for

CISCO WIRELESS PRODUCTS ROUNDUP AND TIMELINE

Product	Description	Scheduled Availability
Cisco IOS® Software 12.2(11)JA for Cisco Aironet 1100 and 1200 Series Access Points	Cisco Structured Wireless-Aware Network feature that provides support for Wi-Fi Protected Access (WPA), fast secure roaming, 802.1X local authentication, and multiple unencrypted VLANs. No charge to upgrade.	June 2003
Cisco Aironet 1400 Series Wireless Bridge	IEEE 802.11a point-to-point or point-to-multipoint bridge. Used for metro-area LAN interconnections at aggregate data rates of 54 Mbit/s across distances of up to 7.5 miles.	June 2003
Cisco Wireless IP Phone 7920	802.11b-compliant wireless IP phone. Can be used with any Wi-Fi network but extends full complement of Cisco CallManager IP PBX in a Cisco environment.	June 2003
CiscoWorks Wireless LAN Solution Engine (WLSE) version 2.0	Cisco Structured Wireless-Aware Network hardware appliance for centralized configuration and management of up to 2500 Cisco Aironet access points. Includes simplified operation and deployment; proactive performance monitoring; security; reporting, trending, and planning features; and a role-based user access model. Version 2.0 is available to Cisco customers with a previous version of CiscoWorks WLSE.	July 2003
Cisco IOS Software Release 12.2(13)JA for Cisco Aironet 1100 and 1200 Series Access Points	Cisco Structured Wireless-Aware Network feature that provides enhanced RF management tools; autodetection and alerts of unauthorized access points and security policy deviations; assisted site surveys; and interference detection reporting. No charge to upgrade.	Fourth Calendar Quarter 2003
CiscoWorks WLSE version 2.5	Cisco Structured Wireless-Aware Network software upgrade that provides capabilities corresponding to Cisco IOS Software Release 12.2(13)JA, above. Available for customers with a Cisco Service Application Support contract.	Fourth Calendar Quarter 2003

WLAN Migration Tips

In trying to grasp the impact of all the recent Cisco WLAN announcements, you might have some migration and interoperability questions. Here are some tips.

Will Cisco continue to support access points running the VxWorks operating system? Yes. However, VxWorks will not gain any of the new wireless-aware enhancements, so you must upgrade to the Cisco IOS Software operating system to leverage those. The Aironet 1200 Series models AIR-AP1200 and AIR-AP1220 are easily upgradable from VxWorks using a free conversion utility and tool available now. A free version of the utility for the Cisco Aironet 350 Series Access Point will be available in the fourth calendar quarter of 2003 with Cisco IOS Software Release 12.2(13)JA.

Should I support the Cisco Wireless Security Suite, WPA, or both? WPA is a component within the Cisco Wireless Security Suite, and you can mix and match capabilities in the suite. Your main consideration is your client base. If you support Cisco client devices only, you can continue using your existing Cisco Wireless Security Suite, which includes Cisco LEAP and other EAP authentication protocols, and Cisco Temporal Key Integrity Protocol (Cisco TKIP) on both your access points and clients. If you support a mix of clients, note that Cisco TKIP and WPA-TKIP (the algorithms that generate dynamic encryption keys) do not interoperate. Cisco suggests supporting the Cisco Wireless Security Suite (including Cisco TKIP) and WPA (with WPA-TKIP) on your access points, using multiple VLANs, which could then communicate with Cisco or Cisco Compatible clients and WPA clients. Cisco and its Cisco Compatible Extensions partners will offer WPA-certified client cards later this year.

Can I upgrade my existing WLSE to version 2? Yes, if you would like the new features but don't need to scale beyond 500 access points. The CiscoWorks WLSE software for the new appliance, which manages up to 2500 access points, can run on the older appliance, but your previous CiscoWorks WLSE will have a 500-access point limit.

What will I have to do to support AES? The WLAN derivative of Advanced Encryption Standard (AES) that is being specified in the 802.11i standard will eventually require access point hardware upgrades. Cisco will provide products to support AES after the standard is finalized.

Why is fast secure roaming important? Certain client-server applications lose their connection when roaming with basic 802.1X roaming. Cisco has developed the fast secure roaming algorithm to allow authenticated client devices to roam securely from one access point to another without any perceptible delay during reassociation. This feature is a Cisco Layer 2 enhancement to the Cisco LEAP authentication algorithm and supports latency-sensitive applications such as wireless VoIP. Cisco has submitted the algorithm to the 802.11i committee as a possible standard and has licensed it to our Cisco Compatible Extensions partners.

Will Wi-Fi-certified clients from other vendors work with Cisco access points? Yes. Cisco products are part of the Wi-Fi Alliance's interoperability test bed, against which the alliance tests other vendor products for interoperability certification. Any Wi-Fi-certified client cards will work with Cisco's Wi-Fi-certified Cisco Aironet access points; however, enhanced Cisco Aironet features, such as the Cisco Structured Wireless-Aware Network and Cisco Wireless Security Suite, are extra tools and capabilities available from Cisco to our customers.

802.11b WLANs extends the feature set of Cisco CallManager IP PBXs—such as call transfer, conference, hold, and pickup—out to wireless users. "We're seeing demand by users who want to improve their productivity by increasing their reachability around campus—particularly in retail and hospital environments," says Kam Toor, senior product manager in the IP Communications Business Unit at Cisco.

The phone supports six extensions for speed dials and local languages in text menus. Like a Cisco wired IP phone, the handset allows users to make and receive calls using their IP PBX extension in any corporate location, regardless of geography. Quality of service (QoS) is supported downstream from an access point to the phone via a queuing mechanism.

"The AP [access point] sets up two queues and places downstream voice in the higher-priority queue," explains Toor. For upstream voice, he recommends using a voice VLAN. Since 802.11b LANs use Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA), it is possible to configure devices on the voice VLAN with a lower contention window, such that voice devices back off the wireless medium for a lower period than data devices in the event of a collision.

The IEEE standard for WLAN QoS, 802.11e, is expected to be finalized in late 2003, says Toor. "When that standard is final, Cisco will support it, and users can upgrade their phones with a centralized firmware push over the air," he says.

For security, Wired Equivalent Privacy (WEP) and IEEE 802.1X with Extensible Authentication Protocol (EAP)-Cisco LEAP are currently supported. "You power up the phone, and it authenticates you like a PC," says Toor. ▲▲

WIRELESS

The Difference Between Business-Class

BY GAIL MEREDITH

CISCO HAS LONG BEEN THE MARKET LEADER in corporate wireless LANs (WLANs) with its Aironet® wireless networking solutions. Now, with its recent acquisition of Linksys, the company is also a leader in WLAN solutions for the home. "Wireless networking in the home is a rapidly growing market," says Charlie Giancarlo, senior vice president and general manager of product development at Cisco. "Linksys is the market leader because its solutions offer the right combination of features at the right price point in a simple plug-and-play package. Cisco backs this with the addition of a strong corporate brand and our own commitment to customer service—consumers know we'll be around to support our customers."

The Linksys acquisition provides a great complement to Cisco's Aironet wireless networking solutions, because each product set offers distinct value to different markets: Linksys offers small-scale, unmanaged network solutions for the home or SOHO customer; Aironet targets any business or enterprise that requires secure, scalable, managed wireless networks. While the underlying technologies are based on the widely deployed IEEE 802.11 wireless standards, the two different solutions offer tradeoffs in price, performance, functionality, and scale that meet the disparate requirements of each market. Understanding these differences can help you choose the best solution for your needs.

Wireless for the Enterprise

The enterprise WLAN is typically an extension of a wired Ethernet LAN, such as a network based on Cisco AVVID (Architecture for Voice, Video and Integrated Data). Ideally, this wireless extension should support the same functionality as the wired network, providing tight controls through rich feature sets, such as security, management, and quality of service (QoS). Enterprise WLANs must also scale from hundreds to thousands of devices. "Viewing the wireless LAN as an extension of the wired LAN helps to simplify operations, management, and security," says Christine Falsetti, director of marketing for mobility solutions at Cisco. Cisco Aironet wireless networking solutions are designed to meet these business-critical requirements.

Security

Security is a major concern for enterprise network managers, especially in light of well-advertised weaknesses in the Wired Equivalent Privacy (WEP) standard. WLAN security is provided for Aironet products by the Cisco Wireless Security Suite, which includes IEEE 802.1X authentication for WLAN clients and advanced encryption through Temporal Key Integrity Protocol (TKIP). "We introduced 802.1X support in the Aironet products more than two years ago," says Chris Bolinger, product marketing manager in the Wireless Networking Business Unit at Cisco. "In fact, you could say that we pioneered the standards-based approach to robust wireless LAN security." Through its Aironet line, Cisco provides support for the broadest range of 802.1X authentication types in the industry, including support for Cisco LEAP, EAP-TLS, and types based on EAP-TLS, such as PEAP. (See "Securing Wireless," First Quarter 2003 issue of *Packet*®.)

"Now, the entire WLAN industry is embracing 802.1X and TKIP, because they are the key components of an industry standard called Wi-Fi Protected Access, or WPA," adds Bolinger. When Cisco Aironet customers use the Cisco Wireless Security Suite they now have the option of securing their WLANs with WPA, or any of the authentication types listed above, depending upon their networking needs.

Cisco Aironet access points also support IEEE 802.1X local authentication, and can be configured to act as a local Remote Authentication Dial-In User Service (RADIUS) server to authenticate wireless clients when the authentication, authorization, and accounting (AAA) server is not available. This provides authentication services for remote or branch office WLANs without a RADIUS server, and backup authentication services during WAN link or server failure. As an alternative WLAN authentication and encryption method, Cisco Aironet products support virtual private networks (VPNs). Cisco also recommends the use of virtual LANs (VLANs) in the wireless edge, separating traffic, applications, and users for security and performance reasons.

Management

The Cisco Aironet series is also a key component of the new Cisco Structured Wireless-Aware Network introduced in June 2003. The Cisco Structured Wireless-Aware Network is a comprehensive

ILLUSTRATION: LEO ESPINOSA
PHOTOGRAPH: © KEVIN TWOMEY

framework for deploying, operating, and managing hundreds to thousands of Cisco Aironet access points in a Cisco network infrastructure. This innovative approach combines the Cisco switch and router infrastructure with the Cisco wireless network, to make one integrated "wireless-aware" network. It extends to the wireless LAN the same level of security, management, ease-of-deployment, scalability, and reliability that customers have come to expect in their wired LANs.

Enterprise networks also need QoS features to separate and control data, voice, and video traffic, especially in Cisco AVVID environments. Cisco Aironet solutions support appropriate Cisco IOS® Software edge control mechanisms to classify, prioritize, and queue packets. New features such as Fast Secure Roaming support rapid reauthentication as wireless IP phones traverse access points in the network. (For more information on Cisco WLAN security and management solutions, see "Reining the RF," page 36.)

Without a home network, each computer would require its own dedicated modem line for Internet access. Sharing one broadband connection over a LAN makes economic sense, and enhances performance and the overall end-user experience. A WLAN provides this connectivity without the costly and awkward task of pulling Category 5 Ethernet cable throughout the home. For these reasons, WLANs have fast become the home network solution of choice, providing a quick, convenient, and easy way to share computer resources, while offering the additional benefit of user mobility. Client devices such as laptop PCs and personal digital assistants (PDAs) stay connected on the wireless network as people carry them from room to room. What's more, with an IP-based network in place, the home user can begin taking advantage of a host of new networked applications targeted to the home.

"IP is the common thread for new services into the home," says Rob Redford, vice president of product technology marketing at Cisco. "The ability to have a broadband connection to every individual is essential to the long-term productivity of every country

WORLDS APART

and Home Wireless LAN Solutions

Home Wireless Networks

Today, wireless networking in the home is a vibrant and growing market. This trend reflects modern families' growing dependence on the PC for work, school, and play. Powerful PCs are increasingly less expensive, resulting in more households with multiple computers scattered through several rooms. A home WLAN enables these computers to share a single Internet connection, usually a broadband connection, as well as share applications, files, and printers. Home WLAN solutions must be affordable and reliable. More importantly, they need to be easy to install and manage by the average consumer, not an IT professional. The sophisticated features of an enterprise-class solution such as Cisco Aironet are unnecessary in most home networks.

in the world. Broadband connections and wireless networking create opportunities for new levels of home automation. For example, Wi-Fi security cameras make it cost-effective for people to monitor activities through the Internet while away from home. In addition, a wireless network can be used for entertainment purposes. Wireless media adapters allow users to stream digital music and pictures stored on their PC to be played through the home entertainment center or TV," says Redford.

Linksys Product Overview

The Linksys product line offers a range of standards-based routers, access points, and device connectivity cards designed to fit any scenario for the home or very small business. Linksys offers easy-to-

Wireless Standards Support

Cisco wireless networking solutions support current and emerging standards. They include:

IEEE

802.11a—defines 54-Mbit/s bandwidth in the 5-GHz band; low usage makes this band least subject to interference.

802.11b—defines 11-Mbit/s bandwidth in the 2.4-GHz band; this is the most widely deployed standard and, therefore, is the most subject to interference.

802.11g—defines 54-Mbit/s bandwidth in the 2.4-GHz band; this standard has wider range than 802.11a signals, but is also subject to 2.4-GHz channel interference.

802.11i—draft security standard that defines Advanced Encryption Standard (AES) encryption and Extensible Authentication Protocol (EAP) device authentication for 802.11 wireless networks.

802.1X—defines port-based authentication in LAN environments; authentication credentials are never transmitted over wireless connections without encryption; EAP authentication types such as Cisco LEAP are components of 802.1X.

Wi-Fi Alliance

Wi-Fi Protected Access—interim technology to strengthen user and device authorization/authentication schemes prior to release of 802.11i standard. It includes support for 802.1X, EAP authentication, and TKIP encryption.

use, value-priced solutions that support 802.11 wireless, broadband, 10/100 Ethernet, and plain old telephone service (POTS) connections.

Linksys products support the widely adopted 802.11b (11-Mbit/s) standards, along with higher speed 802.11a and 802.11g standards (see sidebar). Many products support multiple standards, enabling migration to higher speeds without replacing the box. And because consumers should not need to understand networking, the Linksys Web browser-based management program enables easy setup and basic security.

The basic functionality of Linksys WLAN products includes security based on the WEP protocol for small-scale users. "It's easy to change a WEP key on a single access point," says Steve Troyer, director of marketing for new business ventures at Cisco. "While WEP is insufficient for corporate security, it works well in the home. Telecommuters who need additional protection to access a corporate network from home can use their corporate VPN tunnel."

Linksys offers competitive pricing due to its low operating costs and "SOHO-sized" feature set. "Linksys evaluates, tests, and markets the best designs offered by original design manufacturers [ODMs]. They also provide some very easy installation utilities and superior end-user telephone support," says Troyer. "Cisco does not plan to alter this business model in the future. It works for the consumer market because those customers want reliable and easy-to-use products that do exactly what they want at an affordable price."

The Linksys product suite includes the following types of devices:

- Broadband routers—with four or more wired 10/100 switched Ethernet ports, and wireless connections up to 54 Mbit/s. Single-band 802.11b and 802.11g (draft) solutions inexpensively enable 11 Mbit/s or 54 Mbit/s wireless connectivity, while dual-band options provide an upgrade path from 802.11b to the 54-Mbit/s speeds of 802.11a or 802.11g standards. Each router includes an Ethernet WAN port and multiple Ethernet ports for wired devices such as printers or other PCs. It also provides Network Address Translation (NAT) between the private home network and the public ISP network for security, and Dynamic Host Control Protocol (DHCP) for plug-and-play operations of PCs and other devices. Linksys products have been carefully designed to work seamlessly with a wide variety of broadband services, including all DSL and cable services, with little to no user configuration. The most popular models have integrated wireless access points for greater efficiency.
- Wireless access points—plug-and-play devices that provide instant wireless network connectivity. They are designed to be used in conjunction with existing broadband routers to extend Internet access throughout the wireless network.
- End-user adapters—available in a variety of formats, including PCI, USB, PC card, network adapters, and CompactFlash card, with single- or dual-band options. They interoperate with nearly every IP-enabled consumer device, from desktop and laptop PCs to PDAs and game boxes.
- Wireless Ethernet bridge—standalone unit designed to connect devices with Ethernet ports, such as video game consoles, to support online applications, such as Microsoft Xbox Online or Sony PlayStation networked games.

Something for Everyone

Home and corporate networks have distinct requirements, and Cisco now offers solutions for both. Linksys solutions for home and small-office networks deliver the inexpensive, reliable, and easy-to-use solutions that consumers demand, while Cisco Aironet solutions continue to raise the bar for excellence and integration with rich business-class features, robust security, reliability, management, and scalability. ▲▲

FURTHER READING

- **Home wireless networking: cisco.com/packet/153_6b1**
- **Enterprise wireless networking: cisco.com/packet/153_6b2**

TIPS FROM THE TRENCHES

Cisco customers share wireless experiences.

BY JOANIE WEXLER

HOW ARE ORGANIZATIONS using wireless technologies and what direct benefits can they measure? Not surprisingly, interviews with Cisco customers reflect the continued success of wireless in vertical applications, where returns are quickly achieved and easily calculable. From there, though, network implementers reveal that they are finding additional uses for the technology they hadn't necessarily anticipated.

Pushing out wireless LAN connections to certain knowledge workers, for example, allows executives to respond to critical e-mails or grab nuggets of information during meetings that can seal a deal or move a project forward. In addition, the fluid nature of organizations these days has business partners and consultants frequently coming onsite. So some businesses feel that offering mobile Internet access to visitors is an important capability. Some organizations are offering wireless Internet access as a competitive amenity. And still others are contemplating location-based applications that would not have been possible without wireless mobility.

ILLUSTRATION: LEO ESPINOSA
PHOTOGRAPH: © KEVIN TWOMEY

Let's see how some early wireless adopters are faring.

Furniture Retailer Slashes Inventory Times

For Art Van Furniture, a Warren, Michigan-based furniture retailer, Cisco Aironet® wireless access points have helped the company slash its inventory process from about seven days to a single day. These savings, multiplied across the company's 30 statewide stores, quickly paid for the investment after one triannual inventory cycle, says Mike McDonald, manager of technical services at Art Van Furniture.

Before going wireless, the company handled inventory manually. Each store location features 60,000- to 100,000-square-foot showrooms, each with 3000 to 5000 pieces of furniture plus thousands of small accessories to be inventoried.

Until 2001, explains McDonald, nearly all personnel at each store used pad and paper to record counts of all products. They would then send that information to Art Van's corporate headquarters, where it would be manually keyed into an inventory application.

The application generated hard-copy reports, which would be exchanged between stores and headquarters multiple times for verification and corrections until everyone was satisfied that the information was accurate.

"We figured there must be a better way," McDonald says. "So we barcoded all products on the floor and laser-printed sales tags so they could be scanned." Art Van installed five Cisco Aironet access points and a Cisco Catalyst® 2950 switch for an initial single-store pilot. The company wrote a simple scanning application to run on handheld devices.

The first test of the system enabled one location to be inventoried in about five hours using 15 people, compared to requiring 30 people for seven days previously. The elimination of the back-and-forth manual checking and correction saved still more time, McDonald observes. "Today we book inventory numbers into our general ledger the same night we do inventory," says McDonald. "And we have virtually eliminated all procedural errors."

Art Van now runs approximately 200 access points, including 42 in its main, 1-million-square-foot warehouse, where receiving, delivery-scanning, and bulk-inventory applications are being rolled out. The furniture company uses the CiscoWorks Wireless LAN Solution Engine (WLSE) to centrally configure and manage the access points.

"We have had the luxury of turning those devices [access points] off when not doing inventory so that we don't invite people to challenge its security," says McDonald. "However, new applications like receiving and point of sale mean we'll need the infrastructure all the time, so we'll have to tighten security. Now that the environment is there, we can do lots more with it."

McDonald says the company is discussing, for example, using the 802.11 network for voice applications and invoice approval by store managers. "Store managers could be notified wirelessly that there's a deal pending and they could approve it without leaving their offices," McDonald explains.

Microsoft Takes on Security and Rogue Access Points

Worldwide software titan Microsoft Corp. began investigating wireless LANs in 1999, when the initial business driver was to create an environment for developing wireless-aware operating systems and applications.

The company settled on a Cisco Aironet infrastructure because of its "enterprise network awareness. There were strong management capabilities built into the access points," recalls senior network engineer Don Berry, who was part of the team that oversaw Microsoft's global wireless deployment.

A post-pilot survey of the rollout indicated that each of Microsoft's wireless users—whether knowledge workers or developers—gains between 60 and 90 minutes of productivity per day.

"A senior vice president stated that wireless saves him 1.5 hours per day. It doesn't take too many of those improvements to quickly realize big returns. You can imagine the savings on such a large scale," says Berry.

About 4000 Cisco Aironet access

802.11 Making News in China

For temporary, large-scale government conventions and sporting events, China is realizing the benefits of WLANs. Starting last year, thousands of participants in the Asia-Pacific Economic Cooperation (APEC) in China and the subsequent four-day Bo'ao Economic Forum began using wireless access to on-demand information. Journalists could file stories via wireless Internet access on the spot.

To serve the Bo'ao Economic Forum, Cisco, in cooperation with Hainan Telecom, (a branch of China Telecom) and Cisco-certificated partner Cayee Co. installed Cisco Aironet 802.11b wireless access points and Aironet wireless bridges for metropolitan-area communications. Participants used Cisco Aironet WLAN client adapters for network access. In the news center, 13 access points served about 300 journalists.

The Bo'ao Hainan project followed a separate effort by Shanghai Telecom to install about 20 Cisco Aironet access points in office, residential, and public areas, for secure, generally available wireless Internet service. Services were provided to journalists, for example, at the recent Asia Men's Basketball Tournament.

—*Packet* Magazine Staff

points provide wireless coverage to 14 million square feet of Microsoft workspace in 70 countries. Berry estimates current deployment costs at about 50 to 60 cents per square foot in the US and about US$1 outside the US.

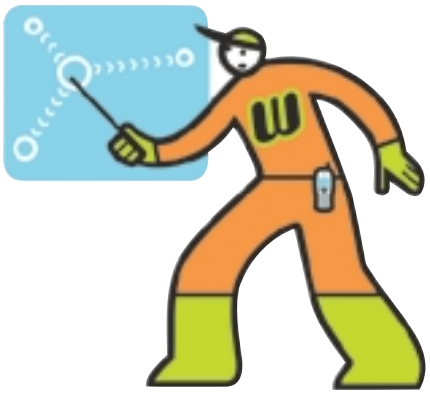

At the time of Microsoft's initial deployment, static Wired Equivalent Privacy (WEP) was the only Layer 2 security available for 802.11 LANs. So the company worked with Cisco, its own Windows development group, and the IEEE standards body to get 802.1X into the standards specifications, along with the Extensible Authentication Protocol (EAP). They then incorporated 802.1X and EAP into the Microsoft XP operating system.

Microsoft is currently using certificate-based EAP-Transport Layer Security (TLS) but is moving to Protected EAP (PEAP), which is based on a username/password combination and is easier to manage, Berry notes.

"When we moved to EAP/802.1X, we cut off a significant portion of our user base for about eight months because they were running previous operating systems that didn't support the security technology," says Berry. "Now with our employees being on XP and 802.1X alternatives for the Windows 2000 and PocketPC operating systems, the deployment has moved companywide."

Microsoft's initial design goal was to deliver performance on par with DSL connectivity, or about 600 to 700 Kbit/s per user, recalls Berry. He says that goal translated into about ten users per access point.

After extensive antenna, coverage, and access point-mounting testing, "We learned to space APs [access points] about 60 feet apart," Berry says.

The biggest challenge? "Rogue access points," says Berry.

He says that because people who know little about wireless are usually the ones plugging in rogue access points, the devices often remain set to the highest power level, which causes network interference. "Corporate policy is that we reserve the right to hunt down rogues and confiscate them," Berry says.

Berry anticipates the next big initiative as being "location awareness of people, places, and things" and says the company is considering wireless voice over IP (VoIP) for field locations.

"We're considering a 'network-in-a-box' for field locations, with a DSL connection, tunneling hardware, and wireless interface" so that mobile field workers could set up shop with complete corporate data and voice connectivity wherever they happen to be, he says.

"Outside forces with laptops, 802.11 connectivity, and cell phones are becoming a significant percentage of the company," Berry observes.

A Necessity at Palm

At wireless handheld manufacturer Palm Solutions Group, there was "simply no way to avoid an 802.11 solution," says Bill Stevenson, chief infrastructure architect.

Stevenson explains that the company, where about 50 percent of employees have 802.11b network interface cards, outsources work on a large scale. Vendors and consultants are frequent visitors who require Internet access, which Palm provides wirelessly. "And we have a small work force that needs to be connected all the time," Stevenson adds.

The company recently moved into new headquarters in Milpitas, California, and took advantage of the move to standardize on a single wireless infrastructure. Previously, an acquisition and other circumstances had resulted in a multivendor environment that was difficult to manage. For example, each access point had to be configured and managed separately.

"We needed to apply standards to enable centralization," says Mike Allison, Palm's global IT director.

So the company had 21 Cisco Aironet 1200 Series access points installed at headquarters and one access point in each of six remote sites, says Allison. The company selected the Aironet 1200 for its capability to add support for 54-Mbit/s 802.11a wireless networks without having to replace the entire unit, says Allison.

He also cites Cisco's power-over-Ethernet capability that saves the company several hundred dollars (US) per access point in wiring and power-outlet installation costs. Cisco Aironet access points are powered over the cabling that connects them to their Cisco Catalyst 6509 Ethernet Switch.

Palm says it has also recently placed an order for the new CiscoWorks WLSE 2.0 (which is part of the Cisco Structured Wireless-Aware Network, see article on page 36), and eagerly anticipates the sophisticated functions coming in version 2.5 later this year.

"We like the security logging and tracking features of the new WLSE version," Allison says. "We can find out how APs are being used, how often they are being used, and how often users roam. All this enables us to anticipate usage and accurately load-balance."

He also says it will be a "great benefit to be alerted to rogue APs" by the CiscoWorks WLSE 2.5 version. "Currently, I have one employee who walks the campus once a week with a scanner to detect them. This will save a lot of time."

For security purposes, Palm has placed all 802.11 devices on a virtual LAN (VLAN) behind a firewall. All access to the corporate network is provided through a Cisco Virtual Private Network (VPN) 3030 concentrator, and Palm employees run Cisco VPN client software.

◆ ◆ ◆

This small sampling of wireless installations reflects the fact that while there might be a single business goal driving an initial wireless installation, enterprises tend to get creative with the technology once the infrastructure is in place. Piggybacking on the network for additional internal applications or to provide amenities to customers allow wireless network implementers to "double dip" to make the most of their investments. ▲▲

JOANIE WEXLER is an independent analyst and editor in the computer networking industry, and a frequent contributor to *Packet* magazine.

MOBILE DATA SERVICES: BEYOND THE HYPE

Considerations When Rolling Out a Mobile Data Service in Your Enterprise

BY COLIN SEWARD

Twenty years ago Microsoft's Bill Gates had a vision: a computer on every desk and in every home. Having pretty much achieved this vision, high-tech industries have set their sights on enabling users to take their computer with them wherever they go and stay connected to the Internet around the clock. The resulting impact in both culture and technology poses enormous challenges for IT departments as the physical boundaries of the workplace further dissolve.

This article looks at some of the considerations in rolling out mobile data services in an enterprise, and how IT at Cisco is approaching the challenge of providing anytime, anywhere, any-device access for its employees.

An Industry of Hype

Like so many new technologies, wireless access and mobile data devices have generated a lot of hype. As any enterprise that has piloted the use of remote access using 2.5G or 3G networks, public wireless LAN (WLAN) hotspots, new personal digital assistants (PDAs), or smartphones will know, these technologies, like many others, are not without their shortcomings.

These technologies are moving along the hype curve, and while they have enormous potential, they can often still be difficult to deploy, manage, and support. Nevertheless, several enterprises are deploying solutions using wireless access technologies and mobile data devices, and doing so successfully.

Have these technologies finally crossed the chasm, as described in Geoffrey Moore's popular 1991 book, *Crossing the Chasm*? And if so, what are the key ingredients that go into successful implementations?

Learning the Hard Way

Two years ago at Cisco, we ran a pilot giving 60 end users a General Packet Radio Service (GPRS) phone and PDA to go with their Cisco laptop. They were instructed on how to use the GPRS phone to access simple corporate Wireless Access Protocol (WAP) applications such as e-mail and the company directory. They were also instructed on how to use the PDA and laptop over a WLAN (IEEE 802.11b) within the office and over GPRS externally to access other applications such as an e-sales portal and e-learning materials. Several problems surfaced with the technology during the pilot; for example, configuration of the PDA was complex and time-consuming for end users, and when the battery drained the configuration was lost. Support calls were difficult to resolve because the Cisco help desk was centralized and unaware of local GPRS problems. Issues around usability and security existed, and improving one invariably impacted the other. Nevertheless, solutions for many of these problems were developed during the pilot, usage continued to rise, and feedback from the end users improved.

After the pilot, the devices and services were left with the end users, and Cisco moved into an adoption test phase. Utilization during a pilot is always artificially high as end users try to make

LEO ESPINOSA

things work to provide feedback. During the adoption test, end users were no longer sent surveys for feedback. Instead, their day-to-day usage patterns were measured to determine whether the new devices and services had been adopted into their working life. After three months, we went back to the end users armed with the metrics and tried to understand what was working, what wasn't, and what we needed to do differently.

Which Device?

In the early pilot, we found that while it was technically possible to deploy a solution with the appropriate levels of security and with applications that appeared to deliver benefits, on some devices the setup was too difficult for end users to maintain. The biggest disappointment was WAP on the cell phone, which, given its subsequent problems in the industry, did not come as a huge surprise.

The next biggest disappointment was using the PDA over GPRS. While this had great appeal in theory, in reality secure access via a virtual private network (VPN) client using token authentication involved juggling three devices and defeated all but the most dexterous and technology-savvy users.

The great success of WLAN access from the laptop perhaps wasn't a surprise, and this has since been rolled out to every Cisco office worldwide. What did come as a surprise? The success of GPRS access from the laptop. Although slow by ADSL standards, end users were finding that once connected from the laptop they could access all their existing office applications. Perhaps, this success showed that it's easier to tune the technology to users than vice versa.

At least for now, the likelihood of an enterprise being able to support *any device*—anytime, anywhere—is slim from a reasonable total-cost-of-ownership perspective.

Based on these findings, our internal strategy divided into two initial tracks. One track, called Extended Remote Access, focused on providing better access from the laptop to our corporate network for users working away from the office or home. The second track, called Cisco Pocket Office, focused on the plethora of new handheld devices springing up within the company. Based on the adoption challenges experienced in the pilot, Cisco Pocket Office looked at providing four main service components: configuration, management, applications, and support.

The first handheld devices we would tackle through Cisco Pocket Office would be employee-owned PDAs, as they had already reached critical mass. But the same model would subsequently be extensible to corporate-owned smartphones and other handheld devices. In the long term, corporate-owned smartphones would likely be the dominant handheld device since Cisco already provides cell phones to many employees.

Mobilizing the Laptop

The number of ways for users to gain access to the corporate network when out of the office is growing: dialup, WLAN hotspots, broadband connections in hotels, guest access from another company's network, and so on. For mobile data access from the laptop using GPRS or Code Division Multiple Access (CDMA) 1xRTT, bandwidth optimization is a key ingredient to improve performance for end users and reduce cost to the company.

Bandwidth optimization can be applied at a number of levels, from automatically shutting down bandwidth-intensive applications when a data session is launched, to client-server HTTP optimization, application-specific optimization (such as plug-ins to Exchange for e-mail), to transformation of the Web site to simplify the content.

Optimization is also beginning to be offered by some mobile network operators as part of their data solutions; however, this only works on unencrypted traffic. If the enterprise is using an end-to-end VPN tunnel, the carrier's optimization will not work, and the enterprise will need to implement optimization behind its own firewall. If the carrier is using the same optimization software as the enterprise, there can be additional complications. In this case, end users might get confused by having two clients on the desktop, or by having the upgrade of one client affect the other.

This scenario can be further complicated if the end user utilizes multiple carriers, each with their own solution! Solutions are starting to emerge that might begin to simplify the tasks for end users and bundle these various services and applications behind one interface or put more intelligence into the network; however, such solutions are still in their infancy. Obstacles other than low bandwidth also exist with 2.5G and 3G connectivity, including cost (particularly when roaming), inconsistent service levels, and coverage. Some level of end-user training and expectation-setting is, therefore, important in deploying 2.5G and 3G to ensure that costs are contained and users are satisfied with the service.

Choosing a Handheld Device

At least for now, the likelihood of an enterprise being able to support *any device*—anytime, anywhere—is slim from a reasonable total-cost-of-ownership (TCO) perspective.

Even ignoring the handheld operating systems that have not reached significant market penetration, an enterprise is still left with at least four operating system choices, and there is no one right way to choose among them.

Palm has a large established user base and plenty of loyal supporters. Desktop engineers accustomed to working with Microsoft will find the transition to working with Windows CE-based devices, such as Pocket PC, Pocket PC Phone Edition, and Smartphone 2002, reasonably easy. RIM, the first manufacturer to develop an enterprise-oriented, end-to-end wireless offering for e-mail, sets the benchmark for ease of use in this area. Lastly, Symbian, although a newer entrant, has the backing of most

handset manufacturers and, over the next year, will start to appear in a lot of midrange as well as high-end handsets.

Each operating system also has its downside, and if end users are left to buy their own handheld devices, it is likely that there will be a variety of types within an organization. Some enterprises might decide to standardize on a device, others on an operating system.

In the Cisco Pocket Office track, we have focused on one operating system initially while acknowledging that we might need to extend this over time because of the variety of devices owned by employees throughout Cisco.

What About Security?

Security is obviously an important consideration both when connecting back to the corporate network and in securing the data on the device. As mentioned, many security solutions reduce device usability and can become a barrier to adoption. So, this is an especially important area in which to strike the right balance.

Enterprises should at least have guidelines for acceptable usage for the new handheld devices and recommendations on how to secure them. Depending on the device type and likely usage, these guidelines might include virus protection, data encryption, password strengthening, a VPN client for secure access back to the corporate network, memory wipe on certain triggers, and so on.

Several security products exist to perform these tasks, but given most enterprises' focus on TCO, the preference is most likely to obtain a product from the vendor providing the same function on the desktop. Unfortunately, many of these vendors are also struggling in working out which operating system to support, and many have chosen to wait for additional signs of growing demand before committing to developing their solutions for these new platforms. As a result, enterprise IT departments that want to move quickly to lock down handheld devices might find themselves having to purchase new licenses and train themselves in supporting a new product.

COLIN SEWARD

DR. COLIN SEWARD is a manager in IT, Cisco EMEA (Europe, Middle East, and Africa). He has worked in IT and consulting for more than 10 years and leads the Mobile Centre of Excellence at Cisco, looking at the internal deployment of mobile technologies, devices, and applications. He can be reached at cseward@cisco.com.

Managing the Device

While there is not as much software on a handheld device as on a desktop computer, installing, configuring, and maintaining any additional software required by an enterprise can be complex and time-consuming for the end user. Because this complexity can be a barrier to adoption, automating the process for end users can be just as important for handheld devices as for desktop computers.

If a small number of handheld devices are to be supported, it might be possible to produce images with the standard configuration for those devices. In Cisco Pocket Office, rather than produce images for the devices, we have used device management software to install and configure each required application.

We have found this process to be more flexible as it can accommodate differences in devices introduced by the manufacturer or mobile operator, and applications already installed by the end user are not overwritten. The device management solution can also be used to perform policy enforcement, so that the IT department can make sure that the security enhancements stay correctly configured and remain enabled.

Whichever approach is taken, it is useful to design this process so it can be done *in situ* by the end user rather than make end users bring their devices into the IT department for configuration. This process not only reduces the operational cost of supporting these devices but for devices where the memory is lost when the battery drains, end users can quickly restore the image themselves.

In time, many of these device management services will probably be offered by mobile network operators or other third parties; however, for the moment, early adopters will likely need to do much of the work themselves.

Applications to Justify the Investment

If handheld devices are deployed to improve a specific business process or for a user base with a specific job function, the enterprise might already have an obvious "killer application." For enterprises with knowledge workers, where the demand for handheld devices has grown from end users adopting them as a personal information manager (PIM) device, a killer application might be less apparent.

While having PIM information on a handheld device is a convenience for many end users, by itself it probably won't justify the IT investment required to secure, manage, and support these devices. The challenge for many enterprises has, thus, been in demonstrating the business justification and finding a convincing return on investment (ROI) before initiating the project.

To build that business case, use of these devices clearly needs to go beyond PIM. An early focus for handheld applications was in synchronizing useful data for use offline, or for some basic form-filling for sales or field service functions. Now, as wireless e-mail gains in popularity, many organizations are looking to it as the killer application. And for many users, this might, in fact, be close to the truth.

◆ ◆ ◆

Have mobile data services really crossed the chasm and become ready for the mass market? Underneath the hype, it is clear that the relevant technologies are far more robust than they were a year ago, and a lot of the issues that hampered successful deployment of such solutions can now be solved. Although mobile data services are still not trivial to piece together, there is enough information available today for enterprises to decide their corporate strategy.

Many organizations, including Cisco, have completed their pilots and are looking at moving some of these services into production. The vision of "anytime, anywhere, any device" is getting a little bit closer every day. ▲▲

Special REPORT

HIGH AVAILABILITY NETWORKING

A Guide to Defining, Measuring, and Improving Availability in Your Network

Part One: DEFINE FIRST, THEN MEASURE

NETWORK AVAILABILITY is more important today than ever before. The heat is on because users expect and need so much more from networks than in the "simple" days of e-mail and file sharing. Today, enterprises maintain on their networks a multitude of mission-critical applications that drive their businesses: videoconferencing with partners, Web-based order entry, payment processing, and customer relationship management, to name a few. The network simply cannot go down.

These types of applications are a cornerstone to making operations much more efficient—and are absolutely critical to an enterprise's viability and productivity. Meanwhile, they also make networks more complex and availability harder to get a handle on. *The crucial concept to remember here is that availability must be defined for the network as a whole.* The probability that users can get to applications is the sum of the availability of individual devices and paths.

By JANET KREILING

Fortunately, the data for analyzing the availability of multipurpose data networks are all there for the asking: the Simple Network Management Protocol (SNMP), ping, trap, and trouble ticket reports most enterprises already collect, as well as from published statistics by the vendors of network devices. Additionally, the software exists to put this data together in useful ways so network managers can see the level of availability their users are experiencing, and also analyze and project problems to be fixed before they affect users.

Having *five nines*—99.999 percent availability, or just a few minutes of downtime a year—has been the traditional aim for service-provider networks. But, points out Jim Thompson, a high availability consultant in the Advanced Services Group at Cisco, "In a multi-application network, availability isn't an issue of how many 'nines' you have. It's having a highly available network where you need it—for your mission-critical applications."

So, if your Web site is your main order entry channel, you want its network infrastructure to be highly reliable. If your company runs 15,000 automated teller machines, you need highly available links to all 15,000. And so on.

Thompson defines availability for enterprises as *a question of what the user sees*. "It's not about whether a given device or link is down," he says. "It's about 'Can I get to the application I need?'"

The user's perspective reflects short-term events in the network. The network manager's perspective should encompass the long term so as to spot and fix potential problems before they affect users.

Metrics and Why You Use What

Several equations describe the basic concepts that define availability:

$$\text{Availability} = \frac{1 - (\text{total connection outage time})}{(\text{total in-service connection time})}$$

OR

$$\frac{1 - \text{downtime}}{\text{uptime}}$$

The definition of a connection is the successful data transfer from end device A to end device B, involving physical connectivity, link-layer protocol connectivity, and network-layer protocol connectivity. Some organizations choose to add a time factor to further define satisfactory availability.

This overall number can be derived primarily from trouble ticket reports on the network. However, overall availability does not give you very much detail. You need to chart availability for all of the devices in your network, the paths, and perhaps the applications.

Availability for devices can be expressed as the following equation for Mean Time Between Failure (MTBF) and Mean Time To Repair (MTTR):

$$\text{Availability} = \frac{\text{MTBF}}{(\text{MTBF} + \text{MTTR})}$$

For example, if MTBF = 38,830 hours and MTTR = six hours for a type of Internet connection device, device availability would equal .999837116 or 99.9837116 percent. Essentially, this equation states that the availability of a device equals the MTBF for that device lessened by the MTTR, whose effect will be slight because the MTTR is quite small compared to the MTBF.

Availability for paths can be expressed as:

$$\text{Availability} = \frac{\text{MTBF}(x)}{(\text{MTBF}(x) + \text{MTTR}(x))}$$

Where x equals a failure due to software, hardware, power, human error, or path convergence time in redundant systems.

These equations are used by network managers not only to actually calculate availability but also to help them understand generally which factors must be tracked to produce the calculation. In fact, Thompson says he usually begins with availability as a known quantity that he has calculated from data collected by the enterprise, and data from Telcordia and vendors. He can then solve the equation for one of the other factors such as MTBF. Actually, all the solving is done by software developed by Cisco that can be customized by individual companies. The software incorporates the equations.

Clearly, the most important data are MTBF and MTTR, which should be collected and analyzed for every device type in the network and every path. To gather data on both devices and paths, organizations typically use:

- Impacted user minutes (IUMs), derived from trouble tickets
- Link and device status, derived from SNMP polling and traps
- Internet Control Message Protocol (ICMP) reachability, derived from ping packets sent as tests

Each metric yields different data, and its use must be tailored to the goals of your organization. *Pings* tell you whether a path from and to given ports is open. But if you define downtime over too short an interval, you might report troubles that don't really exist. *SNMP data* tell you about server uptime; correlate with the ping data and you might find that you have just a port problem, not a server problem. If your network is configured so another port may be used, you don't have an availability problem from the user's point of view. You do have one from the network manager's perspective, however. *Trouble tickets* can tell you what was out, for how long, what users and applications were affected, and what the fix was.

The data is collected by fault management software such as Concord Communications eHealth suite, Hewlett-Packard OpenView. (Cisco supports Management Information Base, MIB, extensions from most of its routers and switches to the fault management software, facilitating data collection from them.) In addition, Cisco Service Assurance Agent, a feature of Cisco IOS® Software, enables availability and response time reporting between Cisco devices using ICMP echo and System Network Architecture (SNA) pings, and can assess availability even of individual services (via their ports). The Cisco Internet Performance Monitor uses this data to provide real-time and historical analysis.

The raw data can then be parsed by additional software. Some companies use customized programs based on core analysis software developed by Cisco.

Calculating Theoretical Availability

Thompson usually begins an analysis of a network's availability by calculating the theoretical availability of three categories of components—hardware, software, and power supplies—in each of the three network segments—*access*, *distribution*, and *core*. Before a company has enough data to plug in its own figures, Thompson uses information from product data sheets, which provides data giving the theoretical MTBF for the devices used in the company's network; theoretical data on software from vendor support organizations; and data on power supplies.

"This approach allows us to look at what is theoretically possible for a company to attain and helps set expectations," he says. "Companies will often start with a very high availability objective, such as five nines, but with the theoretical reliability data in front of them, they can usually see where and how they need to modify that objective."

Putting the theoretical assessment into action, one such calculation for an enterprise produced availability ratings of 99.9702 percent for the access layer; 99.9999 percent for distribution; and 99.9999 percent for the core. The resulting estimated path availability was 99.97 percent. "The math always yields the lowest common denominator," Thompson notes. "It's common sense, that a chain is only as strong as its weakest link." The "six nines" of reliability in the core and distribution portions are due to the redundancy that permeates both areas in this network. For many organizations, a lower rating is acceptable in the access layer because it affects fewer users.

When a company collects enough of its own data, it can replace the theoreticals with actuals. Thompson likes to do this because in his experience "the actual numbers on device availability are often better than the Telcordia tables."

In another network, perhaps operated by a smaller enterprise, the availability numbers in the core and distribution segments might not reach even 99.9 percent. They can show the company where improvements should be made to create a high availability network.

This theoretical availability assessment is important, as it can point the way early to areas that need work. In addition, planners can use it to determine if the network, as functioning or with

Cisco IT Achieves High Availability in San Jose Campus LAN

Spanning 50 buildings and serving desktops, data centers, labs, and manufacturing, Cisco's own network in San Jose, California consists of 900 switches, 200 routers, 250 console servers, 800 Cisco Aironet® access points, and an assortment of content switching devices. Despite its size and complexity, in 2002 the San Jose network approached 99.999 percent availability in areas where the network is 100 percent UPS- and generator-backed. In the second calendar quarter of this year, Cisco's San Jose data centers achieved 99.9986 percent uptime.

How did the Cisco IT San Jose team get so close to five nines—the Holy Grail of availability? Through a careful mix of planning, monitoring, measuring, and implementing processes that enable network uptime.

The team laid the foundation for high availability by determining which devices need to be monitored and how they should be measured. "Availability measurements are not only a measure of service," says Cisco Network Engineer Darrell Root, CCIE® No. 8302. "They are tools that enable you to improve service." Cisco measures the availability of all network devices and analyzes the reports not only to calculate network uptime, but also to prevent future problems.

Equally important is identifying the type of outage. The Cisco IT San Jose team calculates availability from ping statistics, with planned outages separated from unplanned ones. Because users are notified of planned outages, the impact to them is far lower than with an unplanned outage and, so, planned outages are not included in availability measurements.

The Cisco IT San Jose team recommends the following steps to achieve extremely high availability:

- Measure availability and identify the largest outages daily, monthly, and quarterly.
- Ensure that your network is fundamentally stable through implementation of a physical and logical hierarchy.
- Use the outage cause analysis report to prevent similar outages.
- Provision and audit both network and power redundancy.
- Institute a priority notification and management alert process.
- Tie availability measurements to the change management process.
- Build an out-of-band management network.
- Periodically audit router and switch configurations.
- Separate production and alpha networks.
- Monitor and manage which IOS and Catalyst OS (CatOS) releases are deployed.

For more on how Cisco IT San Jose achieves high availability, see the white paper at cisco.com/packet/153_17a1.

planned upgrades, can meet the needs of the business. And the path availability estimates can help predict the cost of downtime and the return on investment likely from improvements to the network.

What's Wrong with that Gateway in Seattle?

A typical record of device availability compiled by the analysis software from one or more of the basic metrics might be a spreadsheet listing all the types of devices down the left and, across the top, items including number of devices; total device outages; total outage duration; shortest, average, and longest duration; and device availability uptime and downtime.

From a spreadsheet, the network manager or consultant could learn, for example, that in a given week, the network's remote broadband access devices were available only 99.738 percent of the time, and branch office routers only 99.842 percent, both missing their targets of 99.9 percent. Four outages among the 20 broadband devices comprised four hours and 20 minutes in total outages, the longest one hour and 34 minutes. In addition, the network's 2700-plus branch office routers recorded over 1700 outages, one lasting more than 18 minutes.

To help locate the problem devices, the analysis software has also produced records on each device category by geographical area. One area alone shows 519 outages from 692 devices. So, there begins an investigation. Because only some of the branch office equipment is under control of the enterprise, it can work on those. For other equipment, it can query the service provider who is managing the routers. Similarly, the geographic reports can help identify problem broadband access devices.

The reports also show that the company's Internet connection devices were out of service four times for a total of 25 minutes 19 seconds. The shortest, average, and longest durations were, respectively, 12 seconds, 6 minutes 19 seconds, and 10 minutes 24 seconds. Total uptime was 99.937 percent, and downtime was 0.0627 percent.

> The crucial concept to remember is that availability must be defined for the network as a whole.

This level of availability might fall within the company's standards, but, Thompson says, he would dig deeper to find out about the 10-minute-plus downtime. He might find, for instance, that the 10-plus-minute downtime shows up on the record for Cleveland, Ohio, which has two Internet connectors. One or both might be the culprit, and testing can reveal which.

This Internet connection example gives a view of availability important to network managers. While the network met its standards for availability of this type of Internet connection, one city experienced a problem that could have affected users, if each device operated independently and both were in service simultaneously. The long outage time was a clear signal to the manager to look further to prevent future outages. If the devices were operating redundantly, users were probably unaffected, but the network manager still needs to investigate.

Network managers can drill down still further to detailed reports on specific devices. For example, a certain type of modem can be associated with an availability percentage of 99.9837 and an unavailability percentage of 0.01628; annual downtime of 85.7 minutes; an MTBF of 38,830 hours; and an MTTR of six hours. The latter is usually set by the enteprise, depending on service-level agreements in its contracts with service providers and its maintenance agreements with the device vendor. Managers can also get reports on groups of devices, for example, those in a particular type of node, to ascertain projected or actual availability.

An enterprise might find that the availability rating on its wireless switches was only 99.22 percent, with 57 devices accumulating 74 hours and 16 minutes of downtime. That might be enough to trigger user complaints. The enterprise might also learn that its 1230 hubs logged 105 hours and 29 minutes of downtime among 200 outages, for an availability rating of 99.94 percent. Again, while within the company standards for availability, the large number of outages should probably raise concern.

The Long View

The reports currently generated by enterprises cover either a week or month—rarely, as long as three months. But, Thompson advises, "You need to trend for at least three months and preferably six months or a year." For some clients, Thompson and the network managers he works with accomplish this manually, by spreading out the reports and eyeballing for trends. Does a router in Phoenix show up repeatedly as being out? Does a class of gateway continually have problems? In time, Thompson hopefully foresees that companies will load the data into a relational database, which can perform the trending automatically.

While that might sound like a costly proposition, particularly for small companies, Thompson suggests that much of an availability assessment does not have to be pricey. There is plenty of documentation available on Web sites such as Cisco.com, for example. And because you are probably already using trouble tickets, ping, and SNMP testing, you might already haive some or all of the fault management software you need, if not the analysis software.

And, he adds, there's a good deal of low-hanging fruit to pick from in improving availability. "In one network, we identified a particular office that showed a number of outages over a three-month period. The problem extended across types of devices and took a while to figure out, but it turned out to be an unreliable power system, which was a relatively inexpensive fix," says Thompson.

Achieving acceptable availability, Thompson adds, "requires intent and thought more than money."

Part Two: GETTING TO THE AVAILABILITY YOU NEED

NOW IT'S TIME TO ACT—THE MOST IMPORTANT part of defining and measuring acceptable network availability. After you've baselined availability and identified problem areas, you should assess four areas to pinpoint and plan remedies: *network design*, *infrastructure*, *operations and maintenance*, and *support*.

"Ninety-five percent of the improvements we suggest are small," says Thompson. "They're generally incremental, but they can add up to a full percentage point of improvement, which is major."

Network Design

The most important aspect of network design, Thompson says, is a rigid hierarchy of core, distribution, and access. "You want the highest availability and resilience in the core—that's where you need five nines," he says. "You may be able to go down a nine in the distribution layer, and perhaps you only need three nines in your access layer." This approach tracks the economics, as you have the fewest devices in the core and the most at the edge, perhaps orders of magnitude more.

Within core, distribution, and access, make sure you understand and maximize the resilience your network incorporates at the device, link, protocol, and application levels. *Device-level resilience* can be provided by redundant processors and uninterruptible power, as well as technologies rooted in Cisco IOS Software, such as Cisco Nonstop Forwarding (NSF) with Stateful Switchover (SSO). *Link-level resilience* can be provided by using technologies such as Redundant Packet Ring, Multi-Link Point-to-Point Protocol (PPP), EtherChannel®, and spanning-tree enhancements such as IEEE 802.1w and 802.1s. *Protocol-level resilience* can be achieved by implementing features such as Cisco Gateway Load Balancing Protocol (GLBP) for first hop redundancy, Cisco NSF/SSO for the routing and connectivity protocols, and routing protocol convergence optimizations. *Application-level resilience* is achieved by applying technologies such as stateful Network Address Translation (NAT) and server load balancing.

Then there are overall design principles such as avoiding single points of failure wherever possible—for example, having only one logical or physical circuit out of a router. "When that link is down, the router is out of business," points out Larry McNiff, high availability consultant in the Advanced Services Group at Cisco. "You want redundant paths to every device. If you have a single point of failure in your core, it's only a matter of time before it comes back to bite you."

McNiff cites one example where network design helped improve availability: a financial services company that needed a better network to support its 15,000 automated teller machines. It had been using an ATM enhanced LAN for the remote machines, but provisioning was difficult as each circuit had to be set up manually. Because the transport equipment was nearing end of life, the fix was to install an Ethernet network built on virtual LANs with Cisco Catalyst® 6500 and 4500 series switches. Switch software automatically provisioned the spanning-tree links. So, provisioning was much less error-prone, availability higher.

Finally, networks designed for high availability should also incorporate some form of out-of-band management to devices in critical areas. Often, a separate management network is implemented that allows operations personnel to reach critical devices via the console or an alternate port. This approach is especially useful for diagnosis and remedial action when the primary pathway is not available, and can significantly aid restoration effectiveness.

Infrastructure

A network design is only as good as the systems it comprises. In the physical infrastructure, both hardware and software play key roles in achieving availability.

Core systems should have full redundancy built in. In a core switch or router, for example, the switching fabric and Supervisor Engines should be designed with zero-loss hot sparing. Interface port modules, multiservice blades, and other components should also have internal spares; these need not be hot but should be capable of being switched into service within seconds. Core systems should also have sufficient throughput and be scalable so they can accommodate spikes in demand, particularly as next-generation IP-based services such as videophone conferencing and public wireless access are added to the network. And, of course, every vital system needs a UPS.

Core, distribution, and access systems can benefit from software features such as the aforementioned EtherChannel, GLBP, and Cisco NSF/SSO—all innovations patented by Cisco. For more on NSF/SSO, see "Speedy Recovery," *Packet*® Second Quarter 2003.

"NSF/SSO technology is especially useful for protecting the network edge devices. Traditionally, core routers protect against network faults using router redundancy and mesh connections that allow traffic to bypass failed network elements," explains Rohit Shrivastava, product manager of high availability technologies in Cisco's Internet Technologies Division. "NSF/SSO provides protection for network edge devices with dual RPs [route processors] that represent a single point of failure in the network design, and where an outage might result in loss of service for customers."

GLBP is a Cisco innovation that uses available network bandwidth effectively by load sharing traffic across up to four routers while protecting the first hop gateway. It enhances availability by setting up multiple redundant paths through a single virtual gateway address.

Easy link provisioning and management should also be part of the network infrastructure. For example, Cisco's patented Portfast, Uplinkfast, and Backbonefast assist in Spanning Tree Protocol (STP) reconvergence. STP prevents loops within a Layer 2 domain. Portfast allows STP to skip the first two of the four phases of its normal procedure in troubled access ports that are linked end to end, shortening time to recovery. Uplinkfast does the same for uplinks. Backbonefast shortens recovery time during an indirect failure—on a link connected to another switch—by allowing the local switch to begin transitioning immediately, instead of waiting for STP timers to expire.

Special REPORT

After you've decided upon infrastructure systems, Thompson recommends using a cookie-cutter approach: "Every branch office, every data center, every network node that performs the same functions should have identical hardware and software. This way, you can make fixes faster and more accurately, cut down on your spares inventory, and compare apples to apples when monitoring your network."

Operations and Maintenance

Once you know how your network is laid out and baselined, you can tweak operations for maximum efficiency. In many cases, you might not need new hardware, just operational changes. "You can also spot anomalies in your network," says McNiff. "Large networks have to be run by exceptions. Otherwise the amount of data would be overwhelming."

One company, for example, found that cutting the size of its routing tables improved availability. As Thompson explains, "They started with 1500 separate route descriptions. They were able to pare that down to 200, easing the workload on the router's CPU significantly while improving troubleshooting response times."

Efficient operations depend on rigorous adherence to best practices as well as network maintenance, which, again, requires monitoring to catch problems before they affect users. Performance analysis and metrics are done through tools such as Cisco Info Center's Internet Service Monitor (ISM) module and Cisco Service Assurance Agent.

It is also smart to map your maintenance contracts to your business and availability goals. A company Thompson worked with had high standards, which they rarely met, for MTTR. The problem was not with the skills of their personnel, but rather with their spares contract. Under the contract, spares would not be delivered until after their MTTR target had elapsed.

"If you target MTTR at one hour, you don't want a four-hour sparing window," he says. "And if your MTTR is 12 hours, you don't need guaranteed delivery in six. You don't need those inventory costs."

What About High Availability for Service Providers?

High network availability is, of course, tantamount to meeting the customers' expectation of getting service with no disruption. But how does the concept of *service availability* factor into the mix? In the case of cable operators who are increasingly deploying IP-based voice telephony services, it ranks in importance equal to that of network availability. Find out more about service availability and how to measure it in the Fourth Quarter 2003 issue of *Packet*, coming in October.

You should also be looking at technical details such as the size of routing tables. Are your MTTR goals what you want? Should you be getting to solutions faster? Thompson offers a final two important words on operations and maintenance: *change management*. Are modifications performed according to stringent standards? Is there an active planning process that includes non-IT personnel in the company?

Support

The final component in improving availability is to evaluate your support structures to make sure they match your needs. Know the capabilities of your own personnel and of your vendor's, and how they fit the needs of your network. Do your people have the skill sets they need? Does your vendor have the tools and the support service people you need? Can they provide the training that your staff needs?

The support available from your vendor can also affect your availability. For example, the Cisco Technical Assistance Center (TAC) Web site (cisco.com/tac) enables users to solve many problems immediately themselves, without even talking with a support person.

"Seventy-eight percent of the queries entered into the Web site are resolved before they get as far as being logged in as cases," says Sameer Khera, director of technical support for the Cisco TAC Web site. The site has huge reserves of knowledge for self-service. Knowledge is delivered through documents and state-of-the-art tools such as Output Interpreter and TAC case collection, a searchable knowledge base of actual Cisco customer cases.

If you can't resolve the problem online, the Cisco TAC Web site is backed up by engineers available by e-mail and phone, and by increasingly higher levels of support up to the likes of McNiff and Thompson, who are both members of Cisco's Network Availability Improvement Support group. They consult with Cisco customers, large and small, and are able to produce intensive analyses of network availability and recommendations for improving it in line with your business goals.

Who knows? After you have identified and implemented improvements to your own network's availability, you might even discover one or more new business opportunities. Thompson recalls a provider whose direct-connect customers were being served by nonredundant links, with an availability rating of 99.95 percent.

◆ ◆ ◆

The relationship between availability baselines and corporate bottom lines is not surprising. Networks are rapidly increasing in complexity and criticality to an organization's success. And as voice and video services are folded into existing data networks, this relationship will become even more profound. All this points to the need for knowing how your network is performing and supporting organizational goals. Armed with metrics such as network availability, decision-making is improved and allocation of resources more precise.

Many enterprise managers and executives are beginning to understand that network availability management is fundamental to successful network operations—just as the accounting department is integral to controlling costs. ▲▲

Enterprise SOLUTIONS

SIMPLIFIED MANAGEMENT: The Waterfront data center is home to the first phase of Deakin University's far-reaching storage-area network solution.

COURTESY OF DEAKIN UNIVERSITY

Storage Strategy

An Australian university efficiently and reliably manages terabytes of information with Cisco storage networking.

BY CAROLE WARNER REECE

WHERE MIGHT YOU BE IF YOU NEEDED to provide 60,000 users spread across five campuses and hundreds of kilometers with access to 35 terabytes of managed storage? Try Australia, where the Information Technology Services (ITS) Division of Deakin University (deakin.edu.au) recently implemented a Cisco storage-area network (SAN) that provides the university with a high-performance storage solution while offering features such as security, scalability, and ease of use.

Deakin University is a government-funded higher education and research institution in the State of Victoria, Australia. Its multiple campuses are located from Victoria's capital in Melbourne to the town of Warrnambool, more than 300 kilometers to the southwest.

In addition to traditional campus-based teaching at these distant locations, the university provides flexible learning options through innovative teaching and online delivery methods that include computer-aided learning, simulations, and videos.

Deakin has a history of technology achievements—it is the first Australian university to embrace UNIX for enterprise applications and the first Cisco customer in Australia to deploy Cisco IP telephony in production.

According to Craig Warren, ITS desktop and network services manager at Deakin, the division faces the challenge of managing the voice and data services that support the entire university community.

"This involves providing IT services for as many people as possible at the university by establishing and maintaining a stable, highly available, resilient infrastructure, and then layering applications and services on this infrastructure," says Warren.

Deakin's underlying IT strategy is to invest in the future and actively embrace innovation and change.

The IT group supports this strategy with a technology refresh and rollover program that involves upgrading systems every few years, says Warren.

Unconsolidated Storage

In 2002, the major applications in the Deakin IT environment included a learning management system, a collaborative content management system, a finance system, and student record systems. ITS also saw an increased use of low-cost redundant clusters of Linux servers for lower total cost of operation for e-mail, file, print, and Web services. At the same time, Deakin had experienced exponential growth of storage requirements for the past several years. Storage needs included student records and correspondence, e-mail, human resource information, imaging and graphics, and special applications used by the various academic departments.

The fundamental storage problem was how to efficiently and reliably manage terabytes of information connected to a multitude of servers. While direct-attached storage devices worked well for a few servers, projected exponential storage growth dictated a need to decrease the complexity of managing and supporting terabytes of separate storage devices. Deakin looked for a reliable, high-performance storage solution to simplify and centralize management, consolidate resources, and support the many applications on UNIX and Linux servers. After reviewing the technology available, the ITS division determined that implementing a SAN would increase the accessibility of data across the university and increase the efficiency of managing that storage.

Until January 2003, Deakin did not have a SAN, but rather a mix of direct-attached storage (DAS), most of which was Fibre Channel Arbitrated Loop (FC-AL) attached and some of which was SCSI. According to Warren, "We decided to build a SAN as a foundation infrastructure that would cater for growth and flexibility, and support our UNIX and IP clusters as well as IP attachment through iSCSI [SCSI over IP] and FCIP [Fibre Channel over IP]."

Deakin's strategy was to consolidate storage onto a Cisco MDS-9000 Series-based SAN with storage servers from IBM. "We chose the Cisco MDS platform for our SAN for several reasons. Cisco was already a strategic partner for networking with Deakin, as well as a partner with IBM. The Cisco solution provided significant features such as iSCSI, FCIP, and VSANs [virtual SANs], as well as extensive diagnostics and management tools. Because storage area networking was viewed as another type of network for ITS to support, Deakin appreciated the FC TraceRoute, FC Analyzer management tools as well as the CiscoWorks integration. "We also thought that the MDS-9509 platform's high performance and high density would provide us with growth and flexibility in the future," says Warren. The MDS storage networking technologies support high-speed Fibre Channel, as well as IP connectivity through iSCSI for individual hosts, and FCIP for interconnecting SANs beyond Fibre Channel distance capabilities. For example, ITS plans to extend the SAN to support geographically diverse secondary storage for disaster recovery in a later implementation phase.

STREAMLINED SOLUTION: Most of Deakin's storage devices are housed in the Waterfront data center where the server infrastructure was consolidated from multiple hosts to a single server.

DEAKIN UNIVERSITY–WATERFRONT DATA CENTER

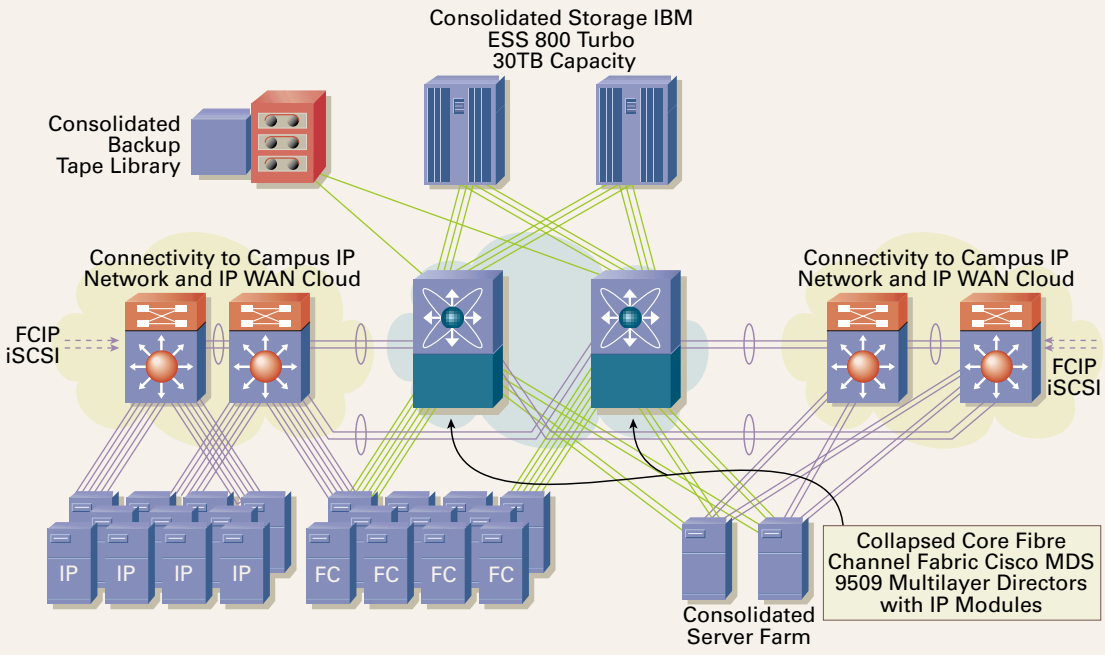

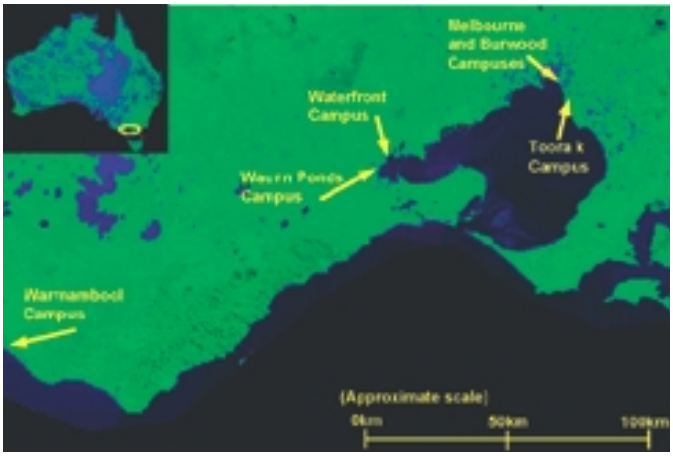

COURTESY OF DEAKIN UNIVERSITY

Implementation Strategy

Deakin implemented its SAN in phases, starting with the Waterfront data center, which supported the greatest number of storage devices. Storage exists to support applications, and the Deakin ITS group consolidated its server infrastructure as part of the migration from DAS to SAN, migrating from multiple hosts at the Waterfront campus to a single server running separate domains. During the implementation, they made greater use of servers running Linux for applications such as e-mail, file, printing, and Web services, using clustering to provide high service availability.

In just three months, Deakin completed migration of more than 90 hosts in the Waterfront data center over to the SAN, one host at a time. The cutover was an ongoing process, with servers migrated over from loop-attach or DAS to the SAN one by one during change-control windows at nonpeak hours. Migrating a server typically involved shutting the machine down briefly to upgrade drivers and firmware; for example, installing newer host bus adapter (HBA) drivers or installing a second HBA for multipathing. Each server then was dually connected to the SAN on one FC port, and to its existing loop or DAS storage on the other port. A mirror was created on the SAN storage of the loop-attached or DAS volumes. The two storage images were run in parallel for a time, and then the mirror was split with the SAN storage becoming the primary data store. After the primary storage on the SAN proved successful, the existing FC port connecting to the old loop or DAS storage was disconnected, and then connected into the SAN, which allowed for multipathing on the SAN.

Phase One Implementation Status

Migration and cutover were remarkably smooth, according to Warren. "The whole configuration of the SAN was simple. We connected it to servers and it worked. We kept testing to try and find something that was not working, but we didn't."

For the current SAN infrastructure, Deakin uses two Cisco MDS-9509's in its production network and one Cisco MDS-9216 in its development and test environment at the Waterfront campus. These devices support data center host-to-IBM ESS connectivity over Fibre Channel without using the campus LAN. The SAN does have connectivity to the LAN and WAN. The ITS group uses Veritas multipathing technology to provide high availability and load balancing to the multiple paths to consolidated storage on the SAN.

To create multiple logical SANs or virtual SANs (VSANs) within the same physical infrastructure, Deakin uses Cisco VSAN technology. From the IBM ESS, five ports connect to the production VSAN on the Cisco MDS-9509's, and one port connects to a development and test VSAN on each Cisco MDS-9509. For security purposes, a third VSAN is configured for unused ports on the Cisco MDS-9509's. This third VSAN isolates new devices connected to the Cisco MDS from appearing on the production SAN. Zoning, which is used within each VSAN, ensures device security.

Consolidation Results

With the first phase of Deakin's SAN successfully implemented, Warren sees that the key benefits of the

Continued on page 89

FURTHER READING

- **Article on "The Scalable, Manageable SAN" (*Packet*, Fourth Quarter 2002): cisco.com/packet/153_7a1**
- **Cisco storage networking white papers: cisco.com/packet/153_7a2**
- **Cisco storage networking product literature: cisco.com/packet/153_7a3**

Industrial-Strength LAN

Cisco Catalyst 2955 Series Switch takes intelligent Ethernet to harsh environments.

BY GAIL MEREDITH

DO YOU NEED A TOUGH NETWORK? IN places where conventional gear cannot survive, outside the traditional office wiring closet or service provider's central office, environmentally "hardened" equipment can withstand temperature and moisture extremes, excessive dust, and vibrations. These environments, such as a power substation or manufacturing plant, are usually mission critical and security sensitive. Until now, manufacturers, utilities, and transportation departments have survived by providing special enclosures to isolate commercial equipment from potentially damaging environmental conditions, but this option can be costly and complex to manage. The prevalence of Ethernet in these environments has been increasing at double-digit rates annually as many industries move from proprietary networking systems to Ethernet and IP-based networks. In many instances, this migration frees users from dealing with comparatively expensive networks that require specialized training and have a limited number of vendors and connectivity.

As the *de facto* LAN standard, Ethernet is widely understood and inexpensive, and offers enticing advantages over proprietary technologies that reduce total costs and enable new functionality for industrial networks. Ethernet supplies the most cost-effective data transmission and plentiful bandwidth, enabling advanced new applications and the potential for real-time data monitoring. Vendors of automation equipment used in manufacturing processes, power generation and transmission, water distribution and treatment, and many other applications that have not traditionally used Ethernet are adding Ethernet interfaces alongside traditional ones, making end-to-end Ethernet feasible in highly complex environments and replacing dozens of side-by-side proprietary network-

RUGGED AND INTELLIGENT: Manufacturing networks are easier to deploy, operate, and manage by migrating to Ethernet—enabled by Cisco Catalyst 2955 switches—from dozens of proprietary technologies.

CISCO ETHERNET IN THE FACTORY

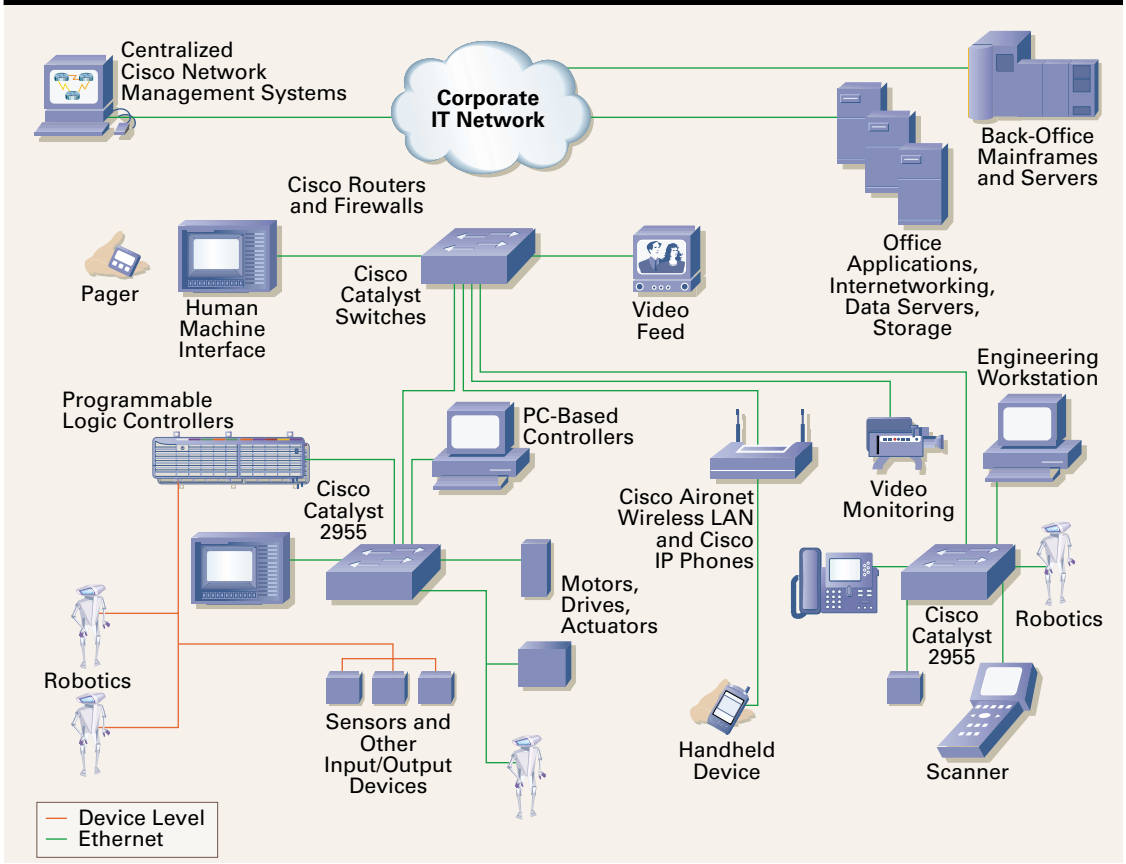

ing technologies. Consolidating everything into one network is much less expensive to deploy and operate and far easier to manage than multiple systems.

Intelligent services such as security, high availability, and quality of service (QoS) are vital to industrial networks. "Intelligent Ethernet is highly critical in industrial deployments and even more relevant than in enterprise networks," says Manrique Brenes, product manager in the Ethernet Access Group at Cisco. "There are two major issues at work: operator safety and plant shutdown. When you shut down a facility, the cost is immense."

The new Cisco Catalyst® 2955 Series Switch is built for industrial networks in harsh environments (see figure). With industrial compliance ratings, no fans to reduce the potential impact of dust or other suspended particles, and a design to withstand extended temperature, moisture, and vibration, these switches are optimized for manufacturing plants, transportation networks, and utilities. Available in three configurations supporting 12 fixed ports each, the Catalyst 2955 Series has a small form factor for housing in traditional industrial enclosures. "Much more compact than the traditional 19-inch rack, the switch's form factor is similar to that of a PLC [programmable logic controller] and other devices deployed in this environment," says Brenes. It is mountable with DIN rails or 19-inch rack (with extenders).

The switches can withstand up to 50 g-forces of trapezoidal shock, endure temperature extremes from -40 degrees Celsius (-40 degrees Fahrenheit) to 60 degrees C (140 degrees Fahrenheit), and tolerate relative humidity extremes from 10 to 95 percent. Included is the full suite of Cisco IOS® Software and environmental Management Information Bases (MIBs) to monitor temperature changes and trigger alarms for conditions such as overheating or failure of one of the redundant power sources.

Cisco IOS security features help protect the network from unauthorized users or applications, and QoS mechanisms enable bandwidth allocation and proper queuing for mission-critical or time-sensitive application traffic critical for manufacturing alongside traffic such as file transfers or Web browsing. This intelligence, combined with greater bandwidth, opens the door to new applications and more effective communications in manufacturing, transportation, and utility networks.

The switch contains many redundancy and high availability features to assure maximum business resilience. However, should a failure call for a controlled shutdown, the switch has multiple alarm mechanisms that can be used to safely stop the automation equipment. Two alarm relays let operators distinguish between critical and noncritical alarms. "They act as a mental bridge between the industrial and enterprise ways of doing things," explains Brenes. "People who are not familiar with network alarm protocols can be alerted of critical events in a manner familiar to them."

Manufacturing

The issues of safety and reliability are vital to manufacturing. A traditional factory floor can include dozens of network technologies designed for highly reliable control of certain machines and robots. Both industry and tool vendors realize the value of offering Ethernet interfaces, and migration toward hardened Ethernet connectivity is under way. Managing one network instead of dozens saves both time and money.

Plentiful switched bandwidth of 10 to 100 Mbit/s per interface replaces legacy shared speeds of 500 Kbit/s or less (shared by multiple users), enabling new applications and services on the factory floor. Ethernet-enabled devices can send usage data directly to management applications, allowing real-time assessment of network status. What's more, manufacturing facilities can leverage Ethernet to interconnect with research and development teams in other locations without compromising security. The enhanced communication capabilities of Ethernet allow engineers to address and solve issues quickly. Companies can offer multiple services to factory workers (voice, video, historic data, access to electronic CAD documents, for example), enabling them to communicate with one another and with management more quickly and effectively, leading to faster improvements and problem solving.

Intelligent switching services increase security and control of access to network resources. The Catalyst 2955 access control parameters allow for filtering and traffic prioritization based on MAC address, IP address, or TCP/UDP—for very granular control of every device. This capability allows for secure remote access into devices for both management and data collection for planning and performance evaluations. Internet Group Management Protocol (IGMP) snooping capabilities in the Catalyst 2955 play a key role in many industrial applications, which are based on multicast models to deliver data. This feature allows the different devices to communicate with each other or bypass uninvolved

Increase Your Internet Quotient

For more on how manufacturers are streamlining their factory operations by standardizing on Ethernet and IP technologies—enabling such applications as real-time, two-way communications into and out of the shop floor and reducing reconfiguration costs—see "Beyond the Assembly Line" in the July/August 2003 issue of *iQ Magazine* at cisco.com/go/iqmagazine.

devices—which is especially valuable for network elements with limited processing capacity (sensors, controllers, monitoring interfaces, and the like), avoiding crippling data overload.

Sasol Synfuels (www.sasol.com), a global petrochemical company headquartered in South Africa, is deploying a Cisco-based intelligent Ethernet solution that the company believes will save them millions over the next three years. "We wanted to improve efficiency to the whole plant and deploy an open technology on which we could continue to build out the network over the coming years," says Johan La Grange, executive for information at Sasol Synfuels. "Cisco understands the needs of networking in a complex environment."

Before migrating to IP, Sasol had 11 separate control rooms in its plant, which ran on private automatic branch exchange (PABX) copper cabling. There was no single way the plant and control rooms could link up to exchange mission-critical process parameters; communication was accomplished by walkie-talkies and telephone. Using Cisco-based intelligent Ethernet, Sasol is now able to communicate efficiently across the plant and control rooms, driving huge cost savings. Sasol has also deployed Catalyst 2948G switches in each control room, which help to automate plant processes such as temperature and valve control.

Transportation

There are two sectors in the transportation industry. The public sector includes airports, highways, bridges, mass transit, and seaports. The private sector includes airlines, railroad companies, trucking, rental cars, and shipping. "There are harsh network environments throughout the transportation industry," says Larry O'Connell, product manager in the Ethernet Access Group at Cisco. "Ethernet offers sorely needed capabilities through greater bandwidth and intelligent services such as security and QoS."

Catalyst 2955 Series switches can support an intelligent transportation system throughout the public sector. For example, a switch at a traffic intersection can gather data from a video camera, temperature sensor, and traffic signal controller, then send real-time information and health diagnostics to centralized management stations. Government agencies can monitor traffic conditions via IP video and make information available to emergency response and law enforcement personnel, and to the public via radio, television, and dedicated telephone-line information.

Another example of Ethernet being deployed for transportation applications is IdleAire Technologies Corporation (www.idleaire.com), which offers services that solve pollution and cost issues in the trucking industry. Diesel rigs congregate at travel centers to rest. An IdleAire hookup keeps engines warm and cabs cool, so rigs don't have to burn diesel fuel. With the Cisco Catalyst 2955 Switch, IdleAire has the flexibility to extend its converged network services, such as IP telephony and Internet access, to truckers in the convenience of their own cab. IdleAire plans to connect 270,000 truck stop parking spaces throughout the US. "We provide truckers with services that they have never had in this fashion before. They can access the Internet and use the phone from the privacy of their cab. They are no longer standing at a pay phone with four or five other people standing around listening to their conversations," says Jon Duren, chief technical officer at IdleAire Technologies.

Utilities

A third area where Catalyst 2955 Series switches are making a difference is utilities—gas, electricity, and water. Both security and reliability are important, along with remote manageability, as many stations are unmanned across networks that span hundreds or thousands of miles. Most utility systems use large control systems known as SCADA (for Supervisory Control and Data Acquisition), which gathers data from control points such as electrical substations, pumping stations, or gas meters to monitor and control mechanisms such as valve relays and pumps.

Ethernet appeals to utilities because it provides more bandwidth at a lower cost than Frame Relay, serial lines, or telephone/modem connections. Extra bandwidth allows utilities companies to gather real-time diagnostics and data. At the environmentally harsh endpoints, Cisco Catalyst 2955 switches communicate with control stations via metro Ethernet networks.

Emerson Process Management Power & Water Solutions (emersonprocess-powerwater.com), for example, has developed Ethernet-based network solutions for its own Ovation control system, which is primarily used in water/wastewater and power generation applications. "The Catalyst 2955 Switch is well suited for these mission-critical types of applications," says Carl Staab, manager of communication technologies at Emerson. "It's almost as if Cisco were looking over my shoulder when they were designing it."

A gas and electric utility based in the US midwest will deploy Ethernet to its substations throughout the next two years, replacing analog serial and telephone line/modem solutions. Among the advantages the utility company expects is the ability to look beyond the remote terminal unit into devices within each substation, a level of control the utility does not have today. ▲▲

FURTHER READING

- **Cisco Catalyst 2955 Series switches: cisco.com/go/2955**

Benefits Stack Up

Cisco StackWise technology innovations define the next generation of stackable switching.

BY DAVID BARRY

THE COMBINATION OF EVER-INCREASING desktop computing power, new bandwidth-intensive applications, and the presence of multiple device types, such as IP phones and wireless LAN access points, is creating increased bandwidth demands on business LANs. In addition, many newer applications, such as voice and video over IP, and a heightened need for security are raising demand for greater intelligence in desktop switches. As a result, IT professionals increasingly view the edge of the network—the wiring closet—as a key point for adding new functionality such as capacity increases, Layer 3 intelligence, and IPv6 readiness. With higher-layer features, network managers can deploy consistent, network-wide intelligent services from the desktop to the core and through the WAN.

With the new Cisco Catalyst® 3750 Series switches, Cisco has markedly raised the bar in the level of performance, intelligence, and automation available in desktop switching. Included for the first time with the new Catalyst 3750 switches is Cisco StackWise™ technology, an innovative 32-Gbit/s-stack interconnect fabric that creates a unified, logical switching architecture by linking multiple fixed-configuration switches. Up to nine stackable switches can be combined into a single unit via special stack interconnect cables that create a 32-Gbit/s bidirectional closed-loop path.

"The stacking capabilities of the 3750, combined with the low price point of the switches, will allow us to make incremental, yet robust, improvements to our network infrastructure," says Jon Snyder, computing and network services at Portland State University, Oregon. "The hardware has room to grow, and we are always looking for a platform that will continue to provide value for a number of years."

In its April 23, 2003 *Competitive Intelligence Report*, Current Analysis states: "Cisco is currently the only vendor that can offer a Gigabit Ethernet stackable, thanks to its 32-Gbit/s stacking backplane. Cisco can offer Gigabit Ethernet technology not available in other stackables and can deliver this technology without compromising performance across the stack. The functionality Cisco is extending into the stackable space is a measuring rod by which other products will be gauged."

One of the most profound benefits of Cisco StackWise technology is the automation of switch management. With a unified, logical switching architecture, a stack of Cisco Catalyst 3750 switches presents a single entity to the network and is managed as a single object. One IP address. One configuration file. One spanning-tree node. One routing hop. This single management concept also applies to activities such as fault detection, virtual LAN (VLAN) creation and modification, security, and quality-of-service (QoS) controls. Each stack has only one configuration file, which is distributed and synchronized to each member in the stack, allowing the switches to share the same network topology, IP address, and routing information.

"We have been testing the Cisco Catalyst 3750 in one of our schools and are very pleased with the performance of these switches. We upgraded the school's backbone to 1000 MB and are using 100-MB connections to the student stations," says Joe Stout, IT manager for the Park City School District in Park City, Utah. "Managing all the switches in the stack as a single IP address has made this upgrade much easier to manage. We hope to make this our standard configuration throughout the [school] district."

"Because the 3750 presents one instance to the network, it dramatically changes how companies have traditionally managed stackable switches," says Scott Boynton, senior product marketing manager in the Desktop Switching Business Unit at Cisco. "A customer only has to Telnet to the master switch in the stack using one IP address. Any changes made to the configuration file or to the IOS® image are immediately propagated to all the other switches in the stack."

The stack behaves as a single switching unit that is managed by a master switch elected from one of the member switches. The master switch automatically creates and synchronizes all the switching and optional routing tables. A working stack can accept new members or delete old ones without interrupting stack performance. When a new switch is added, the master switch automatically configures the unit with

Ask your peers and Cisco experts questions or share your own knowledge about switching and related topics at the *Cisco Networking Professionals Connection* LAN, Switching, and Routing forum: cisco.com/discuss/lan.

STACK AUTOMATION

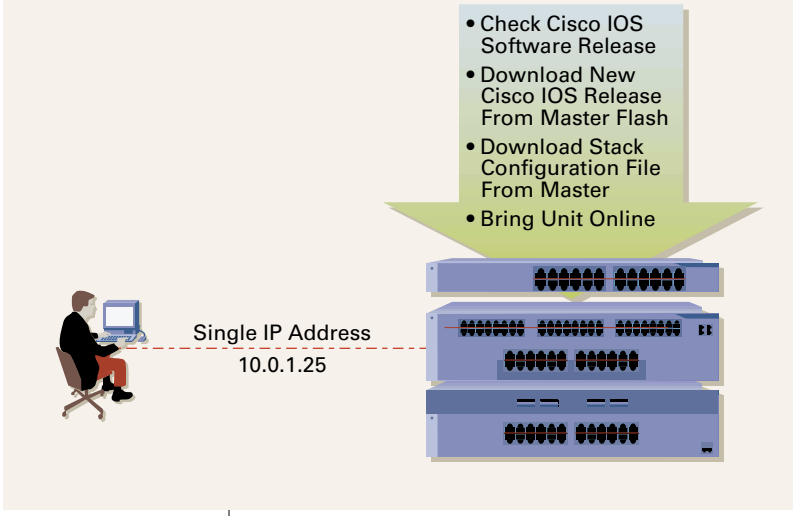

UNIFIED AND LOGICAL: Once plugged into a stack, the Cisco Catalyst 3750 Series Switch is automatically updated by the master switch and brought online. The entire stack has only one IP address, greatly simplifying management of the stack.

the currently running Cisco IOS Software image and configuration on the stack. The network manager does not have to do anything to bring up the switch before it is ready to operate. The stacking cables themselves were created from the results of numerous usability studies, says Boynton. These cables sport several innovations, including the placing of windows on the connectors that allow the installer to easily see where the cable should be plugged into the unit. The screws that tighten the cables have also been redesigned: offset and placed at different angles, making it easy for an installer's fingers to tighten or loosen screws—an often laborious task on most switches.

The Cisco StackWise technology interconnect is fruit of Cisco's ASIC design innovations. While many competitors continue to use off-the-shelf merchant components, Cisco has designed two new ASICs for the Catalyst 3750: the main switching engine and the main processor. They are the key to integration of IOS features in the platform's hardware and to high levels of stack integration. Cisco has also built support for IPv6 into the hardware that is essential to route IPv6 in Layer 3 mode. IPv6 includes a much larger number of available addresses and advanced features for mobility, which demands more from the switch. With support in the ASICs, network administrators will be able to deploy the IPv6 software when it is released in the near future, without compromising switch performance.

Optimized for Gigabit Ethernet, Availability

In addition to the 32-Gbit/s stack interconnect, the Catalyst 3750 provides JumboFrame support. With JumboFrame technology, the Catalyst 3750 can handle 9018 byte packets, much greater than the 1518 byte packet supported in standard Ethernet. This enhancement allows more payload in each packet and is ideal for passing traffic to and from servers. "While many PCs or workstations don't currently support Gigabit Ethernet speeds, it is not uncommon for these platforms to burst beyond Fast Ethernet speeds, to as high as 400 Mbit/s, especially when accessing data on servers," notes Boynton. "On a switch stack with only Fast Ethernet ports, this would throttle back traffic and slow response time. With just five minutes of increased productivity per employee a day, an upgrade to a Gigabit Ethernet switch can provide a return on investment in as little as three to four months."

Gigabit Ethernet ports on the Catalyst 3750 switches also provide companies with a smooth migration path as they evolve from Fast Ethernet to Gigabit Ethernet. As PCs and workstations with new architectures that support Gigabit performance are installed, IT professionals can seamlessly add additional switches with Gigabit Ethernet ports to an existing stack.

The Cisco StackWise interconnect technology is designed with two counter-rotating paths of 16 Gbit/s each. To efficiently load balance traffic, packets are allocated between these two logical counter-rotating paths to create the 32-Gbit/s interconnection. Because of these dual paths from any port to any port within the Catalyst 3750 stack, maximum uptime is ensured, as there is always an alternate path available if a failure occurs in either path. The Cisco Catalyst 3750 supports cross-stack EtherChannel® as well as Cross-Stack Uplink Fast (with subsecond failover), and cross-stack equal cost routes across different switches in the stack. These features augment high availability provided by Hot Standby Router Protocol (HSRP) for router failures. With EtherChannel, the loss of one switch will not affect connectivity for the other switches. And equal cost routes ensure that switches can support dual homing to different routes for redundancy. In the case of a master switch failure, another master switch takes over with minimal disruption within two to three seconds. Because every switch in the stack can serve as the master, a 1:N redundancy mechanism for Layer 3 is created—a capability unrivaled in the industry today.

Offering even higher availability, up to six members of a Cisco Catalyst 3750 stack can be connected through the Cisco Redundant Power System (RPS) 675. When the RPS 675 senses failure, it automatically delivers uninterrupted power to the device, alerts the network administrator of the failure so the device can be restored or replaced during non-interruptive maintenance, and automatically resets into standby mode when the failed unit is replaced with a new unit. ▲▲

FURTHER READING

- **Cisco Catalyst 3750 Series switches: cisco.com/go/catalyst3750**

Optical Advances

New multiservice transport platform capabilities for the Cisco ONS 15454 bring simplicity to metro DWDM deployments.

BY DAVID BARRY

Many new high-bandwidth services have gained in popularity on enterprise networks in the last few years, including Gigabit Ethernet and 10 Gigabit Ethernet, as well as storage protocols such as Fibre Channel, Fiber Connectivity (FICON), and Enterprise Connectivity (ESCON). But because enterprises are seeking to extend these protocols from their LANs across the WAN, pressure has been put on their transport networks or their managed service provider's metro networks that were not designed for these bandwidth-intensive services.

To provide capacity relief, enterprises and service providers have installed dense wavelength-division multiplexing (DWDM) platforms using technology that is typically derived from the long-haul network. Metro demands, however, are much different than those of the long haul because metro networks must deal with multiple services; they must be flexible to handle changing customer requirements; they must support different network configurations; and they must enable services to be turned up and managed with greater speed and efficiency.

In addition, most metro networks must support the traditional time-division multiplexing (TDM) and SONET/SDH services ranging from DS1/E1 to OC-192/STM-64, as well as Ethernet and IP, which are often handled on traditional SONET add/drop multiplexers or next-generation platforms such as the Cisco ONS 15454 Multiservice Provisioning Platform (MSPP).

Cisco has taken the next step in the evolution of the metro optical network with the release of the Cisco ONS 15454 Multiservice Transport Platform (MSTP), which integrates intelligent DWDM capabilities directly into the Cisco ONS 15454. The big advantage of the Cisco MSTP is that it not only delivers significant new capacity, it also greatly lowers the cost of using DWDM by integrating it into an existing proven MSPP platform with broad multiservice support and taking advantage of its network intelligence to simplify deployment and operation of the network.

"This level of integration allows service providers and enterprises to reduce the number of platforms needed to deliver voice, data, Ethernet, and storage services on the metro network," says Ron Johnson, product manager in the Optical Networking Group at Cisco.

Continues Johnson, "Having fewer platforms reduces both capital and operational expenditures because functionality that previously required additional platforms can now be handled by cards or blades in the MSTP. Operationally, management and provisioning are simplified because users only need to provide management expertise and skills for one platform. And when the MSTP is paired with Cisco multiservice provisioning platforms elsewhere on the metro network, providers gain the dramatic benefits of a common user experience and the ability to do end-to-end circuit provisioning."

The Cisco ONS 15454 MSTP takes advantage of the operational simplicity introduced in the MSPP, with user features such as multilayer graphical network, node and card visibility, A-to-Z network-based wavelength service provisioning, and graphical software wizards. These features help to simplify and speed user operations for such tasks as initial ring turn-up, service provisioning, and network node and bandwidth upgrades.

"The MSTP leverages this architecture to introduce a level of operational simplicity unheard of in metro DWDM networks," says Johnson.

DWDM Architecture

The Cisco ONS 15454 MSTP also brings great flexibility to DWDM network design with its "plug-and-play" card architecture. Among the many cards available for wavelength aggregation and transport are a 10-Gbit/s multirate transponder card for OC-192/STM-64, 10 Gigabit Ethernet LAN, and 10 Gigabit Ethernet WAN traffic; a 4x OC-48 10-Gbit/s muxponder card for aggregating four OC-48/STM-16 interfaces; and a 2.5-Gbit/s

multirate transponder card for transport of ESCON, Gigabit Ethernet, Fibre Channel, FICON, and video services such as D1 and HDTV. The Cisco ONS 15454 MSTP supports multiple networking architectures through its flexible node configurability, including terminal, optical add/drop, hub, and line amplifier.

Ask your peers and Cisco experts questions or share your own knowledge about DWDM, the Cisco ONS 15454, and related topics at the *Cisco Networking Professionals Connection* Optical Networking discussion: cisco.com/discuss/optical.

Wide Service Interface Mix

Because metro networks are closer to customer premises than their long-haul counterparts, they require the support for a greater diversity of service interfaces. With these interfaces, enterprises can natively transport a wide variety of services over a common transport network without unnecessary conversion stages and equipment, and providers can apply new tariffs for these services. A wide services mix also simplifies services planning.

Within a single platform type, the Cisco ONS 15454 supports a broad mix of standards-based services, including:

- Aggregated lower-rate TDM services from DS1/E1 over 2.5-Gbit/s and 10-Gbit/s wavelengths
- SONET/SDH wavelength and aggregated services: OC-3/STM-1, OC-12/STM-4, OC-48/STM-16, OC-192/STM-64
- Data services: Private-line, switched, and wavelength based including 10/100BASE-T, Gigabit Ethernet, 10 Gigabit Ethernet LAN Phy, and 10 Gigabit Ethernet WAN Phy
- Storage services: 1-Gbit/s and 2-Gbit/s Fibre Channel, FICON, and ESCON
- Video services: D1 and HDTV

Flexibility was one of the key reasons that the North Carolina Networking Initiative (NCNI) selected the Cisco ONS 15454 MSTP to expand its regional GigaPOP with optical DWDM.

The NCNI is a consortium of North Carolina research-intensive universities, including Duke University, North Carolina State University, and the University of North Carolina at Chapel Hill, and MCNC, a private, nonprofit organization in North Carolina that seeks to enhance economic development in the state.

"While the capacity of DWDM is important to us, flexibility was most important," says Mark Johnson, MCNC/NCNI director of Internet technologies. "We support different services among the many campuses, and we see great economic value by being able to support many different connections for both production and research purposes with as little duplicate infrastructure as possible."

"With the Cisco ONS 15454 MSTP, for example, we can use one interface to support many different signals," says Johnson. "If one of the university departments wants to experiment with digital video HDTV we can use the 2.5-Gbit/s multirate transponder card to run the HDTV signals. But then after a month or so when the experiment is completed, we can use that same interface to run Ethernet traffic or SONET or to try a new protocol we're developing called optical burst switching."

Adds Johnson, "This is a huge cost savings because I don't have to buy an interface card for each type of traffic. Not only are the cards costly, I have to maintain each card. Also, if I had to buy three or four cards, I'd have to maintain spares for each one, adding even more cost. It's not unreasonable to assume that without this capability we might have to cancel certain research projects because it would be too expensive to do it."

Automatic Optical Power Management

Many metro service providers face the challenge of managing metro DWDM network architectures that are typically ring topologies—open ring, multihub ring, and closed ring—with complex, dynamic add/drop traffic patterns. Traditional solutions cannot automatically manage DWDM analog variables such as optical noise, dispersion, and the dynamics of adding and dropping wavelengths.

The Cisco MSTP includes automatic optical power management software that can dynamically monitor and control optical power. The MSTP software reconstructs a model of the provisioned DWDM network, and software algorithms automatically provide networkwide optical power management. The software automatically equalizes channels that are intrinsically unequal, adjusts for optical paths with different insertion losses (add, drop, express and hitless paths), and maintains constant power when wavelengths are added or dropped.

"On legacy DWDM platforms, service providers literally need engineers with PhD's in optics in order to design a metro DWDM network," says Cisco's Ron Johnson. "And it typically is all done on day one, with very minor changes afterward. But this doesn't work in the metro with changing needs."

Adds Johnson, "With the MSTP, you don't have to know your traffic demands on day one and you don't need to have a PhD in optics. We've built intelligence into the MSTP." ▲▲

FURTHER READING

- **Cisco optical networking solutions for service providers: cisco.com/packet/153_8b1**
- **Cisco ONS 15454 Optical Transport Platform: cisco.com/packet/153_8b2**
- **Cisco metro optical strategy: cisco.com/packet/153_8b3**
- **Cisco ONS 15454 MSTP architecture: cisco.com/packet/153_8b4**

Cisco Mobile Exchange

Cisco framework helps mobile network operators combine benefits of mobile data services and public WLANs into a single offering.

BY STEVE HRATKO

Wireless LAN (WLAN) technology based on the IEEE 802.11 standard is no doubt one of the hottest areas in the networking industry. This technology is finding its way into most new laptop computers, personal digital assistants (PDAs), and even wireless phones. Already a staple for many businesses eager to leverage the productivity benefits of WLANs, this technology also continues to gain traction in the consumer/SOHO market where it easily solves the inside wiring challenge of bringing Ethernet to the home.

Now, 802.11-based technology is making its way into the public network to support WLAN hotspots. Typically found in airports, hotels, convention centers, and, more recently, commercial airplanes, hotspots are characterized by locations where large numbers of business travelers congregate and need high-speed data connectivity.

Public WLANs offer significant performance improvements over what can be accomplished with a mobile data service such as General Packet Radio Service (GPRS) or Code Division Multiple Access (CDMA) 1x. But they also have limitations. Foremost among them is that users must be within a few hundred feet of a radio access point to get service. This is typically not a big issue since most heavy data usage occurs in airport lounges, hotels, and other known locations. However, for business travelers on the move, there is no substitute for the ubiquitous connectivity of a mobile data service like GPRS or CDMA 1x.

The challenge for today's mobile network operators lies in how to combine these two very useful services into a single offering. Momentum is already building in this direction as dual-mode PCMCIA cards (for laptop computers and PDAs) and dual-mode cell phones start to make inroads in the market. These devices can access either an 802.11 radio network, if one is present, or one based on a mobile data service. And mobile network operators are beginning to recognize the advantages of an 802.11 service as an adjunct to their mobile data services. For example, 802.11 uses an unlicensed spectrum versus the licensed spectrum used for mobile data services. In hotspots where user density is greatest, an 802.11 service can be used to offload data customers from the licensed spectrum. The users get better performance from an 802.11 service, and mobile network operators free up licensed bandwidth to service more voice traffic. Mobile operators also have experience in pulling together roaming agreements that make public WLAN services viable for business travelers.

Of course, there are technical and business considerations in building a service that combines both radio access technologies. On the technical front: How do operators offer a consistent service across these two very different access technologies, and how do they seamlessly handoff users as they traverse radio access boundaries? The business considerations include how to pay the costs for building out networks, how best to charge for services, and how to reduce the costs of managing the services.

Mobile network operators have an advantage over wireless ISPs in that they already have an infrastructure in place to support mobile data services. But how do mobile operators tap into this infrastructure as they roll out 802.11 hotspot services? The *Cisco Mobile Exchange* framework makes this all possible. By using this framework, mobile network operators can roll out services based on different radio technologies including GPRS, CDMA 1x, and 802.11. Regardless of the radio access technology used, the end-user experience can be the same. The Cisco Mobile Exchange framework even supports the seamless handoff of customers as they cross from one radio technology to another. For instance, if a customer is using an IPSec VPN (IP Security virtual private network) to tunnel back to a corporate network, the customer could change radio technologies and the VPN would stay up. Customers would not be required to restart their session—a very desirable feature for the all-important business client.


STEVE HRATKO

STEVE HRATKO is manager of new product development in the Mobile Wireless Group at Cisco. With more than 20 years of experience in the industry, he has worked extensively with enterprise and service-provider clients (primarily mobile operators) to develop voice and data opportunities. He can be reached at shratko@cisco.com.

Framework Provides Intelligent Integration

The Cisco Mobile Exchange framework integrates the various components that are required to provide an interface between different radio access networks and a common IP infrastructure, such as the Internet. Cisco Mobile Exchange also provides a set of services that are typically needed to build out a mobile data service. In addition, it provides the scalability and redundancy that a mobile network operator needs. Because it is built for IP network services, the framework can easily integrate a public WLAN service as well. Figure 1 shows the key functional elements of the Cisco Mobile Exchange framework. It provides packet gateways—the Cisco Gateway GPRS Support Node (GGSN), Packet Data Serving Node (PDSN), and public WLAN access routers—to terminate the three most common radio access technologies (GPRS, CDMA 1x, and 802.11), along with mobile services that are required for profitability: Mobile IP (to support seamless handoff), service selection, content monitoring, and security.

Mobile IP (Internet Engineering Task Force RFC 2002) enables users to keep the same IP address while crossing over to a different radio access network (which might even be operated by a different wireless provider), thus keeping their session active. This technology is especially important in enterprise applications that use end-to-end VPNs, because it means not having to reestablish the tunnel and reauthenticate into the enterprise.

Service selection is a technology that helps providers brand their services. It allows them to intervene in data flows and determine particular services that subscribers can access. It also provides a common user interface and a uniform billing infrastructure. These capabilities allow service providers to exercise discrete control over service access and enable self-provisioning to reduce operational costs, speed service availability, and boost revenue.

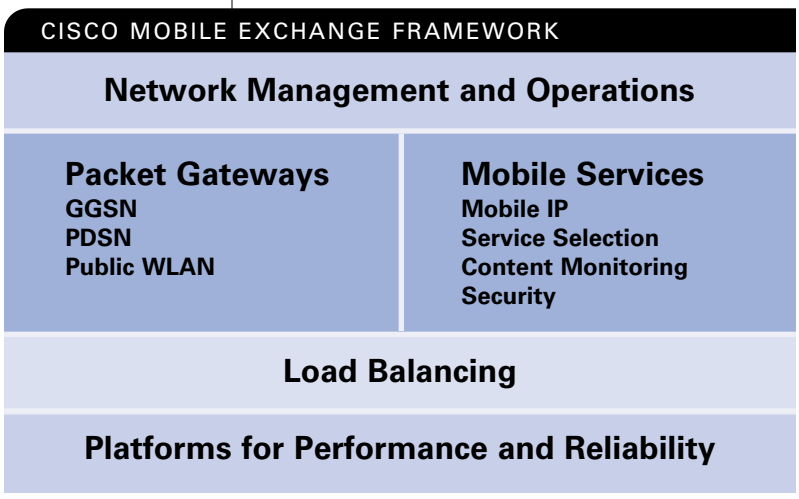

FIGURE 1: The Cisco Mobile Exchange integrates with existing radio access infrastructures through Cisco IOS® Software and application modules for Cisco platforms that many mobile network operators already own.

Content monitoring, alternatively called content billing, examines packets to obtain higher-layer information such as destination IP addresses, URLs, domains, applications, and file names. With this type of information, mobile network operators have great flexibility in how to bill for services.

Security is a critical function for services targeting the business community, and there are several approaches to supporting security. Most businesses bring the security function in house by using end-to-end VPN solutions based on technologies such as IPSec or Secure Sockets Layer (SSL), wherein customers initiate an encrypted tunnel from their client device (laptop computer or PDA) all the way through to the remote concentrator behind a firewall in their enterprise data center. They do not rely on the security capabilities of the mobile network operator. In other situations, customers will look to the mobile network operator to provide security. Cisco Mobile Exchange supports both options.

Finding its way into mobile data networks and 802.11 hotspots all over the world, the Cisco Mobile Exchange framework provides the services described above for each of the different radio access technologies—making it easier for mobile network operators to provide a common set of IP services and a common customer experience.

The Framework in a Public WLAN Service

Cisco Mobile Exchange can support a wide variety of services to address different business opportunities. For mobile network operators that are planning a service based on public WLAN and either GPRS or CDMA 1x, the Service Selection Gateway (SSG) within the Cisco Mobile Exchange umbrella becomes a must have (see Figure 2). The SSG is a very flexible platform with a broad array of capabilities. The four primary functions that it provides mobile network operators are brand-identity promotion, aggregation of different networks (public WLAN and GPRS, for example), selection of services, and support for captive portal applications.

Brand identity is one of the most important assets a mobile network operator can have in today's competitive environment. Regardless of how users access the network (public WLAN or GPRS), the first screen they come to will have the mobile network operator's name on it and a prompt for username and password. This approach helps to reinforce brand identity, provides a common user interface, and a common look and feel for the operator's services. Brand-identify reinforcement is especially important as mobile network operators move into offering public WLAN services as an adjunct to their mobile data services.

Network aggregation in the SSG allows for two different radio access technologies to be merged into a single service. Regardless of where the traffic comes

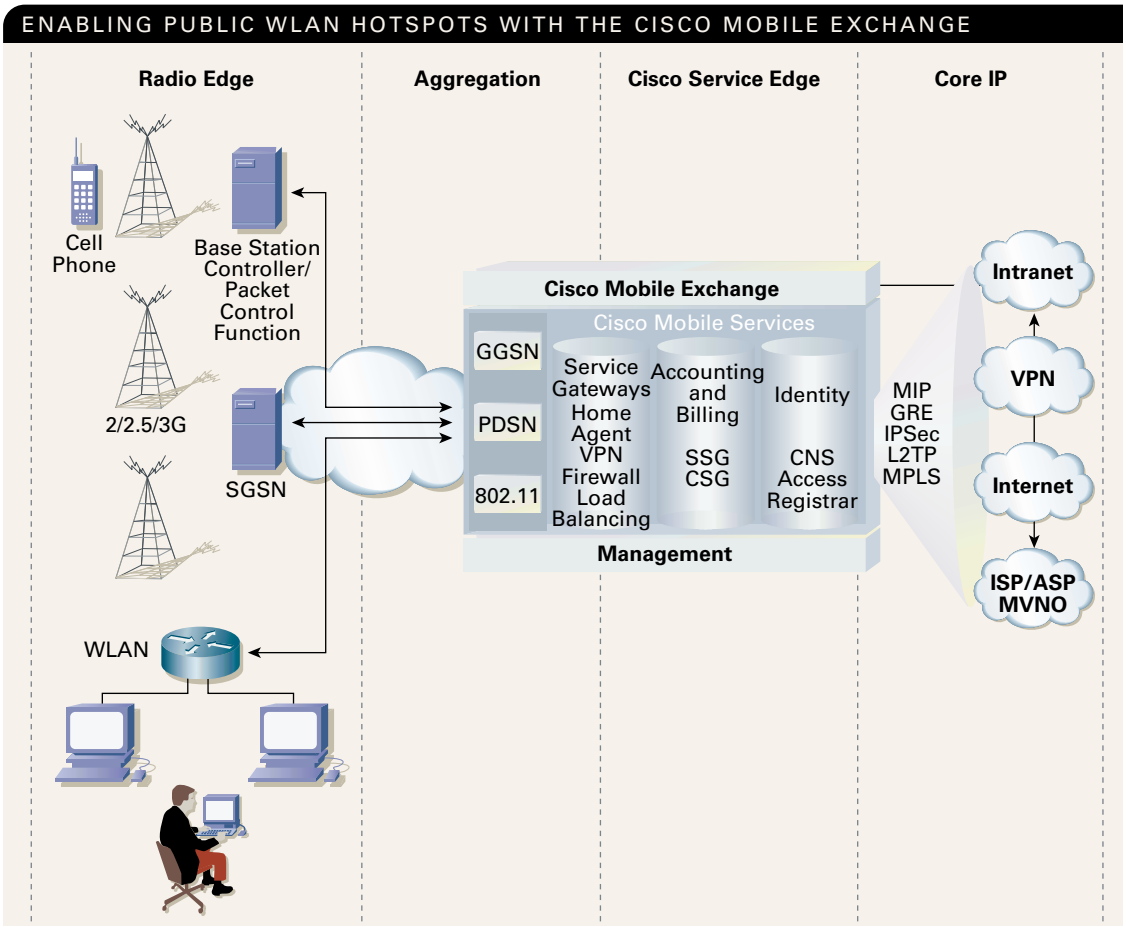

FIGURE 2: The Cisco Mobile Exchange framework combines mobile data services and public WLAN services into a single, coherent offering—for providing seamless handoff across radio technologies and rolling out new services.

from (a GGSN, PDSN, or a public WLAN access router), it is all brought together at the SSG where a common set of services can be offered.

Service selection allows users to have access to a variety of services, including access to the Internet, a corporate VPN (assumes the business does not want this done via end-to-end tunnels over the Internet), and "Walled" or "Open" gardens. Open gardens typically provide content free of charge to users and are usually created and underwritten by local businesses. Walled gardens provide unique content at an additional cost, for example, location-based content that can only be provided by the mobile network operator. Users select their desired destination from the logon screen that is presented. The SSG also has the ability, via authentication, authorization, and accounting (AAA), to verify which services a user is authorized to access.

Captive portal is the starting point for each user's session. It is the screen that comes up when they logon and the first thing users see when they access a mobile network operator's service. From here, the mobile network operator can enable the user to access different services. It is the captive portal that allows operators to put their brand front and center and provides them with the opportunity to sell additional services to users.

For many operators looking to provide value-added data services for their customers, the SSG is the first step but by no means the last. The Cisco Mobile Exchange offers other functional elements that can improve the user's experience and increase the operator's profitability. These include support for prepaid and postpaid billing at the service level (IP address, protocol port) and content level (HTML, FTP, Wireless Access Protocol); aforementioned support for Mobile IP to enable seamless handoff as users move between different radio access technologies; server load balancing to allow these solutions to scale to large numbers of users cost effectively; and VPN technologies to address a variety of security needs. One option involves using VPN technology to tunnel public WLAN traffic from a remote access point back into the mobile network operator's network. VPNs can also be used to connect back into a corporate intranet.

The Cisco Mobile Exchange framework should be an essential part of a mobile network operator's deployment plans. It is the vehicle for merging mobile data services and public WLAN services into a coherent offering, for providing seamless handoff across radio technologies, and for rolling out services that can generate and sustain profitability. ▲▲

IP Communications Made Easy

With IP telephony and a converged network, midsized retailer Edwards Fine Foods finds its communications as sweet as its desserts.

BY JANET KREILING

If you're a baker like Edwards Fine Foods, you want desserts—eight-layer chocolate cake, key lime pie, carrot cake—to die for, but your network alive and leavening your business. With five locations in California, Utah, and their home state of Georgia, many of Edwards' 1000 employees are frequently on the road, creating and nurturing high-quality products that consistently win national awards and doing top-notch marketing to food service companies, restaurants, and retail supermarkets. The Cisco IP communications solution that Edwards installed last September provides dynamic communications—IP telephony, unified messaging, extension mobility, virtual private networks (VPNs) to the home, and more—that help energize its people's work, making them more effective on the road and in their offices.

Good communications for mobile workers is especially important for Edwards, according to Steve Alsop, infrastructure manager for Edwards Fine Foods, because its employees travel more than those of most midsize companies. Because it makes products that have to be fresh and distributes them all over the US, Edwards' manufacturing plants are located where distribution, raw materials, and workforce conditions dictate. The network must reliably and flexibly tie distant locations together and also support mobile workers—the company keeps lean in part by sharing staff members and expertise across locations, and its sales staff traverses the country.

Edwards' parent firm, The Schwan Food Company (schwansinc.com), recently bought Mrs. Smith's Bakeries and combined the two subsidiaries. This created an opportunity to build out the IP network to include a total of nine locations. "The executive staff saw the value in IP telephony, unified messaging, portability, the ease of moves, adds, and changes, and a quick return on investment," Alsop notes.

The ROI for the initial Edwards converged network will be 18 months, he adds; the estimate for the buildout, which includes some new fiber and other infrastructure and will be completed this year, is 26 months. Just the savings in toll bypass charges from IP telephony in the Edwards network adds up to some US$1600 a month. Moreover, Alsop's staff of three put in the Edwards network itself and is now doing the buildout and all maintenance as well.

The initial network, pulled from Cisco's line of IP communications solutions for midsized businesses, is enabled by Cisco AVVID (Architecture for Voice, Video and Integrated Data). The network centers on an integrated communications system installed in Norcross, Georgia—the Cisco ICS 7750—which is equipped with Cisco CallManager call-processing software and Cisco Unity™ unified communications software. From this hub, T1 and Frame Relay links run out as spokes to Cisco 3725 multiservice access routers, which act as voice gateways, at the five locations. To segment bandwidth within the locations and deliver power to each phone, Alsop's team installed 26 Cisco Catalyst® 3524 PWR switches throughout the network.

"Every port can provide power over Ethernet so we can plug in a phone anywhere," Alsop says. On the desks are Cisco 7940G IP phones, and Cisco Aironet® 1200 Series wireless access points are being distributed throughout the manufacturing plants.

The buildout will be more of the same, with the exception of the core switch. Alsop will install a Cisco Catalyst 6513 switch in what will be the new data center for both companies in the former Mrs. Smith's headquarters in Suwanee, Georgia. Cisco Unity and Cisco CallManager are now connected by this switch, which is the new network aggregation point.

Alsop is placing Cisco 3725 routers at the new locations in Suwanee; Pottstown, Pennsylvania; Spartansburg, South Carolina; and Stilwell, Oklahoma, and adding 20 more Cisco Catalyst 3524 PWR switches. With 250 new IP phones in Suwanee alone, the network is virtually doubling in size.

Any Communication, Anywhere, Any Time

"There's not a single piece of our infrastructure that is not Cisco," Alsop says. "We have virtual private network solutions, mobile solutions, remote access solutions—there are all kinds of ways people can get the information they need quickly on the road. Our people can communicate just about anything from anywhere at any time."

Key to this flexibility is Cisco Unity unified messaging, Alsop says. "Our managers have the ability to send voice communications through the e-mail system to everybody on a distribution list—for example, updates on products to their sales staff. They can copy multiple people at multiple locations with just a couple of clicks. If employees are away from their desks, they can get voice and e-mail delivered by voice on IP phones, cell phones, or laptops; if at their desks, in their native form or as voice or e-mail."

Moreover, all voice mail is enabled by the same Cisco Unity system, so employees temporarily at other locations can access their address books, speed dials, or old messages. And any employee, at home or away, can forward voice mails to any other employee. Easy-to-use teleconferencing through Cisco CallManager also brings employees together.

Cisco CallManager helps employees on the move as well. Its extension mobility feature lets them log onto any Cisco IP Phone 7940G anywhere in the WAN and acquire their own extensions, user profiles, configurations, and privileges. "This makes it much easier for us to share people among locations," Alsop says. "An engineer can go to Stilwell, log in to the phone, and get all his or her speed dials, e-mail, voice mail, everything, just as though they were in Atlanta," he continues. Given how much Alsop himself is on the road lately with the installations, he appreciates the convenience.

Whether at home, on the road, or at the office, employees can call anywhere in the nationwide network simply by dialing the extension they want—no dialing of 10 tedious numbers. "We especially wanted dialing by extension among all our facilities," Alsop says.

He benefits personally from another feature of Cisco AVVID: the ability to set up a secure IP VPN to his home. Several employees use VPNs at home, equipped with Cisco 3002 VPN client systems with embedded firewall protection. When working from home, these people have all the features available on the WAN such as extension dialing and teleconferencing.

Edwards's new network delivers reliability as well with the Survivable Remote Site Telephony (SRST) feature of its routers' Cisco IOS® Software.

"We've had several storms this year in the Atlanta area, and the power grid goes down frequently," Alsop says. "There have been several instances where SRST has kicked in. It keeps voice service up for the building via a PRI line to the PSTN." SRST turns the router into a mini-CallManager so it recognizes when an employee wants dial tone and forwards incoming calls to the right extension. "Edwards alone processes about 8000 calls a day so a failure would be catastrophic."

To make sure the benefits of the new network can reach the factory floor, Edwards is installing Cisco Aironet 1200 Series wireless access points that will enable employees to handle tasks such as updating inventory in real time for more efficient planning and supply-chain management.

Easy IP Telephony for Data People

Alsop priced out several different solutions before deciding on Cisco IP telephony. "The traditional PBX

Increase Your Internet Quotient

Small and midsized businesses can realize initial cost savings by moving to a converged network. Other cost savings come from reduced long-distance bills and lower operating costs. New features and benefits offered by IP communications spark greater productivity and enthusiasm among customers and employees alike. To learn more about IP communications, visit *iQ Magazine* at cisco.com/go/iq-iptelephony.

IP Communications—Competitive Necessity

IP communications—which includes solutions such as IP telephony, unified communications, IP-based audio and videoconferencing, and customer contact solutions—is well on its way to becoming a technology companies can't afford to do without. It improves productivity: According to a report on IP telephony productivity from Sage Research published in January 2003, IP telephony ranks second only to VPNs among technologies associated with improved employee productivity. It also saves costs on toll bypass and on maintaining separate voice networks and costly traditional PBX equipment. The question is no longer if a company will install IP telephony, but when. Waiting too long can put a company at a competitive disadvantage surprisingly quickly.

Demand in the US for IP telephony has already grown dramatically—it represented 17 percent of the total customer premises equipment (CPE) lines shipped during 2002, according to an April 2003, InfoTech study, "Enterprise Convergence: The Race of IP Telephony Supremacy." By 2004, InfoTech projects that IP telephony lines will surpass traditional CPE voice lines in number and by 2007 will have 60 percent of the market.

Edwards Fine Foods demonstrates all the reasons. Cisco IP communications:

- Improves support of mobile and remote staff
- Lays a communications infrastructure for quick access to future communications
- Standardizes the communications infrastructure
- Reduces, contains, or avoids expenses
- Improves IT staff and employee productivity

IP communications translates to a converged voice and data network with a wealth of applications to come. Among them are instant messaging, distance learning, teleconferencing with data sharing, and desktop videoconferencing. Cisco, with a wealth of expertise in creating converged networks, has documentation and experienced personnel to help companies get started.

solutions were vastly more expensive because of their proprietary nature and the cost of installation," he says. "And we are not big enough to have someone on staff to work on and administer a traditional PBX."

Then there are the savings in lines dedicated to voice, especially for extension dialing across the network. "Without Cisco IP telephony, we'd have to buy separate T1s and use them as tie lines. And then we wouldn't be able to do anything else with those lines, just use them for extension dialing."

With Cisco IP telephony, he points out, "You use the pipes you have anyway for e-mail, manufacturing systems applications, and other data. There are a mountain of reasons to have a data pipe, and only one to have a voice pipe. With IP telephony, you are using lines that are already there. That is really sweet."

Edwards terminates all types of network connections on the Cisco Catalyst 6513—voice, data, fiber, copper.

And, he adds, "Cisco AVVID is so flexible that we can do so many things, like hanging home office phones off our corporate network."

Alsop decided to install the network himself. "I'd had a lot of experience with Cisco equipment," he says. "The Cisco manuals are complete and easy to use. We called the Technical Assistance Center [TAC] only six times in three months for install clarifications. Cisco CallManager, Cisco Unity, SRST—if the manual didn't have something we needed to know, we consulted the Web site." Alsop says he saved the US$80,000 budgeted for a consultant to install the initial network, and expects to save another US$180,000 budgeted for the buildout. Even given his own expertise, Cisco made the job easier, starting with the parts list, Alsop says. "The parts list was perfect, even though it had dozens of parts. The job only went off as well as it did because Cisco was so involved. The TAC was also an amazing resource."

He adds that IP telephony was easy for his staff to learn. "It's just another data stream. My staff picked it up in a week or two. IP telephony just makes sense to data people." ▲▲

PUTTING THE IP IN HIPPOCRATES

By **RHONDA RAIDER**

Network technologies elevate healthcare and improve safety while driving down costs.

Improved healthcare at lower cost is now being achieved worldwide, thanks to innovative uses of wireless LANs, IP telephony, and IP communications. "Recent changes in healthcare have inspired providers to look to their networks for innovative solutions," says Frances Dare, director, Global Healthcare for the Internet Business Systems Group at Cisco. "Clinician shortages around the world have made productivity a critical issue. As more drugs go generic, pharmaceutical companies need to develop drugs faster and more effectively. Countries are imposing more stringent security and privacy regulations. And the threats of SARS [Severe Acute Respiratory Syndrome] and bioterrorism have raised awareness of the importance of improved public health surveillance."

Today hospitals and consumer health organizations worldwide are meeting these challenges and driving costs down with wired and wireless networks.

Untethering Caregivers

"The key value proposition of wireless is to reduce medical errors, increase caregiver productivity, and simplify workflow," says Anita Pandey, healthcare business development manager for the Wireless Network Business Unit at Cisco.

Clinicians at Sharp HealthCare of San Diego, California, save hours each week by accessing patients' electronic medical records (EMRs) using the hospital's wireless LAN, which is based on Cisco Aironet® 1200 and 350 series access points throughout the campus. Using personal digital assistants (PDAs) or mobile carts equipped with Cisco Aironet PC adapters, doctors and nurses can access EMRs from the patient's bedside, sparing themselves multiple trips to wired terminals every day. A side benefit: more time with patients.

"Our physicians remain at the patient's bedside even while ordering tests and issuing prescriptions," says William Spooner, chief information officer for Sharp HealthCare.

EMRs also enable more flexible workflows. "With paper-based systems, workflow is serial: first register the patient, then find a bed, then take the medical record," explains Pandey. "A bottleneck at one step can impede immediacy and quality of care. With an EMR, certain processes can be performed in parallel, and the precise order of those particular processes no longer matters so much—so in effect the quality of care is less beholden to process."

In fact, EMRs were the primary incentive of wireless for NorthBay Healthcare System, a not-for-profit care organization in Solano County, California. NorthBay Healthcare has already deployed 20 wireless applications supporting more than 1000 people.

"Wireless has enabled us to develop a much more comprehensive clinical information system," says NorthBay chief information officer Paul Alcala. "Nurses now have the flexibility to document patient status at any point in care delivery: at the bedside, outside rooms, or outside nursing stations. They save hours each week by not having to walk to a terminal."

Reducing IT Costs

Like many healthcare providers, NorthBay Healthcare experienced immediate cost

savings after deploying its wireless LAN. "By equipping caregivers with wireless devices instead of installing a terminal in every room, we'll need 20 percent fewer access devices this year," says Alcala. "In addition, for our new family practice clinic opening later this year, we're expecting US$30,000 yearly savings in infrastructure support costs by eliminating moves, adds, and changes."

Increasing Accuracy of Care and Patient Safety

Many hospitals are turning to their IT groups for ways to improve patient safety. Medication-related errors cost an estimated US$2 billion each year. And worldwide, between 44,000 and 98,000 people die in hospitals each year as a result of preventable medical errors, according to a 1999 report issued by the Institute of Medicine.

Wireless applications help prevent errors by bringing up-to-date, real-time decision-support data to the point of care. For example, a doctor who enters a prescription on a wireless device at the patient bedside—an application known as Physician Order Entry (POE)—can receive instant alert notifications about potential adverse drug events. Staff at Lindenlohe Orthopedic Hospital in Germany can access decision-support data such as X-rays through the hospital's IT systems anywhere in the hospital—even in the aseptic operating theater—thanks to 135 Cisco Aironet access points installed in the ceiling of the waiting room that adjoins the operating room. A surgeon can verify a previous diagnosis while carrying out an operation and use a laptop in the operating room to access and update the patient's records on the database. As a result, a patient's accurate surgical diagnosis is available as soon as the patient arrives in the recovery room or intensive care unit.

NorthShore Long Island Jewish (LIJ) Health System, the third-largest nonsecular, not-for-profit healthcare system in the US, recently won the prestigious Long Island Software Award, for its Web-based EMR system, called HealthPort. Physicians can log in securely over the Internet for a unified view of all patient information from disparate systems such as lab, pharmacy, radiology, or insurance. "When a physician is called in the middle of the night to make a decision, it's invaluable to be able to access the information immediately, without looking for a chart or wondering if a nurse is providing the correct information," says Rick Carney, chief information officer at NorthShore-LIJ.

Another aspect of patient safety is ensuring availability of hospital systems, especially communications. Over the last year, the local phone carrier for NorthShore-LIJ has experienced several outages. Because NorthShore-LIJ recently converted to voice over IP, hospital clinicians and administrators continued to communicate with each other without interruption—with no compromise to patient care.

COURTESY OF ST. JOSEPH'S HEALTHCARE HAMILTON

TELEROBOTIC SURGERY: Using real-time video and robotic controls, doctors at St. Joseph's Hospital in Canada perform acid reflux surgery on a patient 400 kilometers away

Ensuring Security and Privacy

Network solutions like Sharp HealthCare's bar-code application help hospitals address new regulations related to the security and privacy of patient data. The Health Insurance Portability and Privacy Act (HIPAA) in the US, Personal Information Protection and Electronics Document Act (PEPIDA) in Canada, and emerging regulation in the European Community are spurring healthcare organizations to secure their networks. The security features in the Cisco Aironet products were among the reasons that Lifespan, a not-for-profit network of five acute-care hospitals based in Providence, Rhode Island, chose Cisco solutions for its wireless infrastructure. In designing its infrastructure, Lifespan relies on the SAFE Blueprint for security from Cisco to examine the entire network infrastructure, both wired and wireless, to eliminate weak links.

"The entire SAFE Blueprint architecture, of which the Cisco EAP [Extensible Authentication Protocol], or LEAP, is a key part, gives us high certainty that patient confidentiality will be securely guarded," says David Hemendinger, chief technology officer for Lifespan.

Serving Larger Files to Different Clients

To provide access to EMRs and images to more people using different types of clients, IT groups are upgrading their infrastructures. Maimonides Medical Center of Brooklyn, New York, takes advantage of a Cisco storage-area network solution to extend its Picture Archiving and Communication System (PACS) application, which delivers digitized X-ray and CAT scan images over the network.

"To gain the bandwidth and storage capacity that we needed to extend PACS to other clinical departments, we created a storage-area network linking two server farms one mile apart, each equipped with a Cisco MDS-9216 Multilayer Fabric Switch," says Jeff Kaplunovich, manager of network services at Maimonides. For backup and disaster recovery, the two switches can communicate either using a 155-MB OC-3 ATM link over the hospital's SONET ring, or 2-gigabit Fibre Channel connecting the two server farms point to point.

For Gundersen Lutheran Hospital, of LaCrosse, Wisconsin, the challenge was to make its Web-based portal, which provides access to EMRs, pharmacy information, and scheduling, available to clinicians using PDAs. Gundersen's solution: use the Cisco CTE-1400 Transformation Engine to render the Web pages for the unique display requirements of the client. The solution recognizes specific Web-enabled wireless

devices, such as PDAs and mobile and IP phones, and transforms the content according to the device's display characteristics.

Speeding Response with Two-Way Voice

To respond more quickly to patient requests, hospitals are introducing wireless voice solutions that provide instant two-way messaging capabilities over IEEE 802.11X networks—and without the disturbance of overhead pagers and nurse call systems. "In the best hospitals in the US, the average response time to locate a caregiver at patient request is about seven minutes," says Pandey.

"Using a Cisco wireless IP phone can reduce response time dramatically—even to a few seconds if the caregiver is nearby when the request is made."

NorthShore-LIJ's converged network, based on Cisco AVVID (Architecture for Voice, Video and Data), also supports telemedicine applications. Among the most valuable is the telepathology application, which enables pathologists at different hospitals in the system to consult with each other in real time while viewing the same images.

Driving Errors from Claims Processing

The same network intelligence that prevents medication errors can be applied to charge capture for claims processing to save money. "Healthcare providers lose billions annually due to claims-processing-related errors, in many cases entering the wrong codes, which delays payment," says Pandey. "When claims processing is moved to the network, the hospital can associate a menu-driven code with the type of care so that the correct code is entered automatically. We're hearing about return on investment of 400 to 500 percent."

Empowering Healthcare Consumers

Other types of healthcare organizations are using the Internet to empower consumers to make more informed choices. For example, in the US, consumers can visit Web sites such as HealthGrades (healthgrades.com) to locate care providers by various search criteria. Recently, HealthGrades began providing wireless access so that people away from home can use wireless phones or PDAs to locate nearby emergency service providers or clinics.

And in Mexico and throughout Latin America, Internet technology is helping educate people in rural areas.

"Wired and wireless networking technologies make it possible to bring healthcare to under-serviced communities," says Dr. Tracey Wilen-Daugenti, a senior business manager at Cisco and a Stanford University fellow. Wilen-Daugenti developed a consortia of private and public partners including Cisco, Macromedia, Reuters Digital Foundation, Stanford Center for Research in Disease Prevention at Stanford Medical School, and Tecnologico de Monterrey Medical School and Virtual University, to offer free, preventive health online courses to more than 350 communities in Mexico and the southwestern US. The public can access the online health courses from community learning centers in houses and churches.

A professor at the Tecnologico de Monterrey Medical Center, Doctor Gabriela Villarreal Levy has a further vision: to create an online diabetes center

The Business of Healthcare

Wireless LANs and unified communications play critical roles when working behind the scenes in healthcare.

- *Pharmaceutical companies* are beginning to deploy investigator relationship management (IRM) applications, a flavor of customer relationship management (CRM), to strengthen their relationships with clinical trial investigators. "Only 52 percent of physicians participate in two or more clinical trials because of the administration burden," says John Grant, manager, Pharmaceuticals Practice, for Cisco's Internet Business Solutions Group in Europe, the Middle East, and Africa. "With unified communications in the contact center, the pharmaceutical companies can identify the physician caller to the agent, present a history, and make intelligent recommendations for service—such as offering to send supplies or drugs that have likely run low." Grant notes that pharmaceutical companies also are taking advantage of their networks for electronic data capture for clinical trials—potentially speeding approval by weeks or months—and for convenient, cost-effective video-based training of investigators.
- *Payers* such as health insurance providers are using Cisco content networking to deliver video training for their contact center employees. For example, as part of this transformation, Empire Blue Cross implemented a new LAN architecture. It also deployed services to streamline routine processes and develop stronger relationships between senior management and its employees, and between the company and its customers. These new services run across a multicast network, a beefed-up intranet, a new set of external Web portals, and a Web-enabled call center.
- *Hospitals* track assets—and even patients—with radio frequency ID tags (RFID) and a wireless infrastructure. "An organization that has invested in a wireless infrastructure can use it for ancillary but high-impact services such as supply-chain management and inventory management," says Pandey. Preventing loss of 40 misplaced infusion pumps can lead to savings of US$200,000 or more. And being able to track patients helps prevent infant abductions and improve safety of Alzheimer's patients.

St. Joseph's Hospital and Bell Canada Put Telerobotics to Work

In recent years, surgeons have turned to robotic devices to conduct more accurate, less invasive operations. These devices have gained widespread acceptance for treating a variety of urologic, general, and chest conditions. Yet, while the benefits of medical robotics are clear, the ability for surgeons to travel to remote clinics or patients to drive hours to a large hospital has hindered use.

Enter telerobotics, which makes robotic surgery possible over vast geographic distances using a network. "By bringing a less invasive surgical procedure to patients in remote areas, it's possible to reduce the complications and costs," says Dr. Mehran Anvari, a professor of surgery at McMaster University in Hamilton, Ontario, Canada.

In February 2003, Anvari conducted pioneering telerobotic surgery over an IP network. Using a three-armed robot manufactured by Goleta, California-based Computer Motion (which this year merged with Intuitive Surgical), Anvari, working from St. Joseph's Hospital in Hamilton, completed an acid reflux surgery on a patient 400 kilometers away at North Bay General Hospital in North Bay, Ontario. Since then, he has conducted more than a half dozen other procedures.

From the beginning, Bell Canada, working with Anvari, developed a scalable and reliable system that could serve the healthcare provider's needs. The two hospitals are linked by Bell's Virtual Private Network Enterprise (VPNe), which creates a private and secure connection over a 12-megabit IP backbone. Relying on Cisco Multiprotocol Label Switching (MPLS) technology using Cisco 12000 and 7500 series routers and Cisco IOS® Software, the Bell VPNe can overcome any failure within 50 milliseconds. The MPLS network automatically finds another transmission route if a network connection fails.

Because Bell VPNe is a Layer 3 solution, it's possible to dial up bandwidth on demand. In addition, Bell Canada has built redundancy and quality of service (QoS) into the entire network and developed a packet prioritization model to ensure that the data controlling the robot receives a higher priority than video and voice data. If the network fails completely, a qualified physician finishes the surgery using conventional techniques.

Bell Canada spent almost nine months testing the network. "Although we used off-the-shelf components, it was essential to build an infrastructure that would support the specific and demanding needs of medicine," says Harvey Stein, senior director of Managed Solutions Architecture for Bell Canada. "This is a whole new class of application that combines real-time video and real-time robotic controls to allow doctors to do things that were never possible before."

In the coming months, The Canadian healthcare system plans to expand the use of telerobotics to two additional sites, and Anvari envisions at least six more hospitals and 20 remote medical clinics deploying the system. "This technology is changing people's lives," he says.

To read more about both telerobotics over the network and wireless in healthcare, visit *iQ Magazine* at cisco.com/go/iq-stvincents.

—Samuel Greengard

to combat this growing national health threat. The Web portal will provide users with real-time access to an online nutritionist and health educator, a Reuters Health news feed in Spanish, and animated Macromedia Flash modules for illiterate populations.

"Through the community learning center model and the new online diabetes center, we have an opportunity to close the gap in healthcare between rural and urban areas in Mexico," says Villareal Levy.

IP on the Rise

From rural Latin America to New York City, from hospitals to universities to consumer health portals, IP technologies are elevating the quality of care while reducing costs.

"In the last 18 months we've seen an increase in the uptake of networking solutions in healthcare," notes Chris White, area vice president for US Sales at Cisco. "Patients and doctors receive better information and providers are paid more accurately. We're seeing technology applied for the good of medicine." ▲▲

FURTHER READING

- **Business strategies and solutions, healthcare: cisco.com/packet/153_10a1**
- **Article on the "Hospital of the 21st Century": cisco.com/packet/153_10a2**

Technically Speaking

Telephony Signaling over IP

BY HELEN ROBISON

Cisco began as a startup devoted to solving the problem of any-to-any data internetworking. The problem that once encumbered data networks—networks using different protocols that prevented them from exchanging information with each other—similarly plagues voice networks. So, it should be no surprise that Cisco is a leader in the effort to solve the problem of any-to-any voice signaling.

In the public switched telephone network (PSTN), this problem is solved with expensive interworking switches at the edge of each different signaling network—at country boundaries, for example. These PSTN switches and the associated infrastructure are expensive, and the links across which the signaling packets are exchanged are of relatively low bandwidth (56 or 64 kb per link), which is often not effective for utilizing the processing capabilities of the database or application servers in the PSTN Signaling System 7 (SS7 or C7) network. Consider that each country has at least one national SS7 signaling variant (the Cisco PGW 2200 PSTN Gateway has a library of more than 80 SS7 variants) and you can begin to glimpse the enormity of the problem.

The Internet Engineering Task Force (IETF) Signaling Transport working group (SIGTRAN), of which Cisco is a key member, set about developing an open architecture with functional and performance requirements to support voice signaling over IP. The open architecture of the SIGTRAN protocol suite (also variously called IPS7 and SS7 over IP, or SS7oIP) is particularly appealing to potential implementers. IP is used for its flexibility to be carried over any packet transport media. Stream Control Transmission Protocol (SCTP) is the common transport layer protocol used because of its many feature improvements over TCP or UDP. PSTN Call Control layers and above (for example, SS7 MTP3, ISUP, SCCP, or ISDN Q.931) are typically terminated at the PSTN data link or network layer and backhauled over IP and are called the User Adaptation (UA) layer.

SIGTRAN protocols such as MTP3 User Adaptation Layer (M3UA) and SCCP User Adaptation Layer (SUA) allow the application vendor (for example, Short Message Service Center (SMSC), IP-based Home Location Register (IP-HLR)) to develop only the application layer and not have to support the large number of SS7 variants. By reducing interworking complexities and integration problems, the time for developing and marketing these new applications can be much faster. Using SIGTRAN or SS7oIP also eliminates the signaling and bearer bandwidth limitations of the PSTN network, allowing application servers to support heavy traffic requirements. Further, hybrid PSTN and IP networks can be supported by using a Signaling Gateway (SG).

Distributed MTP3

One way that Cisco is using the SIGTRAN protocol suite is with the distributed MTP3 (DMTP3) feature. The purpose of DMTP3 on a Signaling Gateway or Signaling Transfer Point (STP) is to provide SS7 network node redundancy. It allows a signaling linkset to be distributed between two physical nodes. To the adjacent network node, this pair appears as a single SS7 node, eliminating single points of failure. Cisco IOS® Software for the ITP-SG provides support for the DMTP3 feature. DMTP3 in the ITP allows SS7 signaling linksets to span two ITPs, enabling a point code to be shared between them. To enable links in a linkset to be shared across two ITPs, *distributed link*s are created. A distributed link is defined as a link in a linkset that is not physically attached to the local device. However, it is accessible using the DMTP3 peer. Traffic directed to the distributed link must travel from one peer to the other across a reliable IP-SCTP network, typically using a point-to-point, high-speed dedicated medium.

For more information about SS7oIP and the Cisco ITP, read "SS7 over IP: A Signal for Savings" (*Packet* Fourth Quarter 2002) at cisco.com/packet/153_11a1.

Distributed MTP3 and SS7oIP

A new software feature on the Cisco PGW 2200 PSTN Gateway supports the M3UA and SUA SIGTRAN protocol adaptation layers. This enables the PGW 2200 to interwork with the ITP-SG using SS7oIP. A large complex can be formed with a pair of ITPs and a farm of PGW pairs, allowing a higher call volume and a higher number of sustained calls. Using the DMTP3 feature on the ITP pair enables the ITPs and PGW farm to share a single SS7 point code, making the node appear as a single large entity and allowing consolidation of point codes and signaling links. Further, the PSTN Gateway supports voice over IP (VoIP) H.323 and Session Initiation Protocol (SIP) signaling, enabling numerous voice signaling protocols to interwork. PSTN and VoIP voice signaling protocols can both be supported over high-bandwidth IP networks and interconnect with the PSTN network.

These features are just a few of the ways Cisco is working toward support of any-to-any voice signaling. ▲▲


HELEN ROBISON

HELEN ROBISON is a senior technical marketing engineer in Service Provider Solution Engineering at Cisco. She can be reached at hrobison@cisco.com.

New Product Dispatches

Core Routing

Cisco 12000 Series Router: New Line Cards

Two new Cisco 12000 Series IP Services Engine (ISE) line cards enable service providers to integrate services for ATM and Ethernet environments with Layer 2 transport. The four-port OC-12/STM-4 ATM ISE line card transports ATM traffic over a 10-Gbit/s IP Multiprotocol Label Switching (MPLS) infrastructure with advanced ATM traffic management and sophisticated IP-to-ATM quality-of-service (QoS) functions. This card is ideal for ATM WAN connectivity, low-speed access aggregation, and business DSL aggregation. The four-port Gigabit Ethernet ISE line card provides comprehensive IP MPLS features for metro Ethernet aggregation, network peering, and IPv6 networking.

cisco.com/go/12000

Edge Routing, Access, and Aggregation

Cisco 3220 Mobile Access Router

As part of the Cisco 3200 Series Mobile Access Router, which provides seamless mobility while moving across wireless technologies, the new Cisco 3220 model is targeted, in both price and configuration, at high-volume deployments in public safety and commercial transportation environments. The preconfigured Cisco 3220 Router provides three Fast Ethernet ports, two synchronous serial ports, and one asynchronous serial port bundled with a Cisco IOS IP Security (IPSec) Triple Data Encryption Standard (3DES) PLUS software image.

cisco.com/go/3200

Cisco 10000 Series Router: New Engine and Line Card

The Cisco 10000 Series Performance Routing Engine (PRE-2) offers Cisco's adaptive network processing technology and 6.2-million-pps processing. The PRE-2 delivers more than 60,000 subscriber sessions for a service provider's broadband aggregation networks. The Cisco 10000 Series 1-Port OC-48c/STM-16 packet-over-SONET/SDH (POS) line card provides superior connectivity to the core network with line-rate capacity of 2.488 Gbit/s and POS technology. Internet service providers (ISPs) can use this card to substantially increase data volume on an optical transport infrastructure. Cisco also understands service providers' need for comprehensive network management and offers a range of carrier-class element managers including the new Cisco 10000 Manager for 10000 Series routers (see the new products under "Network Management," page 87).

cisco.com/go/10000

Cisco 7600 Series Router: New Models and Optical Services Modules

The new Cisco 7613 Router provides 13 slots, and the Cisco 7609 Router provides nine slots, enabling service providers to add services at the network edge while increasing efficiency. Both models provide 30 million pps forwarding rates and up to 256 Gbit/s total throughput for WAN connectivity from DS0 to OC-48/STM-16 and LAN connectivity from 10 Mbit/s Ethernet to 10 Gigabit Ethernet. The Enhanced 4-Port Gigabit Ethernet Optical Services Module (OSM) enables metro Ethernet services by aggregating switched Layer 2 traffic and transporting it over a scalable IP/MPLS backbone. The 1-Port Channelized OC-12/STM-4 to DS0 OSM and the 12-Port channelized T3 to DS0 OSM deliver high-density T1/E1 or DS0 private line services with up to 8000 queues per line card.

cisco.com/go/7600

Cisco 7304 Router: NPE-G100 Network Processing Engine

The Cisco 7304 NPE-G100 Network Processing Engine delivers superior general-purpose processing performance, flexibility, and high availability to powerful applications and services at the network edge including IPv6, multicast, and Any Transport over MPLS (AToM). The NPE-G100 offers 1 GB of SDRAM and three 10/100/1000 Ethernet ports for cost-effective, multigigabit connectivity as a high-end, managed customer premises equipment (CPE) or enterprise WAN gateway router.

cisco.com/packet/153npd1

Cisco 3700, 3600, and 2600XM Series Routers and 2691 Router: New Voice/Fax Network Module

For enterprises and service providers implementing IP communications, IP telephony, or toll bypass, the IP Communications Voice/Fax Network Module (NM-HD) and associated voice interface cards (VIC2) deliver mixed densities of analog and digital voice connectivity (from 2 to 48 channels) together with data and integrated voice/data access

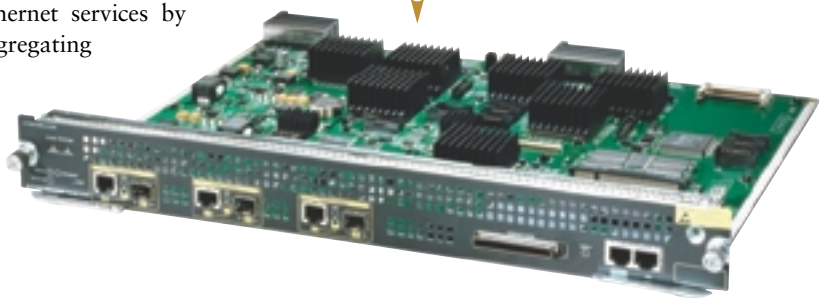

connectivity. The module and line cards are available for selected Cisco 3700, 3600, and 2600XM series models and the 2691 Router. The VIC2 cards provide support for any two combinations of standard voice protocols for connection to analog phones, fax machines, the public switched telephone network (PSTN), key systems, and PBXs.

cisco.com/packet/153npd2

Switching

Cisco Catalyst 3750 Series Switches

Cisco Catalyst® 3750 Series switches feature Cisco StackWise™ technology, a 32-Gbit/s stack interconnect technology for building a unified, highly resilient switching system. Cisco StackWise technology connects up to nine individual Cisco Catalyst 3750 Series switches into a single logical unit that accommodates voice, video, and data applications. The Catalyst 3750 Series also delivers advanced features for failover, ease of use, quality of service (QoS), and IPv6 readiness. For more on the Cisco Catalyst 3750 Series and Cisco StackWise technology, see page 67.

cisco.com/go/catalyst3750

Cisco Catalyst 2955 Series Switches

The new Cisco Catalyst 2955 Series Switch provides wire-speed Fast Ethernet and Gigabit Ethernet connectivity for deployment in harsh environments, such as manufacturing plants, intelligent transportation systems (public-sector highways and mass transit, for example), and power and utility stations. The Catalyst 2955 Series uses industrial-grade components, a compact form factor, convection cooling, and relay output signaling to extend intelligent services such as enhanced security, high availability, and advanced QoS to areas that cannot be served by traditional commercial-grade Ethernet switches. For more on the Catalyst 2955 Series Switch, see page 64.

cisco.com/go/2955

Cisco Catalyst 2950 LRE 997 Switch and Cisco 576 Customer Premises Equipment

Designed for central office deployments, the new Cisco Catalyst 2950 LRE 997 Switch enables service providers to extend metro Ethernet services over existing phone and legacy wiring. The switch features DC power and ETSI 997 band plan compliance with 24 Long-Reach Ethernet (LRE) ports, two Gigabit Ethernet ports, and two Small Form-Factor Pluggable ports. The Cisco 576 LRE CPE is a simple, 1-port bridge that translates LRE into Ethernet and provides one Fast Ethernet port and RJ-11 connectors for an analog phone and for the LRE link.

cisco.com/go/lre

Cisco Catalyst 6500 Series Four-Port 10 Gigabit Ethernet Module

The four-port 10 Gigabit Ethernet module for the Cisco Catalyst 6500 Series enables cost-effective migration from four or more gigabit EtherChannel® links to 10 Gigabit Ethernet in enterprise campus and service provider metro Ethernet deployments. Featuring a 40 Gbit/s switch fabric interconnection to the Cisco Catalyst 6500 Series Supervisor Engine 720, the 10 Gigabit module delivers full nonblocking performance with two ports and is up to 98 percent nonblocking when all four ports are enabled. Other features include 10-ms buffers and eight queues with multiple thresholds. The module supports hot-pluggable Xenpak long-range and extended-range modular optics for transmission up to 40 km over single-mode fiber. A choice of daughter cards that support centralized and distributed forwarding provide flexibility and investment protection. In the future, an accelerated Cisco Express Forwarding daughter card will be offered for additional scalability.

cisco.com/packet/153npd3

Security & VPNs

Cisco VPN 3000 Series Concentrator Software Version 4.0 and AES Encryption Module (SEP-E)

Among the many new features of Cisco VPN 3000 Series Concentrator Software Version 4.0 are Lightweight Directory Access Protocol (LDAP) authorization, Active Directory authentication, and RADIUS-based user filters (access control lists). Version 4.0 also includes support for the new AES-module (SEP-E), which provides Advanced Encryption Standard (AES) encryption in hardware, as well as support for Data Encryption Standard (DES) and Triple DES (3DES), for a fivefold increase in performance compared to AES in software.

cisco.com/packet/153npd4

Cisco VPN Client Version 4.0

Cisco VPN Client Version 4.0 includes an enhanced user interface for Microsoft Windows and Mac OS X operating systems. VPN virtual adapter support in this version increases application compatibility under Windows 2000 and XP, including support for Microsoft NetMeeting/H.323.

cisco.com/packet/153npd5

Cisco PIX Security Appliance Family: New Models

Three new models in the Cisco PIX® Security Appliance Family offer tailored functionality for specific applications. The Cisco PIX 501-Unlimited enables unlimited user connections in a security appliance ideal for teleworkers and small businesses. The Cisco PIX 515E-6 FE provides six ports of Fast Ethernet with 190-Mbit/s firewall and 140-Mbit/s VPN throughput for enhanced protection of multiple demilitarized zones or server environments that require physical network separation. The Cisco PIX 525-2GE delivers integrated dual Gigabit Ethernet and dual Fast Ethernet ports for flexible deployment in high-throughput environments.

cisco.com/go/pix

Cisco 7200 Series Router: VPN Acceleration Module 2

The VPN Acceleration Module 2 (VAM2) for Cisco 7200 Series routers provides high-performance encryption, compression, and key-generation services for IP Security (IPSec) VPN applications. The module supports AES and both DES and 3DES for up to 260 Mbit/s of encrypted throughput, as well as hardware-assisted Layer 3 compression services. Designed for large enterprises and small service providers, VAM2 supports site-to-site and remote-access IPSec VPNs.

cisco.com/packet/153npd6

Cisco 3700, 3600, and 2600 Series Routers: New Intrusion Detection Module

The new IDS Network Module implements intrusion detection system (IDS) capabilities in Cisco 2600XM, Cisco 3660, and Cisco 3700 series routers. This module uses a separate processor to free the router CPU from processor-intensive intrusion detection tasks and can monitor up to 45 Mbit/s of traffic in T1/E1 and T3 environments.

cisco.com/packet/153npd7

Cisco 2600 Series Routers: VPN Advanced Integration Module

The new VPN Advanced Integration Module (AIM-VPN BPII) can be added to current Cisco 2600, 2600XM, and 2691 routers. This module offers DES/3DES and new AES support with hardware-assisted compression services. The module's hardware-based encryption services deliver up to 20-Mbit/s 3DES performance in the Cisco 2621 Router and up to 22 Mbit/s in the Cisco 2651XM Router.

cisco.com/packet/153npd8

Optical

Cisco ONS 15302 and ONS 15305 Multiservice Access Platforms

The Cisco ONS 15302 and ONS 15305 Multiservice Access Platforms extend a multiservice network to the customer premises, allowing service providers to cost-effectively support data services and existing voice traffic over SDH infrastructures. The Cisco ONS 15302 delivers low-cost access with support for 12 E1 ports, four 10/100BASE-T ports, and an STM-1 optical uplink. The Cisco ONS 15305 supports access or aggregation with 63 E1 ports, six E3 ports, eight 10/100BASE-T ports, two Gigabit Ethernet ports, and scalability for STM-1, STM-4, and STM-16 optical uplinks.

cisco.com/go/ons15300

Cisco ONS 15530 Multiservice Aggregation Platform: New Cards

Two new cards for the Cisco ONS 15530 Multiservice Aggregation Platform help service providers make more efficient use of capacity in fiber-optic networks. The 8-Port Fibre Channel/Gigabit Ethernet Aggregation Card uses wavelengths efficiently by aggregating up to eight Fibre Channel, fiber connectivity (FICON), or Gigabit Ethernet signals on a single 10-Gbit/s wavelength. The 2.5-Gbit/s ITU Trunk Card increases fiber carrying capacity for storage and data applications by aggregating multiple enterprise system connection (ESCON), Fibre Channel, FICON, and Gigabit Ethernet services over a single International Telecommunication Union (ITU) wavelength.

cisco.com/go/ons15530

Storage Networking

Cisco MDS 9000 Family: IP Storage Services Module

The IP Storage Services Module is an 8-port line card for the Cisco MDS 9000 Family of Multilayer Directors and Fabric Switches that supports Small Computer Systems Interface over IP (iSCSI) and Fibre Channel over IP (FCIP) simultaneously on each Gigabit Ethernet port. This module enables interconnection of remote storage-area network (SAN) islands with FCIP and extends SAN connectivity to IP-enabled servers using iSCSI. The modules can be configured for either short (550-m connectivity) or long (10-km connectivity) wavelengths. For a case study involving deployment of the Cisco MDS 9000 Family, see "Storage Strategy," page 59.

cisco.com/packet/153npd9

Cisco 7400 and 7200 Series Routers: FCIP Port Adapter

A new port adapter for Cisco 7200 Series and Cisco 7401 routers supports FCIP, enabling enterprises to interconnect SANs over long distances securely and reliably. The Cisco FCIP Port Adapter delivers connectivity at speeds from fractional T1 up to OC-3. The port adapter also provides advanced networking services such as data compression, encryption, access control lists, firewalls, and quality of service (QoS).

cisco.com/packet/153npd10

Wireless

Cisco Aironet 1400 Series Wireless Bridge

The Cisco Aironet® 1400 Series Wireless Bridge offers a cost-effective alternative to leased lines for connecting multiple LANs in a metropolitan area. The bridge supports line-of-sight applications for point-to-point or point-to-multipoint configurations. Features include industry-leading performance with IEEE 802.11a data rates up to 54 Mbit/s, ease-of-use features that provide for simple installation and ready troubleshooting, a rugged enclosure for harsh outdoor environments, and support for integrated or optional external antennas. The Cisco Aironet 1400 Series Wireless Bridge is covered in greater detail on page 39.

cisco.com/packet/153npd11

Network Management

CiscoWorks Security Information Management Solution

CiscoWorks Security Information Management Solution (SIMS) 3.1 collects and analyzes security event information across the enterprise network for display on a real-time, graphical dashboard. Features of this software include event monitoring for a multivendor security environment; extensive reporting for network operators and administrators; threat assessment information for fine-tuning security policies; forensics tools to investigate attacks; event correlation to rapidly detect emerging attacks; and traffic utilization and trend analysis.

cisco.com/go/sims

Cisco Security Device Manager 1.0

The Cisco Security Device Manager (SDM) is an intuitive, Web-based device management tool embedded in Cisco 830, 1700, 2600XM, 3600, and 3700 series routers. Smart wizards enable users to quickly and easily deploy, configure, and monitor a Cisco access router without knowledge of the Cisco IOS® Software command-line interface (CLI). Configurations include LAN and WAN interfaces, firewalls, and virtual private network (VPN) features. SDM also offers a one-click router lockdown and an innovative security auditing capability to check and recommend changes to router configurations based on industry-recognized ICSA Labs recommendations.

cisco.com/go/sdm

Cisco IP Solution Center Version 3.0

Cisco IP Solution Center version 3.0 is the latest addition to the Cisco Internet OSS Domain Manager family. An integrated, carrier-grade management solution for Cisco routing, switching, and security product portfolios, Cisco IP Solution Center supports end-to-end IP services.

Cisco IP Solution Center applications include Multiprotocol Label Switching (MPLS) VPN management; Layer 2 VPN (ATM over MPLS) and metro Ethernet management;, security management application for IP Security (IPSec), VPN, firewall, and Network Address Translation (NAT) management; and policy-based quality of service (QoS).
cisco.com/packet/153npd12

CiscoWorks VPN/Security Management Solution 2.2

CiscoWorks VPN/Security Management Solution (VMS) 2.2 is an integral part of the Cisco SAFE Blueprint for enterprise network security and protects the productivity of organizations by combining Web-based tools for configuring, monitoring, and troubleshooting VPNs, firewalls, and network- and host-based intrusion detection systems (IDSs). CiscoWorks VMS 2.2 provides VPN configuration management, firewall management, surveillance, device inventory, and software version management features. Supported in VMS 2.2 is threat-protection technology Cisco gained in its April acquisition of Okena Inc., which looks at user behaviors and takes steps to stop actions that a company deems objectionable before they cause harm.
cisco.com/go/vms

CiscoWorks Wireless LAN Solution Engine Version 2.0

CiscoWorks Wireless LAN Solution Engine (WLSE) version 2.0 is part of the new Cisco Structured Wireless-Aware Network framework for deploying, maintaining, and monitoring a secure, fully integrated, enterprise-class WLAN infrastructure. CiscoWorks WLSE 2.0 provides support for centralized configuration and management of up to 2500 Cisco Aironet access points, and includes simplified operation and deployment; proactive performance monitoring; reporting, trending, and planning features; and a role-based user access model. For more on the Cisco Structured Wireless-Aware Network framework, see page 39.
cisco.com/go/wlse

Cisco Secure Access Control Server Solution Engine

The Cisco Secure Access Control Server (ACS) Solution Engine runs a 2.66-GHz processor in a 1-rack-unit form factor, and provides the same basic features and functions as Cisco Secure ACS for Windows in a dedicated, application-specific appliance package. Its hardened security posture, plug-and-play setup, and reliability benefits are aimed at improving total cost of ownership with reduced setup and simplified management. An integral component of Cisco's identity-based networking services (IBNS) solution, Cisco Secure ACS is the integration and control layer among all enterprises users, administrators, and network infrastructure resources. For more on Cisco IBNS, Cisco Secure ACS, and the new Cisco Secure ACS Solution Engine, see page 27.
cisco.com/go/acs

Cisco 10000 Manager

The Cisco 10000 Manager software delivers carrier-grade element management for Cisco 10000 Series routers in IP network-edge deployments. Detailed troubleshooting tools include fault detection on the router's cards and interfaces, as well as threshold and performance monitoring to maintain service availability. The Cisco 10000 Manager supports redundancy of Cisco 10000 Series routers with capabilities such as monitoring the active and standby Performance Routing Engines. This product also includes CiscoWorks Resource Management Essentials (RME) to simplify time-consuming tasks for administering network elements.
cisco.com/packet/153npd13

Voice & Video

Cisco Wireless IP Phone 7920

The Cisco Wireless IP Phone 7920 provides comprehensive IEEE 802.11b (Wi-Fi) wireless voice communications with Cisco CallManager and Cisco Aironet® 1200, 1100, 350, and 340 series access points. Designed for in-building use, the phone supports numerous calling and user convenience features while also leveraging the security, quality of service (QoS), and management advantages of a Cisco IP communications solution. The Cisco Wireless IP Phone 7920 is covered in greater detail on page 39.
cisco.com/packet/153npd14

Cisco IOS Software

New IPSec VPN Features

New IP Security virtual private network (IPSec VPN) capabilities in Cisco IOS® Software include IPSec/Multiprotocol Label Switching (MPLS) integration; IPSec accounting; Dynamic Multipoint VPN (DMVPN) tiering, load balancing, and self-healing; and Easy VPN failover. New trust and identity features include RSA key import/export; access control based on a certificate security attribute; source interface selection for HTTP traffic; Secure Sockets Layer (SSL) server; N-Tier certificate authority chaining; integration of Public Key Infrastructure (PKI) with authentication, authorization, and accounting (AAA); and secure RSA private key.
cisco.com/packet/153npd15

ABOUT NEW PRODUCT DISPATCHES

Keeping up with Cisco's myriad new products can be a challenge. To help readers stay informed, *Packet*® magazine's "New Product Dispatches" provide snapshots of the latest products released by Cisco between April and July 2003. For real-time announcements of the most recently released products, see News@Cisco at newsroom.cisco.com/dlls/.

Storage Strategy, Continued from page 61

SAN are better storage utilization and management, as well as an architecture that will support the university's projected storage growth. "The SAN will also be the foundation for our next-generation disaster recovery strategy," he adds.

Continues Warren, "Based on our ROI calculations, we expected to reduce the amount of time spent on disk subsystem and storage management by 60 percent. We are currently seeing these sorts of savings, and the staff members involved have been redeployed onto other challenging projects especially aimed at improving our capability to deliver even better teaching and learning outcomes."

Future Plans

Based on current storage and past growth, Deakin anticipates that its storage requirements will exceed 35 terabytes by the end of 2004. "We can't afford even a single student assignment to be lost, let alone the latest research results. So in our strategic planning it was deemed that all data generated by the university was corporate, and all data should be stored on the university's most reliable, available, high-performing storage platform—the SAN," says Warren.

The plan is to have all data stored on the SAN, from the file servers, to student records (including digitization of all student correspondence), to rich media curriculum content (e-learning).

Eventually, Deakin plans to implement secondary storage for disaster recovery at the Burwood campus. It will also support tertiary storage and backup in the Warrnambool and Burwood campuses. They plan to implement two Cisco MDS-9506's at each of these campuses in the coming months, migrating existing loop-attached (FC-AL) disk arrays and direct-attached storage to the SAN. Deakin is currently testing with iSCSI on an IP Services line card on the Cisco MDS-9216 that connects to the development and test VSAN on the Cisco MDS-9509's. The university plans to use its Cisco MDS IP Services line cards for both iSCSI and FCIP services.

Spanning the Distance

The wide geographic distribution of the Deakin campuses makes iSCSI and FCIP services particularly appropriate. With hundreds of kilometers between the campuses, most intercampus network connectivity is currently in the form of fractional E3 carried on microwave, over which no Fibre Channel connectivity is possible. The ITS group anticipates that a converged IP network that supports storage traffic will expand connectivity options for its campuses and provide a resilient, cost-effective network infrastructure.

For example, the hosts within the Geelong campus will simply connect to the hosts at Waterfront campus approximately 10 miles away using iSCSI directly across the IP network. According to Warren, they are looking at implementing FCIP tunnels across converged IP infrastructure between the Warrnambool and Waterfront campuses, and between the Burwood and Waterfront campuses to allow for replication between the sites.

The majority of data backups will be performed by tape silos at the locations where the storage is located. ITS will also support remote backups using a remote-campus tape silo to backup disk at another campus. They expect to reduce the burden that backups put on Deakin's overall data network by using the SAN instead of the data network.

Overall, a SAN enables efficient, simplified management of Deakin's storage and provides a robust, reliable data utility for the university community. A SAN can allow organizations to support their current storage requirements with a flexible architecture that will support growth. SAN technology is extremely effective within a campus and can be extended beyond the LAN with line cards, iSCSI, and FCIP. ▲▲

Cache File

SNIPPETS OF WISDOM FROM OUT ON THE 'NET

Mobile Phones Download DNA

Biologists can now seize inspiration whenever it strikes, thanks to a new database that sends genetic information to mobile phones. As reported by *Nature.com*, genomics researcher Björn Ursing of the Karolinska Institute in Stockholm, Sweden, has compiled a Wireless Application Protocol (WAP)-accessible database of all the 117 organisms that have had their genome sequenced. WiGID—the wireless genome information database—contains species' names, scientific descriptions, and reference to their published genetic sequence. It also has information on genome size in DNA letters, chromosomes, and genes.

Portland Top in Hotspots

The Portland, Oregon, metropolitan statistical area ranks as the top market for wireless Internet accessibility in the US, according to a new survey from Intel (intel.com/pressroom/archive/releases). Portland was followed by San Francisco, Austin, Seattle, Orange County, Washington, San Diego, Denver, Ventura, and Boston. Intel's rankings are based on the number of hotspots, the level of Internet penetration, and the number of mobile phone operators offering wide-area Internet access.

OVER HALF OF JAPANESE POPULATION ONLINE

More than 50 percent of the Japanese adult population is online, reports AsiaBizTech. New data from the Ministry of Public Management, Home Affairs, and Telecommunications reveals that in 2002, the number of people using the Internet in Japan increased by 13.5 million to 69.4 million—equivalent to a 54 percent penetration rate. More than 80 percent of Japanese households are now online, compared to 79.1 percent of businesses. In terms of broadband, approximately 29.6 percent of Internet users connect to the Net via a high-speed connection.

Sweden Ranks First in E-Readiness

According to recent research from the Economist Intelligence Unit, Sweden has overtaken the US to become the leading nation in e-readiness, which is defined as the extent to which a country's business environment is ready for Internet-based commercial opportunities. Among the ranking factors are connectivity and technology infrastructure, business environment, consumer and business adoption, social and cultural infrastructure, and legal and policy environment and support services. Denmark took second place in the e-readiness rankings, while the Netherlands, US, and the UK tied for third position. Finland and Norway took sixth and seventh place, respectively, and Switzerland came eighth with Australia ninth. Canada and Hong Kong tied for number ten.

CYBER QUOTE

"We reject KINGS, PRESIDENTS, AND VOTING. WE believe IN ROUGH consensus AND RUNNING code."

—David Clark (from a talk given to the IETF in 1992)

THE 5th WAVE

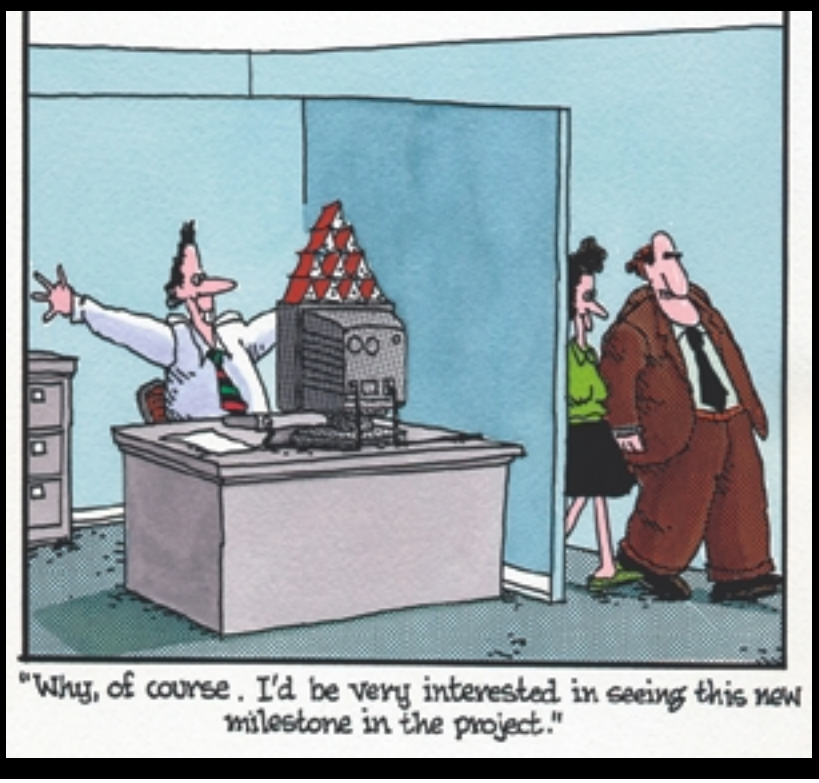

©*The 5th Wave*, www.the5thwave.com