

PACKET®

CISCO SYSTEMS USERS MAGAZINE

FIRST QUARTER 2004

Self-Defending Networks 24

Protecting the Network
from Human Nature

57 Extending SANs

52 High Availability
Campus Networks

39 Securing IP Voice

69 Troubleshooting
MPLS VPNs

CISCO SYSTEMS

cisco.com/packet

PACKET

CISCO SYSTEMS USERS MAGAZINE

FIRST QUARTER 2004 VOLUME 16 NO. 1



▼ ON THE COVER

The Self-Defending Network

The volume of threats to network security is rising, and so is the damage potential. While effective security technology is readily available, many companies still do not adequately protect themselves. This article provides an overview of the security landscape and offers insight on how organizations can get started with an effective security approach assisted by new technologies from Cisco. **24**

▼ FEATURES

A Winning Game Plan

Policies, processes, people, and products are the four critical ingredients that make up the best practices for successfully securing networks. A proactive security posture creates a synergy of expertise and technologies that can protect your business. Find out how. **29**

Security Advocates

Cisco's most customer-facing security team, the Product Security Incident Response Team (PSIRT), reacts quickly to product security issues and works hard to keep customers' trust. . . . **36**

Securing Voice in an IP Environment

As voice technologies meet the world of IP, many IT, telephony, and business managers are concerned about security. Defense-in-depth strategies for making IP voice as secure as any other mission-critical application are laid out here. **39**

FOCUS ON NETWORKERS

69 Troubleshooting MPLS VPN Networks

Drawing from his popular presentation at Networkers 2003, Cisco MPLS expert Rajiv Asati offers practical tips for troubleshooting Multiprotocol Label Switching virtual private networks (MPLS VPNs).

PACKET

TECHNOLOGY

17 **ROUTING: Locking Down IOS**

Cisco reinforces Cisco IOS® Software security with new features and best practices.

21 **SECURITY: Mitigating Network DoS Attacks**

This article presents solutions to proactively mitigate denial of service (DoS) attacks using the built-in protection of Cisco IOS Software for dynamic routing protocols.

ENTERPRISE SOLUTIONS

45 **Security Advances in Mobile IP**

Logging in to Mobile IP is more secure and easier to configure with new Cisco IOS Software features.

49 **Deployment Diary**

The conclusion to a firsthand account of how New York University deployed a novel Tetrahedron Core network.

52 **High Availability for Campus Networks**

Why and where do you deploy the Gateway Load Balancing Protocol (GLBP)? Design options network operators can use to achieve high availability and resource efficiency.

57 **Extending SANs**

Dense wavelength-division multiplexing (DWDM) and storage-area networking (SAN) extensions from Cisco provide disaster recovery for high-end synchronous mirroring applications.

SERVICE PROVIDER SOLUTIONS

61 **Triple-Play Cable**

Time Warner Cable's voice over IP (VoIP) telephone service signals the start of a major technological shift in the cable industry.

64 **Proactive Protection**

As more service providers offer converged services, including voice, the importance of network defense escalates. Now service providers can take a multifaceted approach to protect their infrastructure.

73 **The IP Revolution in Mobile Messaging**

IP-enabled Signaling System 7 (SS7) applications yield significant gains in bandwidth and greatly reduced operations costs.

75 **Ringing Up Profits**

New capabilities on the Cisco ONS 15454 ML-Series Line Card help Canada's Cogeco Cable target the lucrative business market.

SMALL AND MIDSIZED BUSINESSES

77 **Paradise Reconnected**

At the Eden Project, a popular tourist attraction and educational center in Britain, a Cisco IP communications solution helps to cut costs and expand communications capabilities.

▼ DEPARTMENTS

1 **From the Editor**

Who's Afraid of a Little Girl?

5 **User Connection**

Internet Training Centers ■ CCIE® 10-Year Anniversary ■ Cisco Certifications Community ■ Center for Internet Security ■ Cisco Powered Network Operations Symposium

9 **Tech Tips & Training**

Top Security Questions ■ Identity-Based Networking at a Glance ■ Wireless Security Update ■ Reader Tips

83 **Technically Speaking**

Cisco engineer Jason Dachtler discusses compatibility and interoperability issues for voice call transfer and forwarding in Cisco IOS Software.

80 **Netizens**

A glimpse into the world of a CCIE Lab proctor.

85 **New Product Dispatches**

A roundup of what's new from Cisco over the past quarter.




39

Securing IP Voice



61

Triple-Play Cable



36

Security Advocates

▼ IN EVERY ISSUE

Mail 3

Calendar 5

Networkers 6

Tech Tips 10

Advertiser Index 89

Cache File 90

The 5th Wave 90

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, NICOLE MAZZEI
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

STAN MUSELIK
COVER PHOTOGRAPH

**SPECIAL THANKS TO THE FOLLOWING
CONTRIBUTORS:** ROGER DELACRUZ, JANICE KING,
RHONDA RAIDER, HELEN ROBISON, KELLY VINCENT

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

Aironet, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, Cisco IOS, Cisco Networking Academy, Cisco Press, the Cisco Powered Network logo, the Cisco Systems logo, Cisco Unity, IOS, IP/TV, iQ, *Packet*, PIX, and StackWise are registered trademarks or trademarks of Cisco Systems, Inc., and/or its affiliates in the USA and certain other countries. All other trademarks mentioned in this publication are the property of their respective owners.

Packet copyright © 2004 by Cisco Systems, Inc. All rights reserved. Printed in the USA.

No part of this publication may be reproduced in any form, or by any means, without prior written permission from Cisco Systems, Inc.

This publication is distributed on an "as-is" basis, without warranty of any kind either express or implied, including but not limited to the implied warranties of merchantability, fitness for a particular purpose, or noninfringement. This publication could contain technical inaccuracies or typographical errors. Later issues may modify or update information provided in this issue. Neither the publisher nor any contributor shall have any liability to any person for any loss or damage caused directly or indirectly by the information contained herein.

This magazine is printed on recycled paper.



10%
TOTAL RECOVERED FIBER

Reprinted with permission from *Packet* magazine (Volume 16, No. 1), copyright © 2004 by Cisco Systems, Inc. All rights reserved.

FROM THE EDITOR

Who's Afraid of a Little Girl?

IF YOU DON'T HAVE AN EFFECTIVE NETWORK security plan and system in place, that girl on our front cover could be your worst nightmare. While it's true that threats from external sources are on the rise—78 percent, up from 57 percent in 1999, according to a 2003 CSI/FBI security study—threats from inside the organization can be up to ten times more expensive to repair. Internal threats and exposure can come from anywhere, anytime—an employee deploying a rogue access point, a techie misconfiguring a router, even the boss's daughter, as depicted in the new Cisco security ad campaign (see inside cover), can take down the network. Anyone can inadvertently harm your organization simply by visiting a Website, downloading a file, or opening an e-mail attachment.



The number of threats is not only increasing at an alarming rate, thanks to self-propagating worms and viruses, these destructive forces can spread across the globe in a matter of minutes. The advent of "point-and-click" hacking tools, and readily available Internet worms, has ushered in the era of the tech-illiterate cracker—now anyone can maliciously assault your network. Worse still, the attacks made by the "pros" are getting ever more sophisticated. Blended threats—combining worms with viruses with trojans—widespread system hacking, network denial of service (DoS), and attacks targeting applications or embedded in the data itself call for an entirely new approach to network security.

Cisco calls this approach the "Self-Defending Network." By viewing security as an innate and pervasive feature of the network infrastructure, not as an "add on," Cisco is building networks that protect themselves. You can read all about it, starting on page 24.

But technology alone isn't the answer. Best practices are a critical part of any security solution. Without them, it doesn't matter what technologies you're using. In "A Winning Game Plan," page 29, you can read all about the four P's of effective security—policies, processes, people, and products—and the different strategies necessary for securing your campus, data center, branch offices, and remote workers.

The most powerful defense of all against harmful network threats is information. The hacker community is constantly sharing information about the latest vulnerabilities and methods for exploiting them. As the governors and protectors of our networks, we must do the same: sharing fixes, technology solutions, and industry-specific best practices whenever possible.

We at *Packet*® hope to do our part with this issue. All in all, 37 pages are dedicated to network security: best practices for both enterprises and service providers; means of completing a security assessment for your organization; how to form an incident response team; and the latest information on protecting your wired infrastructure, wireless LANs and IP telephony networks. Check out the "Further Reading" links to increase your knowledge with more in-depth network security coverage.

If you only check out one link, I urge you to visit cisco.com/powernow/packet. Armed with better information, maybe we can all be a little less scared of that girl on our front cover.

David A. Ball

Editor-in-Chief
Packet
daball@cisco.com

Mail



Baffled by Hidden Gems

I really liked your "Hidden Gems" Tech Tips article [Third Quarter 2003], but none of the tricks described in the article seem to work for me. I am running Cisco IOS® Software Release 12.2.7b on a Cisco 7206 Router. Do you have any idea why the tricks don't work?

—Mark Coutu, Télécommunication Transcontinental, Ville St. Laurent, Quebec, Canada

Following is a response from "Hidden Gems" author Mark Basinski. —Editors.

You do not see the features because your Cisco IOS image is based on the 12.2 mainline IOS software release train. All the "hidden gems" mentioned in the article were included in the 12.2T release train, which is noted in the article.

Cisco IOS Software has several release trains, each of which targets the requirements of a specific target market. For example, the mainline train focuses on maximizing stability, and therefore does not allow the introduction of new features into the code base. The T train allows the introduction of new features and functionality—thereby continuing to enable customers to introduce new applications and advanced technologies into their networks, such as wireless, IP communications, and integrated security, as they emerge.

The mainline and T trains have a unique relationship with each other (see illustration at cisco.com/packet/161_2a2). At the start of Cisco IOS Software Release 12.2, the 12.2T train splits off and the new features, including the "hidden gems," are added. The 12.2T train eventually serves as the beginning of the 12.3 mainline release train.

To access the "hidden gem" features, you must run either 12.2T releases (refer to the article for the specific releases) or if you prefer to use mainline IOS code, the features are also now available in Cisco IOS Software Release 12.3.

Promoted in India

I have noticed that since I began subscribing to *Packet*® a year ago my knowledge and skills have improved. As a result, I have been promoted to senior systems engineer in my organization for the global network operations center of Patni Computer. How can I get copies of the 2001 and 2002 issues of *Packet*?

—Vikram Patil, Patni Computer Systems Ltd., Mumbai, India

You can find archived back issues of *Packet* online at cisco.com/packet161_2a1. The issues are available in PDF format suitable for printing. —Editors

Hunting for VoIP Troubleshooting Info

I have read and reread the article, "Voice Quality" [Fourth Quarter 2003]. I enjoyed the article and especially appreciated its focus. It seems that voice over IP [VoIP] can be very profitable, but if not done correctly at all levels, issues arise that go beyond just VoIP for a cable-based service. Because Comcast is about to begin testing, I want to know if Cisco has compiled any troubleshooting information that applies to Internet service provider service application of VoIP.

I provide technical support to the project in Denver and have been researching and designing troubleshooting for our

upcoming tests. Although I have searched on Cisco.com and also found other information resources on the Web, I would be interested to know if Cisco has additional information.

—Julia Ann Gilkey, Comcast Corporation, Denver, Colorado, USA

Following is a response from "Voice Quality" author Helen Robison. —Editors.

A team of voice quality experts is currently working on best practices that we hope to make available soon. Also, several significant software features have been added to Cisco IOS Software Release 12.3T to improve the retrievability of voice quality metrics and diagnostics for voice gateways. For now, I recommend Chapter 5, "Defining and Measuring Voice Quality," in *Integrating Voice and Data Networks* (Cisco Press). You can also try searching on a specific topic such as "echo" in the Cisco online documentation (cisco.com/univercd/home/home.htm), where you should find some very useful materials.

CORRECTION

The Reader Tip "Changing Dynamic IP NAT Configuration" (Fourth Quarter 2003, page 13) contained an incomplete command. The correct command to clear all active translation entries is `clear ip nat translation *`. Without the ending asterisk (*) the command does not work. We apologize for the error. —Editors

SEND YOUR COMMENTS TO PACKET

We welcome your comments and questions. Reach us through e-mail at packet-editor@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

Internet Training Centers in Developing Countries on the Rise

CISCO AND THE INTERNATIONAL Telecommunication Union (ITU) have announced plans to establish 20 new Internet Training Centers in developing countries. The expansion is part of an ongoing collaboration between the two organizations to provide Internet education and greater access to information technology throughout the developing world. The new training centers will be located in ministries of communications, or their equivalents, and will offer the Cisco Networking Academy® Program curriculum.

The centers are part of the Internet Training Centers Initiative (ITCI), created by ITU in 2001, to provide students and professionals in nonindustrialized countries with access to affordable and relevant technology training.

“Governments in developing countries recognize the importance of having skilled professionals to help them bridge the digital divide, and always welcome the opportunity to have adequate IT training facilities to train their own staff responsible for telecommunications policy,” says Hamadoun Touré, Director of the ITU Telecommunication Development Bureau. “One of our goals for the ITCI is to strengthen Internet skills on a large scale. We believe extending the initiative to include government officials can only benefit the economy of each country.”

Since its inception in 2001, the ITCI has established 57 centers at educational institutions. The goal is to train a minimum of 50 students per center. More than 2500 students have enrolled in the centers; 30 percent of those enrolled are females.

Cisco is a pioneer partner of the ITCI, and has established the Cisco Networking Academy Program for students to acquire

Internet technology skills. So far, 181 students have graduated from the CCNA® associate-level curriculum, with an overall employment rate of 83 percent.

Curricula

ITU and Cisco are extending the curriculum in 20 of the best-performing centers to enhance the IT competencies of students. The curricula, sponsored by Hewlett-Packard, Panduit, and Sun, will include IT Essentials 1: PC Hardware and Software; IT Essentials 2: Network Operating Systems; Fundamentals of Voice and Data Cabling Systems; Fundamentals of UNIX; and Fundamentals of Java Programming. Targeted centers will implement one of these courses in addition to the CCNA curriculum.

“The Cisco Networking Academy Program succeeds everyday in teaching students around the world the skills they need to join the information age and to help their communities,” says Tae Yoo, Cisco vice president of Corporate Philanthropy. “We are proud to extend our association with ITU to give developing

countries access to the same training opportunities available in a thriving urban community.”

One of the chief goals of the ITCI is to encourage greater female participation in learning IT skills. In furtherance of this goal, in 2002, Makerere University in Kampala, Uganda, began offering the Cisco Networking Academy Program to students in its Department of Women and Gender Studies through the partnership between ITU and Cisco.

Ugandan student Anita Sanyu Mago-Sempa, who works for a media services company, says, “My company is expanding into computer networking, an upcoming area where there is still a gap in such services offered. As a start, I will be able to network my company’s computers.”

Find Out More

For more information about the ITCI, visit cisco.com/edu/itci. For more information about the Cisco Networking Academy Program, visit cisco.com/edu/academy. ▲▲

Cisco Worldwide Events

FEB. 10–11	CALL CENTER WORLD 2004	BERLIN, GERMANY
FEB. 10–13	CISCO PARTNER SUMMIT	HONOLULU, HAWAII, USA
MARCH 1–4	VOICECON 2004	LAKE BUENA VISTA, FLORIDA, USA
MARCH 8–11	NETWORKERS AUSTRALIA 2004	BRISBANE, AUSTRALIA
MARCH 18–24	CEBIT 2004	HANNOVER, GERMANY
APRIL 5–8	STORAGE NETWORKING WORLD SPRING 2004	LAS VEGAS, NEVADA, USA

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

CCIE Program Celebrates 10 Successful Years



TEN YEARS AGO 31 networking pioneers agreed to take part in an emerging certification program to prove their ability to configure, test, and troubleshoot real-world networking equipment. These 31 individuals became the first to earn the elite CCIE® expert-level networking certification from Cisco, which recently celebrated its 10-year anniversary. Now widely considered to be the industry's most rigorous IT certification program, more than 10,000 CCIE certified individuals worldwide are involved in some of the most complex networking projects being implemented.

“Certifications provide organizations assurance that IT projects will be completed on time and on budget,” says Cushing Anderson, program director for IT education and certification research at IDC. “Complex technical environments require a higher level expertise for IT professionals

and the CCIE is widely seen as the pinnacle of IT certification programs.”

In 2003 *Certification Magazine* ranked the CCIE as “Best Hands-On Program” and “Most Technically Advanced Program” in the publication's Certification Top 10 Lists. The CCIE program was also voted number one by IT professionals in the CertCities annual survey of the “10 Hottest Certifications for 2003.”

Training and Availability

Hands-on experience is the one of the best ways to prepare for CCIE exams. Although no formal prerequisites exist, candidates are strongly encouraged to have at least three to five years of IT job experience before attempting certification. Authorized Cisco training for the CCIE is available from a global network of Cisco Learning Partners and the Partner E-Learning Connection.

For more details on the CCIE program, visit cisco.com/go/ccie. ▲▲

Networkers Online

A wealth of detailed technical content on networking technologies and solutions is now available from the premier online user group for networking professionals: Networkers Online. If you plan to attend Networkers 2004, you can take advantage of a variety of educational opportunities and prepare for the conference by joining Networkers Online at cisco.com/packet/161_3f1. The cost of a subscription is US\$150 and can be applied toward registration at Networkers 2004 (US conference only).

Cisco Certifications Community Available Worldwide

AFTER A SUCCESSFUL PILOT PERIOD, the Cisco Certifications Community Web portal is now available to Cisco certified network professionals worldwide. The portal, which attracted 20,000 member subscribers in its pilot phase last year, provides an interactive discussion 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.

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 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.

For the October 2003 launch, the Certifications Community provided three hours of free online training on security best practices, a CCNA and CCNP game on networking skills, and a Networkers toolbox with technical information and



profiles of successful CCIE experts.

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

Benchmarking Tools Available from Center for Internet Security

MANY COMPANIES LOOK TO information-security consultants and vendors to help identify vulnerabilities and improve overall network security. But the experts don't always agree. When companies develop their own policies and device-configuration standards, they can spend significant time, effort, and resources to do so, yet they can't always be sure they have implemented best practices.

The Center for Internet Security (CIS) seeks to address these problems by working with government agencies, educational institutions, consulting firms, and industry experts worldwide to gain consensus on acceptable baseline security controls for operating systems and networking products. CIS also provides organizations with tools to benchmark and improve the security of their computers and network devices.

Updated versions of the CIS consensus benchmarks and tools for the Cisco IOS® Router and Linux and HP-UX operating systems are available for download free of charge from cisecurity.org. An updated version of the free CIS scoring tool

software for Windows operating systems is also available.

CIS benchmarks specify security configuration settings and actions that strengthen a system's defenses against malicious attacks. CIS scoring tools provide a quick and easy way to evaluate systems and network devices, comparing their security configurations against the CIS benchmarks and producing automatic, easily understood reports.

The latest version of the Cisco IOS Router Audit Tool (RAT) includes more than 10 updates and improvements. A complete listing of these developments is available at cisecurity.org.

Formed in 2000, CIS is a nonprofit consortium of public and private sector participants that develops detailed operational security practice documents, or benchmarks, and makes them available for download, free of charge. CIS methods and tools measure, monitor, compare, and improve the security status of Internet-connected systems and network devices. CIS benchmarks and tools are updated regularly based on user feedback.▲▲

Cisco Powered Network Operations Symposium Expands to Europe

EUROPEAN STRATEGIC SERVICE providers now have the opportunity to attend the Cisco Powered Network Operations Symposium on their own turf. In its sixth year, the symposium is a must-attend annual event that focuses on technical topics such as mobility, voice, cable, IP, Metro Ethernet, Multiprotocol Label Switching (MPLS), optical, storage, subscriber services and managed services, wireless, and security. Courses address the topics from the perspective strategic service provider engineers worldwide need to further their knowledge and ability to deliver high-performance, reliable services over Cisco networks to their business customers. Until this year, the annual symposium has been held in the US only, with an average of 500 attendees, 33 percent of whom are from outside the US. Now, in addition to the US event scheduled for March 14–19, a second event will be held in Paris on September 5–10 and is expected to attract 250–300 participants.

Continued

Roundup of Activities

The European symposium's main focus will be mobility and wireless technologies. Attendees at both the US and European events will be able to receive industry updates and participate in basic, mid-level, and advanced courses, design clinics, hands-on sessions, and technical forums in specific technological areas such as mobility, voice, optical, security, cable, and network management. And, as an added benefit during the symposiums, technical members can increase their Cisco Career Certifications levels by sitting for a written or lab exam at no cost.

Partnering activities include a Partner Showcase with exhibits from Cisco Learning Partners, Cisco Powered Network Solutions Ecosystem partners, and Cisco technology groups.

For more information on the Cisco Powered Network Operations Symposiums in both the US and Europe, visit cisco.com/go/cpnsymposium2004 (accessible to Cisco Powered Network program members only). ▲▲

Cisco Adds Rich-Media Conferencing to IP Communications Portfolio

CISCO HAS ACQUIRED LATITUDE Communications, Inc. (latitude.com) of Santa Clara, California.

Founded in 1993, the company has 183 employees and more than 400 customers worldwide, and is a leading provider of enterprise conferencing products with its MeetingPlace audio and Web conferencing solution. The acquisition advances Cisco's leadership in IP communications by adding rich-media conferencing, which combines voice, video, and Web conferencing, to the Cisco AVVID (Architecture for Voice, Video and Integrated Data) product portfolio.

The Latitude acquisition will enable Cisco to speed the delivery of intelligent multimedia conferencing solutions that take advantage of dynamic network information—such as presence and location data about network users—to improve workplace productivity. Basing these products on industry standards,

such as Session Initiation Protocol (SIP), H.323, and Extensible Markup Language (XML), ensures solutions that integrate voice and video conferencing, as well as emerging technologies such as instant messaging and data collaboration.

Latitude MeetingPlace currently integrates with leading enterprise desktop scheduling applications such as IBM/Lotus Notes and Microsoft Outlook, as well as with data collaboration and instant messaging solutions such as IBM/Lotus Sametime. Latitude MeetingPlace also offers significant integration with Cisco CallManager, enabling users to schedule, attend, and manage meetings using the display on Cisco IP phones. Cisco and Latitude also intend to integrate MeetingPlace with Cisco IP/VC for video-conferencing capability.

Latitude's business will become part of Cisco's Voice Technology group. ▲▲

Tech Tips & Training

Combating Internet Worms

Q&A from Cisco Webcast delivers actionable advice on combating worms.

FOLLOWING IS A SAMPLING OF the top 100 questions and answers generated during Cisco's popular September 2003 Webcast seminar, "Combating Blaster and Other Internet Worms." To see all 100 questions and answers, visit cisco.com/packet/161_4a1.

Network-Based Application Recognition (NBAR) was effective against Code Red. Is it appropriate for use against current worms?

NBAR is an effective tactical tool to block malicious packets while you are patching, or otherwise establishing, defenses against a worm. It does require some type of match value, which is unique to the worm. For example, for Code Red, we can use an HTTP match on default.ida. With Blaster, we look for SQL packets of a specific length.

A feature built into Cisco routers, NBAR allows traffic to be marked based upon application- or service-specific bases, and then dropped, shaped, and policed using various quality of service (QoS) or access control list (ACL) mechanisms.

Is there a way in Cisco IOS® Software to filter Internet Control Message Protocol (ICMP) packets containing specific payload data?

Assuming that you are using an extended ACL, you can create ACL entries for specific ICMP message types.

Network Address Translation (NAT) shows ICMP on port 1024 (and other ports). How do I find what ICMP type that is?

You cannot tell the ICMP type from the output of `show ip nat translation`. The output from that command lists only the protocol. The port that is shown in the output is actually an identifier within the ICMP

packet that does not directly correspond to the ICMP type, such as echo or echo-reply.

How do you protect your network from infected machines connected through a virtual private network (VPN)?

The best thing to do with a VPN is to terminate the tunnel in front of the firewall and then block the ports described in the Cisco Product Security Incident Response Team announcement—135, 444, and UDP 69. This should prevent the worm from spreading further. Make sure to apply the Microsoft patches.

Using Cisco, is there a way to register machine addresses on the network, so that new machines that attempt to connect to the network are prevented from doing so?

There are multiple answers to this question. You can use port security on switches to lock specific MAC addresses to specific ports. There are several recent features within Cisco Catalyst® switches to prevent other Layer 2 attacks, such as Dynamic Host Control Protocol (DHCP), General Attribute Registration Protocol (GARP), and Spanning Tree Protocol (STP) attacks. But, you may wish to consider IEEE 802.1X authentication, which requires hosts and users to authenticate themselves to their network ports (or access points in the case of wireless devices) prior to virtual LAN (VLAN) assignment, thus enabling the port to forward traffic to the rest of the network.

How can we tell if backdoors have been set up on our systems?

There are various software packages that can determine if a backdoor has been set up on your system. One simple way is to use a port scanner to scan your system

and look for any open ports that are unusual. This approach requires a good knowledge of the ports that should normally be open on your system. Other software to consider is Nessus, which can identify potential backdoors. You can also investigate the chkrootkit open source tool at chkrootkit.org.

What effect does putting ACLs on your peripheral routers have on your network if the traffic going through them is too high?

If the traffic rate is too high, the CPU usage on the router will increase, which slows down the router. In most cases, this doesn't happen because ACL code runs in ASICs on the input queue. If you have slowing, you need to apply NBAR from your Internet service provider to rate-limit traffic.

In terms of mitigation methodology, should you identify the worm as step one, before moving to containment?

No. You should contain the infection as quickly as possible. Identification of the worm is the next step, once you have stopped the infection from spreading in your network or beyond.

Would diligent patching procedures have helped prevent Blaster and its variants from spreading? What products, aside from Windows Auto Update, can help with patching?

Diligent patching is always recommended, but is often administratively cumbersome. If only one new patch comes in each week, can your network productivity sustain the downtime required to apply and reboot the patch? One strong argument toward implementing Cisco Security Agent is its ability to manage and modify profiles as they happen, to protect servers

from time-zero, while allowing you to adopt a routine, rather than reactive, patching schedule.

We used multilayer switching flows to monitor ICMP traffic and track the Nachi worm, but this is a manual process. Is there a way to automate this?

Cisco NetFlow, in conjunction with tools or products such as Arbor Networks' (a Cisco development partner) anomaly detection system, is useful in this arena. For more information, see cisco.com/packet/161_4a2.

I have Cisco 4000, 2500, and 1600 Series routers and Cisco IOS Software Release 11.3 and 12.1.

Can I use any features built in to these routers to prevent worms?

Yes. Basic access lists and rate limiting are your best options. If the Cisco IOS Software release you have doesn't support rate limiting or NBAR, try to negotiate with your service provider to limit the traffic for you. For applying router-based fixes, see cisco.com/packet/161_4a3 and cisco.com/packet/161_4a4.

Would I want an ACL on my edge router, as well as the core routers?

You want an ACL on your edge routers. You generally want to place ACLs close to the traffic source. Core routers should be configured to route traffic as quickly as possible, and ACLs might interfere with that.

Can QoS be used to prevent a Blaster type of worm from debilitating the network?

QoS could possibly be used, but for faster recognition of bad traffic and less impact on your CPU, NBAR and rate limiting are a better option.

Do you need to buy CiscoWorks VPN/Security Management Solution (VMS) to use all Cisco Security Agent functions?

Cisco Security Agent is a centrally managed product and is not standalone. The management console that you will need to control it with is through CiscoWorks VMS.

What is a good resource to learn about port numbers and what they do?

A good place to start is the registry on the

Internet Assigned Numbers Authority (IANA) Website: iana.org/assignments/port-numbers.

Is it true that VPN tunneling can bypass a hardware firewall, so you could be protected from the Internet but infected by someone inside the company over the tunnel?

Yes, Cisco Security Agent works with the Cisco VPN client via "Are You There" (AYT). You can configure the VPN not to enable the tunnel if Cisco Security Agent is not present on the remote system. Cisco Security Agent will protect the remote system.

If you block ports 133, 135, etc., will there be any impact on Windows active directory?

Yes. It is imperative to only filter these ports when there is normally no business need for them to exist. To mitigate worms such as Blaster in cases where these ports must be open, other technologies, such as antivirus and HIPS, must be used.

Can Blaster be transmitted through a VPN connection?

Yes. You want to ensure that you have protection on the endpoint if you are going to allow it in via VPN.

In what way does a worm work differently than a virus? Which one is more deadly?

A virus is a piece of code that attaches itself to another document or program and executes when that document or program is opened. A worm is a typically self-contained program that can infect other systems on its own and then copy itself over and continue the infection. Like their biological equivalents, viruses require "vectors"—something to carry them from one system to another. As to which is more deadly, that depends on the action each takes. Some viruses are coded to erase hard drives once they are activated. Worms usually don't want to inflict too much damage on a system because they need the system as a platform to continue infecting other systems. However, as in the case of SQL Slammer, the worm's infection rate could be so fast that it causes link congestion problems.

It is not practical for ISPs to block ports. How can such organizations deal with worm attacks?

Blackholing or null-routing, including Border Gateway Protocol (BGP)-triggered uRPF blackholing (which allows blackholing based on source address), are techniques that scale to large ISP networks. *Cisco ISP Essentials*, by Barry Greene and Philip Smith, is an excellent resource in this arena. For more on the book and other ISP resources, visit ispbook.com.

Is there any way to identify worms by bit patterns?

Yes. Each worm exploits a particular vulnerability on a system. If you look for the byte patterns of that exploit in your intrusion detection system (IDS) signatures, you will be able to identify a worm and have an alarm raised on your monitoring console.

How can we get the most current security information from Cisco during a virus event?

The Cisco Product Security Incident Response Team provides the most current information regarding network security issues and fixes as they apply to Cisco products. For more details, see cisco.com/packet/161_4a5.

Would some of these worms cause your computer to try to connect to the Internet by itself and keep you from downloading antivirus updates?

This definitely would cause the computer to connect to the network by itself as it infects other machines. We haven't seen the Blaster worm stop anyone from downloading viruses, but there may be a variant that does this. If you think you have this problem, the best thing to do is delete `nsblast.exe` and then also delete any reference to `msblast` from your registry.

How is it possible to differentiate between SYN packets from a distributed DoS attack and valid SYNs from browsers and other TCP applications?

The use of Cisco IDS and Cisco NetFlow, in conjunction with an anomaly detection system such as Arbor Networks Peakflow products, allows detection of many types of DoS traffic, including SYN floods. ▲▲

Wireless Security Update

How to Combat New Attacks on IEEE 802.11-Based WLANs

NEW SECURITY EXPOSURES in certain IEEE 802.11-based wireless LAN configurations surfaced last year. Cisco recommends various fixes for these threats, which include the following:

- Offline dictionary attacks on password-based authentication algorithms that use the Microsoft Challenge-Handshake Authentication Protocol (MS-CHAP). One such algorithm is Cisco LEAP.
- Denial of service (DoS) attacks against Cisco IOS® Software-based wireless devices caused by repeatedly sending malformed URLs that force devices to continually reboot.
- Exposure of encryption keys in plain text in networks that use static 802.11 Wired Equivalent Privacy (WEP) encryption keys

instead of dynamic, rotating Temporal Key Integrity Protocol (TKIP).

Let's examine each of these threats and their associated fixes.

Offline Dictionary Attacks

Most MS-CHAP password-based authentication algorithms, such as Cisco LEAP, are vulnerable to dictionary attacks in the absence of a strong password policy. During offline attacks, an intruder captures the MS-CHAP challenge-response messages exchanged between the user and the access network. Then, the attacker tries to match these messages against those in a precomputed dictionary to discover the user's password and use those credentials to gain unauthorized network access.

Amid the various handshaking steps during authentication, the wireless client uses MD4 hashing on an access point's 8-byte challenge to its request for authentication. This produces a 16-byte hash, which is padded with five nulls to produce 21 bytes.

The latest offline dictionary-attack tools focus on this step. The last 2 bytes of the hash are padded with nulls, so there are only two characters to decipher in the last 7 bytes. For this dictionary tool to work, though, the word must be in the attacker's dictionary and the MD4 hash precomputed. So using non-dictionary passwords greatly mitigates your vulnerability.

Cisco recommends using strong passwords and changing them at least every three months.

Strong passwords have the following characteristics:

- They contain at least 10 characters
- They mix uppercase and lowercase letters
- They contain at least one numeric character (0-9) or non-alphanumeric characters (example: !#@&)
- They contain no form of the user's name or user ID
- They are not words found in any dictionary
- They are randomly generated or created by using the guidelines above

HTTP GET Patches and Workarounds

Certain releases of Cisco IOS Software for Cisco Aironet® 1100 and 1200 Series access points and Cisco Aironet 1400 Series bridges (see table) allow remote attackers to reboot devices by sending a specially crafted URL. Repeated exploitation of this HTTP GET command can lead to prolonged service interruption.

The vulnerability has been fixed in Cisco IOS Software Release 12.2(11)JA1 or later. Cisco is offering free software upgrades to address these vulnerabilities for all affected customers. Visit the Cisco Software Center to download this upgrade: cisco.com/packet/161_4c1.

There are also two workaround alternatives. One is to use `access-class` or `access-list` commands to limit access to legitimate hosts. Another is to disable HTTP and use Secure Shell (SSH) to administer the Cisco Aironet device. There are many free and commercial versions of SSH software available.

Here is an example of using `access-class` commands:

```
ap(config)# ip http access-class 10
```

```
ap(config)# access-list 10 permit host 10.0.0.1
```

Only host 10.0.0.1 is allowed to access the wireless device. All other hosts are prohibited.

Similarly, here is an example of mitigating the exposure by configuring access control lists (ACLs):

```
ap(config)# access-list 111 permit tcp host 10.0.0.1 host 10.0.0.50 eq www
ap(config)# access-list 111 deny tcp any host 10.0.0.50 eq www
ap(config)# access-list 111 permit ip any any
```

In this example, the host 10.0.0.1 is the only one allowed to access the device at 10.0.0.50 on port 80. No other traffic destined to port 80 on the access point is allowed through. You must change host IP addresses and the ACL number to suit your configuration. This ACL must be applied to all interfaces and block all IP addresses assigned to the affected device.

To use the other workaround—disabling HTTP and enabling SSH—use this example:

```
ap(config)# no ip http server
ap(config)# ip domain name <your-domain>
ap(config)# crypto key generate rsa
```

The name for the keys will be `ap.your-domain`.

Choose the size of the key modulus in the range of 360 to 2048 for your general-purpose keys.

```
How many bits in the modulus [512]:
1024
% Generating 1024 bit RSA keys ...[OK]
ap(config)# line vty 0 4
ap(config-line)# transport input ssh
```

Now you can connect to the Cisco Aironet device using SSH client software.

Keeping Encryption Keys Private

Encryption keys in IOS-based Cisco Aironet networks are also potentially vulnerable under the following circumstances:

- You have a Simple Network Management Protocol (SNMP) server deployed
- You are using static WEP encryption keys
- You have enabled an option called `snmp-server enable traps wlan-wep`. Note that this option is disabled by default on Cisco Aironet products.

The `wlan-wep` trap notifies the SNMP server of events relating to WEP keys. Cisco Aironet devices will also transmit the values of any static WEP keys being used on the network as clear text to the SNMP server in the trap message when the WEP key is changed, or when the access point is rebooted.

An attacker intercepting the SNMP traffic could obtain WEP key values stored on the access point and snoop on encrypted wireless communications.

Cisco recommends installing its patch for Cisco IOS Software Release 12.2(13)JA1 as soon as possible. Customers unable to get the patch can disable the `snmp-server enable traps wlan-wep` option or switch to another encryption method, such as Cisco TKIP (CKIP) or WPA TKIP. ▲▲

FURTHER READING

- **Dictionary Attacks on Cisco LEAP - Cisco TAC Notice:**
cisco.com/packet/161_4c2
- **Dictionary Attacks on Cisco LEAP - Cisco Aironet Notice:**
cisco.com/packet/161_4c3
- **Cisco Security Advisory: HTTP GET Vulnerability in AP1x00:**
cisco.com/packet/161_4c4
- **SNMP Trap Reveals WEP Key in Cisco Aironet Access Point:**
cisco.com/packet/161_4c6
- **Cisco SAFE: Wireless LAN Security in Depth:**
cisco.com/packet/161_4c7
- **Exploiting Known Security Holes in Microsoft's PPTP Authentication Extensions (MS-CHAPv2):**
cisco.com/packet/161_4c8

HTTP GET VULNERABILITY: AFFECTED PRODUCTS

Hardware Model	Cisco IOS Software Release(s)
Cisco Aironet 1100 Series Access Point	12.2(4)JA, 12.2(4)JA1, 12.2(8)JA, 12.2(11)JA
Cisco Aironet 1200 Series Access Point	12.2(8)JA, 12.2(11)JA
Cisco Aironet 1400 Series Wireless Bridge	12.2(11)JA

Cisco recommends a free upgrade to Cisco IOS Software Release 12.2(11)JA1 to circumvent wireless DoS attacks.


Why Should I Care About Identity?

The majority of unauthorized access to traditional networks and resource misuse comes from internal sources. The ability to identify users and devices attempting to access the corporate network is the first step of any security solution. In addition to using identity-based networking to solve security concerns and allow network access, validating the identity of users and devices enables network administrators to provision services and allocate resources to users based on their job functions.

What Are the Problems to be Solved?

A comprehensive network security policy must keep the outsiders out and the insiders honest. The policy should:

- Prevent external hackers from having free rein in the network.
- Allow only authorized users into the network.
- Prevent network attacks from within.
- Provide different layers of access for different types of users.



To be truly effective, the security policy must accomplish these goals in a way that does not disrupt business or make authorized access prohibitively difficult.

What Is IEEE 802.1X?

IEEE 802.1X is a standard for authenticating network clients (or ports) on a user ID or device basis. The 802.1X standard applies to end devices and users (known as supplicants) trying to connect to ports and other devices, such as a switch or access point (the authenticator). Authentication is achieved with backend communication to an authentication server, such as a Cisco Secure Access Control Server (ACS).



Extensible Authentication Protocol


Extensible Authentication Protocol (EAP) is a flexible protocol used to carry authentication information. The authentication information can include user passwords or predefined security keys. EAP typically rides on top of another protocol, such as 802.1X or Remote Authentication Dial-In User Service (RADIUS), which carries the authentication information between the client and the authenticating authority.

IDENTITY-BASED NETWORKING At a Glance


What Does Identity Do for Me?

Once you know who is on the network, you can apply policies on a per-user basis. This approach provides a solid, comprehensive security solution that enhances the usability of the network. Following are examples of the advantages of an identity-based security solution.

Prevents Unwanted Access



Limits Access to Networked Resources



Without Identity
Access to human resources (HR) databases and other sensitive content is open to all employees


With Identity
Using 802.1X with extensions, you can specify which networked resources the user may access (for example, only managers have access to HR information)

User-Based Service Provisioning



Working with Authentication Servers

IEEE 802.1X is only half of the identity story. The information carried by 802.1X must be authenticated by an authentication server. This can be done with name and password validation using a RADIUS or TACACS+ server, or with digital signatures confirmed by a third-party validation service such as Public Key Infrastructure (PKI).



RADIUS

RADIUS is a protocol used to communicate between a network device and an authentication server or database. RADIUS allows a network device to securely pass communication of login and authentication information (username/password), as well as arbitrary value pairs using vendor-specific attributes (VSAs). RADIUS can also act as a transport for EAP messages. In addition to the protocol, RADIUS also refers to the actual server.

Public Key Infrastructure

Public Key Infrastructure, or PKI, is a method of providing identity authentication between two parties via a trusted third party. A PKI certificate is "proof" of identity, signed by the trusted third party. It is the network equivalent of having a valid passport that is trusted by the customs agents of all other countries. Just as a passport is signed by the passport office, stating your verified identity and citizenship, a PKI certificate is signed by a Certificate Authority stating your verified identity and network associations. And unlike passports, PKI certificates cannot be forged.



Reader

Packet® thanks all of the readers who submitted technical tips this quarter. 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.

Configuration

TIP Shortcut Commands

Cisco IOS® Software provides a shortcut tool that enables you to create aliases for frequently used commands. This saves you time when you configure or troubleshoot. In global configuration mode, use the **alias** command to create a shortcut command. For example, for the command **alias exec sii show up interface brief**, you can create the alias **show ip interface brief by sii**.

You can also create shortcuts for commands in other modes such as global, interface, and line.

— Albert Elias, Synergy Professional Services, Dubai, United Arab Emirates

Editor's Note: This is a useful tip. For documentation, refer to the section "Create and Monitor Command Aliases" in the configuration guide for Cisco IOS Software Release 11.3 at cisco.com/packet/161_4e1.

Routing

TIP Testing IP Routes

Typically, you check learned routes using the **show ip route <ip-address>** command. However, if you have to regularly check for dozens of routes, it is convenient to create a text script consisting of a **show ip route** command for each specific route to be tested. Add the specific subnet mask as a comment to each command line, as follows:

```
show ip route 135.42.16.0!/25
show ip route 135.42.16.128!/25
show ip route 135.42.18.0!/26
show ip route 135.42.18.64!/26
show ip route 135.42.18.128!/26
show ip route 135.42.18.192!/26
```

Then you can easily paste the script into the router Telnet session. Only by doing so can you verify that the route being returned (each returned route entry always lists the subnet mask) is the specific route you are looking for and not, for example, a summarized route being sourced somewhere else in the network by intent. Furthermore, with the test script

available as a flat text file, you can easily reapply the test on it at various locations within the network, ensuring that the routes are truly being distributed and learned as intended in the network design.

— Bart Van Damme, EDS, Terneuzen, The Netherlands

Editor's Note: Instead of using the comment mask with the exclamation point symbol (!), you should specify the mask in the command. For example:

```
show ip route 135.42.16.0 255.255.255.0
```

This way, if there is a summarized route in the network that is not of the specified netmask, it will say "% Network/Subnet not in table" Automating your management tasks is always a good practice. For a useful tool, consult the Expect (TCL) scripts at <http://expect.nist.gov/>.

TIP OSPF Adjacencies

To find out the uptime of Open Shortest Path First (OSPF) adjacencies, use the **show ip ospf neighbor <interface> detail** command. For example:

```
router1#sh ip ospf ne g0/2 det
Neighbor router2.xyz.com, interface address 191.15.255.2
  In the area 10 via interface GigabitEthernet0/2
  Neighbor priority is 1, State is FULL, 24 state changes
  DR is 191.15.255.1 BDR is 191.15.255.3
  Options is 0x42
  Dead timer due in 00:00:08
  Neighbor is up for 1d08h
  Index 5/5, retransmission queue length 0, number of
  retransmission 1
  First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
  Last retransmission scan length is 0, maximum is 1
  Last retransmission scan time is 0 msec, maximum is 0 msec
```

You can filter it for "up for" to get a quick view, as follows:

```
show ip ospf neighbor <interface> detail | include up for
```

— Amer Kaleem, UBS Investment Bank, Chicago, Illinois, USA

Editor's Note: This is a good tip, and you can extend this method to other routing protocols such as IS-IS, EIGRP, IGRP, and RIP, although the commands will be different. Following are the commands you use to find the uptime for other common routing protocols:

BGP: **sho ip bgp summary**

EIGRP: **sho ip eigrp neighbors**

IS-IS: **sho isis neighbors de | inc state changed**

Tips

Submit a Tip

Help your fellow IT professionals and show off to your peers by submitting your most ingenious technical tip to packet-editor@cisco.com. Who knows, you may see your name in the next issue of *Packet*. When submitting a tip, please tell us your name, company, city, and country.

Troubleshooting

TIP Locating Problems in the Network

My network comprises several Cisco Catalyst® 6500 Series Layer 3 switches and many Catalyst 3500 Series Layer 2 switches, all using the Cisco IOS Software. To quickly find the exact location and network port of a workstation that generates problems such as worms, viruses, or loss of connectivity, I use the following method:

1. Check for the workstation's IP address, and obtain its MAC address using the command **sh arp | include 172.16.23.4**. I shorten the command using an **alias** command (such as **alias exec arpsh sh arp | include**). For example, I shorten the above command to **arpsh172.16.23.4**.

```
Cat6k#arpsh 172.16.10.227
Internet 172.16.10.227      3    4c00.1095.1e37  ARPA
Vlan30
```

2. After you have the MAC address, determine from which port the MAC address came. To look for the MAC address and its associated port, use the **alias exec macsh sh mac-address | include** command.

3. Use the **macsh 4c00.1095.1e37** command:

```
Cat6k#macsh 4c00.1095.1e37
30 4c00.1095.1e37  dynamic No  -- Gi2/8
```

Now you know the MAC address comes from Gigabit Port 2/8. Because this leads to the associated Cisco Catalyst 3500 Series Switch, you must connect to the switch using Telnet.

4. To find the IP address of the switch, use the **sh cdp neigh gi2/8 det** command:

```
Cat6k#sh cdp neigh gi2/8 det
-----
Device ID: XX-cat3k
Entry address(es):
  IP address: 172.16.163.8
Platform: cisco WS-C3524-XL, Capabilities: Trans-Bridge
Switch
Interface: GigabitEthernet2/8
<output suppressed>
```

5. Telnet to the switch as follows:

```
Cat6k#telnet 172.16.163.8
Trying 172.16.163.8 ... Open
Password:
XX-cat3k>en
Password:
XX-cat3k#
```

6. Configure the alias in Step 2 in the switch, and use the **macsh** command to find the problem port. As shown in the following result, the workstation is located in port Fa0/13.

```
XX-cat3k#macsh 4c00.1095.1e37
4c00.1095.1e37      Dynamic      30 FastEthernet0/13
```

—Affan Basalamah, Institut Teknologi Bandung, Bandung, Indonesia

Tech Tips

Manage Cisco devices with the SoftWare Image Manager (SWIM) tool of CiscoWorks Resource Manager Essentials.

The SWIM tool is useful for software management and upgrade for most Cisco devices, including the Network Analysis Module (NAM) for Catalyst® 6000 switches. This document discusses issues related to NAM upgrades and provides tips on how to use SWIM to upgrade your NAMs.

cisco.com/packet/154_4d1

Learn about router resources available on the Cisco Technical Assistance Center (TAC) Website. Sign up for the free Cisco TAC Web seminar, "Using the Cisco TAC Website for General Router Issues," or view an on-demand session. cisco.broadshow.com/tac/ (available to all users, but this standalone site requires registration separate from Cisco.com)

Understand the most common problems related to Active Directory integration with Cisco CallManager. This troubleshooting guide addresses common issues in the field,

including installing the Active Directory plug-in, Cisco CallManager and user pages with Active Directory integration, and application problems related to Active Directory. cisco.com/packet/154_4d2

Get information about the latest Cisco TAC training. TAC Training offers the same rigorous courses that Cisco TAC engineers take as part of their ongoing technical development. Training sessions include course materials related to the implementation of various debugging tools, as well as the necessary troubleshooting steps to resolve networking issues. The sessions are short, self-paced instruction and practice modules that use audio and video on demand, animation, and hands-on simulation labs to explain technical concepts. cisco.com/packet/154_4d3

For more Tech Tips from the Cisco TAC Website, visit cisco.com/public/support/tac/.

Technology

Locking Down IOS

Secure networks start with protecting routing devices.

THE NETWORK LANDSCAPE IS GETTING riskier as hackers continue to grow more cunning. At the same time, communications networks—including the “untrusted” public Internet—have become inextricably tied to day-to-day business operations.

Because business dependence on networks is greater than ever, it's essential to reinforce the networking software infrastructure of private and public networks against attacks. Because Cisco IOS® Software forms the foundation of so many communications networks, Cisco has embarked on a formal, multipronged program to further protect the network operating system.

One aspect involves strengthening the operating system with more stringent testing, software-development training, and engineering procedures (see sidebar, “Cisco Bolsters Internal IOS Development Efforts”). At the same time, Cisco continually releases enhanced IOS security features to help network operators combat risks that might threaten their networks (see chart, “Key IOS Infrastructure Security Enhancements,” page 19).

Let's take a look at a few recent IOS security enhancements that network operators can use to fight attacks today.

One-Touch Security Best Practices

The increased complexity of networks combined with the growing number of possible threats adds to the number of steps required to configure a secure network device. As a consequence, the likelihood of an incomplete or incorrect configuration is greater. Such errors can lead to a network exposure.

To ease configuration tasks and thus reduce errors, Cisco has incorporated a “one-touch” device lockdown process into Cisco IOS Software. Called Cisco *AutoSecure*, this tool enables network operators to quickly secure a device without having thorough knowl-



edge of Cisco IOS Security features. AutoSecure applies best-practice recommendations from Cisco experts and the US National Security Agency (NSA), explains Charles Goldberg, IOS security product line manager.

Router configuration generally includes setting proper device parameters, creating filters, and enabling and disabling certain services to protect a router or switch's forwarding, control, and management planes. AutoSecure, a command-line interface (CLI)-based tool, automates these tasks.

In addition, *Cisco Secure Device Manager (SDM)*, a Web-based management tool embedded in Cisco routers, has a one-touch security lock-down and auditing feature accessible through a GUI. By clicking one button, network administrations can check a configuration against an AutoSecure configuration for compliance.

Cisco Bolsters Internal IOS Development Efforts

Cisco has modified its internal development processes to ensure that the base IOS code is less vulnerable to exploitation. The internal program has multiple facets:

- **Proactive vulnerability testing.** Rigorous testing is at work to unearth vulnerabilities to common attacks prevalent in the industry. This effort includes software vulnerability assessment testing, as well as system testing targeted at hardening IOS features.
- **IOS developer training.** Training courses and software development guidelines to edu-

cate the engineering community about secure design considerations and coding techniques inherently thwart many potential attacks and prevent vulnerabilities waiting to be discovered by the hacker community.

- **Refined engineering processes.** Cisco has developed testbeds that look for security vulnerabilities. Developers run their new IOS software scripts against the testbed code. This process is now a formal part of Cisco's general engineering methodology to eliminate security exposures in the initial development process.

- **Forwarding plane protection.** The forwarding plane is responsible for switching packets across a router or switch. Attacks on the forwarding plane can flood a device and network with malicious traffic and affect network services. To minimize forwarding-plane attacks, AutoSecure automatically activates features such as anti-spoofing and the Cisco IOS Firewall (in cases where devices run IOS Firewall-enabled images).

AutoSecure also creates access control lists (ACLs) that block bogon and IETF RFC 1918 addresses. *Bogon* IP addresses are those that haven't been assigned by a formal Internet address registration authority, such as the Internet Assigned Numbers Authority (IANA). It is safe to block traffic from those addresses because the addresses by definition aren't in legitimate use.

RFC 1918 IP addresses are those reserved for private intranets and extranets. So networks not part of those private networking communities should not be receiving traffic from RFC 1918 addresses.

- **Control/management plane protection.** The control and management planes in a router are responsible for processing incoming packets and for exchanging router-reachability information using standard routing protocols. They are also responsible for Simple Network Management Protocol (SNMP) management and Telnet remote-access services. Because the control and management planes are integral to router functions, any compromise to their security could result in a network outage.

To prevent this, AutoSecure automatically disables management features that are typically unnecessary or unused. For example, AutoSecure takes the following best-practice steps:

Disables unnecessary global services—Among these are Finger user lookup services; legacy TCP and UDP

small-server features such as echo, chargen, and discard; the BOOTP broadcast protocol; and HTTP.

Disables often unnecessary, per-interface services—These include Proxy Address Resolution Protocol (ARP) and Internet Control Message Protocol (ICMP) redirects, ICMP "unreachable" messages, and ICMP mask reply messages.

Enables services that help secure necessary global services—Among these are password encryption, TCP synwait-time, and TCP keepalives.

Adds secure access to the router—Automatic functions include securing communication between routers and management devices using Cisco Secure Shell Version 2 (see chart, page 19) and Cisco Secure Copy.

AutoSecure can be deployed in one of two modes. *Interactive mode* prompts the network administrator with options to enable and disable different services and security features. *Noninteractive mode* automatically executes the AutoSecure command with recommended Cisco settings.

By allowing network operators to automatically set security parameters with a single command at the Cisco CLI or with an SDM button click, Cisco AutoSecure reduces the configuration complexity, time, and errors associated with securing network device access and services.

Protecting the Route Processor

Control Plane Policing is a new feature that protects against overloading a device's control plane with unnecessary traffic. This feature combats denial-of-service (DoS) attacks that flood a device's route processor (RP) with traffic. Under attack, a high rate of packets presented to the RP could force the network control plane to spend too much time processing malicious traffic. This

unnecessarily takes processing cycles away from legitimate production traffic. Worms and other rapidly replicating applications also pose a threat, says Michael Lin, product manager for Cisco IOS Software. As worms scan and target large blocks of IP address space, they often generate traffic that requires processing by the control plane. “Gone unchecked, this traffic can severely reduce the performance of a network device,” Lin says.

To tame malicious control plane or management plane traffic, Control Plane Policing works with Cisco quality-of-service (QoS) classification mechanisms to identify, then limit or drop, specified traffic types when they are above a given threshold. The network operator sets the threshold and specifies the parameters for how much to limit a traffic flow or whether to drop it entirely. So even during an attack, the RP is protected to keep the network available.

Configuring Control Plane Policing involves four distinct steps:

1. Define a packet classification criterion via the Cisco Modular QoS CLI (MQC) interface. MQC is a common-language framework for simplifying QoS configuration. Below are the CLI commands for defining packet-classification criteria:

```
router(config)#class-map <traffic_class_name>
```

```
router(config-cmap)#match <access-group | protocol>
```

2. Define a service policy:

```
router(config-pmap)#policy-map <service_policy_name>

router(config-pmap)#class <traffic_class_name>

router(config-pmap-c)# police cir <rate>
conform-action transmit

exceed-action drop
```

3. Enter control-plane configuration mode:

```
router(config)#control-plane
```

4. Apply QoS policy:

```
router(config-cp)# service-policy {input|output} <service_policy_name>

input Assign policy-map to the input of an interface

output Assign policy-map to the output of an interface
```

The show access-lists command displays a count, per

KEY IOS INFRASTRUCTURE SECURITY ENHANCEMENTS

Cisco IOS Feature	Description	IOS Version Support
AutoSecure	Automatically disables exploitable IP services and activates IP services that help defend a device or network under attack from malicious levels of traffic that can cause DoS attacks	12.3 Mainline, 12.2(18)S
Control Plane Policing	Protects the route processor from unnecessary or malicious levels of traffic, including DoS attacks	12.3(4)T, 12.2(18)S
Silent Mode	Suppresses response messages from the router's control plane to limit network reconnaissance information available to hackers	12.3(4)T
Raw IP Traffic Export	Allows copies of inbound and outbound packets to be sent out a LAN interface to efficiently capture packets with analysis or IDS tools	12.3(4)T, 12.2(22)S
Login Enhancements—Password Retry Delay	Delays potential dictionary attacks and provides other ways to thwart unwanted device access	12.3(4)T, 12.2(22)S
Image Verification	Replaces manual process of validating the integrity of all downloaded IOS software images with an automated method	12.3(4)T, 12.2(18)S, 12.0(26)S
Memory Threshold Notifications	Mitigates low-memory router conditions by sending alerts when available memory has fallen below a configured threshold	12.3(4)T, 12.2(18)S, 12.0(26)S
CPU Threshold Notification	Triggers a syslog notification when a specified percentage of CPU resources for a given process exceeds or falls below a certain threshold for a configured time period	12.3(4)T, 12.2(22)S, 12.0(26)S
Secure Shell Version 2 (SSHv2)	Enhances previous versions of SSH for remote network management by concealing password length, making dictionary attacks more difficult. Resolves SSHv1 vulnerability to man-in-the-middle attacks during user authentication	12.3(4)T, 12.1(19)E

SOFTWARE REINFORCEMENTS. Some IOS features auto-protect devices while others simplify the configuration of security parameters, making security more likely to be implemented effectively.

ACL entry, when traffic matches a particular class. This can be used to ensure that traffic is being classified as desired. The **show policy-map control-plane** command helps verify Control Plane Policing configuration and provides dynamic information about the actual policy applied. This includes rate information and the number of packets (and bytes) that conformed or exceeded the configured policies.

*CBQoS*MIB, the primary accounting mechanism for MQC based policies, provides SNMP MIB support for monitoring and managing Control Plane Policing. This feature will be available in March 2004 in Cisco IOS Software Release 12.3(7)T, as well as in the 12.2S train in the second calendar quarter of 2004.

Keeping Hackers in the Dark

One requirement for hacking a system is reconnaissance—gaining information about the network. Hackers conduct reconnaissance by requesting system messages, such as the status of packet delivery. These messages provide information such as a device's IP address.

To reduce the amount of information a hacker can gather about a network, the *Silent Mode* feature in IOS stops certain informational packets from being generated by the router. For example, ICMP messages and SNMP traps that are normally generated by the router are suppressed. Silent Mode leverages the familiar MQC.

Enhanced Offline Monitoring

The *Raw IP Traffic Export* feature allows a router to export unaltered IP packets to LAN or virtual LAN (VLAN) interfaces. Now monitoring probes on routers no longer need to be inline with the network device. This feature eases the deployment of protocol analyzers, intrusion detection systems (IDS), and other monitoring devices that can detect attack behavior on the network.

IP Traffic Export allows the placement of monitoring devices next to any device in the network with an Ethernet interface or to direct all exported traffic to a VLAN dedicated to network monitoring. This setup enables more network segments to be monitored for security purposes using a single device.

Reinforced Firewall

The IOS security features mentioned here lock down a device to attack. In addition to these security measures, it is important to carefully check traffic in transit to avoid unauthorized network entry and data theft. To this end, the Cisco IOS Firewall recently gained enhancements. The IOS Firewall is a stateful traffic-filtering access-router software feature that inspects data flows for protocol conformance and violations. The software, which performs at WAN interface speeds of up to 198 Mbit/s, has gained the following:

- **Additional voice protocol support.** Stateful filtering of Session Initiation Protocol (SIP) and Skinny

Client Control Protocol (SCCP) has been added to the firewall's support of H.323 Version 2 traffic. Left alone, voice protocols dynamically open available ports for control channels. If inspection is not performed on actual protocol loads, attackers can hijack a session. However, the Cisco IOS Firewall allows only specific ports to be opened, particular to a single session, to keep the network safe.

- **Authentication proxy capabilities.** The router locally challenges an action before granting permission for it to occur. In the case of Proxy HTTPS—the HTTP protocol encapsulated in Secure Sockets Layer (SSL)—for example, the router presents a local challenge to a user's PC for a user ID and password (which are encrypted) before allowing that user to create a tunnel to an outside network. Similarly, the Cisco IOS Firewall supports authentication proxies for Telnet and FTP applications.

- **ICMP message checks.** The Cisco IOS Firewall can be programmed to selectively allow certain ICMP messages through the router. For example, the firewall would likely be programmed to block messages that might provide a hacker with information about network topology.

- **Content/URL filtering.** Organizations that want to exclude access to certain Internet content or sites need a way to program their WAN access routers to selectively block access to URLs. URL filtering capabilities can be enabled in the firewall in two ways: by manually programming the router to block certain specific URLs or by using the Websense and Secure Computing/N2H2 URL database software to support larger-scale filtering.

◆ ◆ ◆
These IOS security steps represent an ongoing effort for strengthening security in the network software foundation. Cisco is reinforcing the inherent security in the operating system with more stringent testing, development, and engineering practices. In addition, users are advised to leverage the latest features to protect networking devices against threats already lurking. These steps include mitigating DoS attacks, suppressing reconnaissance information available to would-be hackers, and closing the management plane to outsiders. ▲▲

FURTHER READING

- **Cisco IOS Software Release 12.3(4)T: New Security Features and Hardware:**
cisco.com/packet/161_5a1
- **Cisco AutoSecure data sheet:**
cisco.com/packet/161_5a2
- **National Security Agency guide to locking down Cisco routers:**
nsa.gov/snac/cisco/index.html

Mitigating Network DoS Attacks

IOS offers built-in protection for dynamic routing protocols.

BY ZAHEER AZIZ AND TIM SZIGETI

NOT ALL DENIAL OF SERVICE (DOS) attacks are designed to overload servers; some target the network infrastructure itself. These types of attacks deny service by saturating link bandwidths, exhausting router and switch CPUs, or spoofing control plane traffic.

Link saturation and CPU exhaustion DoS attacks can deny dynamic routing protocols the bandwidth necessary to maintain neighbor relationships. When a router loses a neighbor, it flushes all the routes learned from the downed neighbor. It then must switch traffic destined to those flushed routes to an alternate route or drop the traffic all together. Because it has to continually recompute new routes—withdrawing downed routes and updating new ones—the CPU’s resources are eventually exhausted. Either way, the result is denial of service.

DoS attacks that spoof control plane traffic hijack dynamic routing protocol traffic and maliciously reset neighbor relationships or update neighbors with false information, again resulting in denial of service.

Following are solutions to proactively mitigate such detrimental scenarios—using tools already available in Cisco IOS® Software.

Link Saturation/CPU Exhaustion Attacks

Cisco IOS has long supported an unparalleled set of QoS features. However, few customers have realized the advantages of enabling QoS in the context of mitigating DoS attacks. In this role, QoS technologies become absolutely critical to maintaining network availability and security.

Cisco defines 11 classes of traffic that require specific service levels: routing, interactive video, streaming video, mission-critical data, call signaling, transactional data, network management, bulk data, scavenger, and best effort. This doesn’t mean that organizations must implement an 11-Class QoS model today. With QoS design, a proactive approach (explicitly protecting your important traffic) is more effective than a reactive approach (trying to identify and squelch bad traffic). Therefore, the first step in deploying QoS to protect a network against DoS attacks would be to explicitly protect voice, routing, call signaling, and mission-critical data traffic at a minimum.

Interior gateway protocols such as RIP, OSPF, and EIGRP are protected by an internal Cisco IOS mechanism called PAK_Priority. PAK_Priority marks these protocols to DSCP CS6 and gives them preferential

treatment, which varies slightly according to platform. PAK_Priority also marks BGP traffic to CS6, but doesn’t give it any preferential treatment beyond this marking. Whenever BGP is deployed it is recommended to explicitly provision class-based weighted fair queuing (CBWFQ) for routing traffic. CBWFQ provides user-defined traffic classes based on match criteria including protocols, access control lists (ACLs), and input interfaces.

Scavenger-class traffic is of special interest in our discussion. Scavenger class is based on an Internet2 draft outlining a “less than best effort” service, identified by the DSCP value CS1 (DSCP 8). Non-business, entertainment-oriented applications such as KaZaa and Napster, as well as gaming traffic, are well suited to such a service class. Scavenger traffic will be permitted as long as all other more important classes are being adequately serviced. In the event of congestion, the scavenger class is the first to be dropped and squelched.

A sample QoS policy configuration that can be used on WAN edges, CE/PE, or within a service provider core is shown below. Note: the best effort class (class default) must include a “bandwidth” statement in order to protect it from scavenger traffic. It is assumed that such traffic has already been correctly classified and marked to the recommended DSCP values. For more details on these classes and recommended markings, refer to the “Service Provider Quality of Service Design Guide” at cisco.com/packet/161_5b1.

Example 1: LLO/CBWFQ for WAN Edge, CE/PE, or as DiffServ within Service Provider Core

```

policy-map EDGE
class VOICE
  priority percent 33      ! Voice gets 33% LL
class ROUTING
  bandwidth percent 3     ! Minimal BW guarantee
                           for Routing traffic
class CALL-SIGNALING
  bandwidth percent 3     ! Minimal BW guarantee
                           for Call-Signaling
class MISSION-CRITICAL-DATA
  bandwidth percent 35    ! Mission-Critical class
                           gets 35% BW guarantee
  random-detect           ! Enables WRED for
                           Mission-Critical Data
class

```

Continued on page 23

To learn more about the built-in security features of Cisco IOS Software, and to view other configuration examples, check out the “Cisco IOS Security Configuration Guide” at cisco.com/packet/161_5b2.

DoS Attacks, Continued from page 21

```
class SCAVENGER
  bandwidth percent 1      ! Scavenger class will
                           choke first

class class-default
  bandwidth percent 25    ! Class-Default gets 25%
                           min BW guarantee

  random-detect           ! Enables WRED on class-
                           default

!
```

From the model policy above, voice, call-signaling, routing and mission-critical data will all be protected in the event of a DoS attack. DoS traffic will likely be marked DSCP 0 (provided trust boundaries have been correctly set and remain uncompromised). Therefore, with such policies alone, the major impact of a DoS attack would be limited to dominating the best effort class only.

The policy would be even more effective if you identify DoS-type traffic and mark it as scavenger traffic (DSCP CS1). Such identification and marking may be performed by an ingress policer, either at the campus access edge or at a service provider's ingress edge. Such a policer could be defined so that normal traffic is passed untouched, but traffic that seems abnormal (as defined by volume) is marked to CS1. This suspicious traffic is not automatically dropped by the policer, since it may be legitimate (albeit unusually high volume) traffic. However, it will be constrained to the scavenger class and limited on a link-by-link basis accordingly. In this manner, even best effort traffic is protected against DoS attacks.

An example of a campus policer that marks abnormally high-volume traffic as scavenger traffic is shown below. In this example, traffic in excess of 10 Mbit/s is marked to CS1 (DSCP 8) on a Cisco Catalyst® 2950 Series Switch. (Note: the syntax will vary slightly among Catalyst platforms and versions of Cisco IOS Software; however, the basic concept remains the same.)

Example 2: Catalyst 2950 Access-Switch Ingress Policing to Mark Abnormal Traffic Volume to DSCP CS1 (Scavenger)

```
policy-map POLICE-DOS-TYPE-TRAFFIC
class class-default
  police 10000000 8192 exceed-action dscp 8
```

ZAHEER AZIZ, CCIE #4127, is a Cisco technical marketing engineer for MPLS networks, and co-author of "Troubleshooting IP Routing Protocols," a Cisco Press book. He can be reached at zaziz@cisco.com.

TIM SZIGETI, CCIE #9794, is a member of the Enterprise Solutions Design team at Cisco. He works closely with customers and engineering to develop scalable, tested solutions for next-generation enterprise infrastructures. He can be reached at szigeti@cisco.com.

As with any security policy, it's good to have multiple lines of defense. Therefore, similar policers can also be deployed at aggregation points in the network. What constitutes "normal/abnormal" traffic volume will vary according to context. Trend analysis is highly recommended when setting such policing policies.

Defense Strategy against Spoofing Attacks

Spoofing IP source addresses is a fairly common technique used during DoS attacks. However, hackers may also try to spoof control plane traffic such as routing updates. This article will only address protection techniques against spoofing of control plane traffic.

Spoofing of control plane traffic could allow a hacker to gain access to the network routing topology to further attack the end host. The network could also be disrupted when spoofing successfully resets the neighbor relationships of dynamic protocols, or if incorrect information is injected into the network. Spoofing control plane traffic is a difficult—but not impossible—task, as each dynamic protocol may be using transport layer (TCP) that has its own sequence number and acknowledgement scheme to detect and receive correct packets. Nevertheless, if a hacker has access into the network, and has the ability to sniff TCP packets, then he can forge malicious packets that will resemble valid protocol packets. However, several sequence number predictability theories exist that can predict the TCP sequence number, thus allowing hackers to copy false packets on the wire.

Authenticating peering routers is the most common technique used against such attacks. Cisco IOS offers two authentication techniques to protect dynamic protocols against DoS attacks: plain text password and message digest algorithm version 5 (MD5). In both cases, peers are configured with "keys" or passwords that negotiate with each other to authenticate trusted peers. In plain text authentication, keys are sent on the wire whereas in MD5, a message digest is sent rather than the actual keys. Protocols that support plain text authentication include IS-IS, OSPF, and RIP version 2. Protocols that use MD5 authentication include OSPF, RIP version 2, BGP, and IP Enhanced IGRP.

Example 3: MD5 Configuration between BGP

```
router bgp 109
  neighbor 145.2.2.2 password v61me0qke13
```

The keys at both ends of BGP peers should be identical for MD5 authentication to work.

DoS attacks against the network infrastructure are a growing threat to network security. By proactively safeguarding dynamic protocols with the QoS and authentication methods described above, you'll be a step ahead of the hacker. ▲▲



THE SELF-DEFENDING

FASTER SPREADING, MORE MALEVOLENT

BY
RHONDA
HELDMAN
RAIDER

NOT SO VERY LONG AGO, VIRUSES WERE MERE NUISANCES, BARELY registering on the concerns of senior executives. Today, a combination of self-propagating threats, collaborative applications, and interconnected environments has transformed security—or its absence—into the stuff of headlines. Case in point: in January 2003 the SQL Slammer worm took down a sizable number of networks worldwide, and nearly felled Korea Telecom Freetel's network.


“Hackers targeted individual computers in the '80s, individual networks in the '90s, and the global infrastructure today,” says Jeff Platon, senior director for security product and technology marketing at Cisco. “As a result, security has shifted from annoyance avoidance to a business-critical issue. Senior IT managers have arrived at the conclusion that they need different methods and architectures to combat the new breed of threats.” In fact, end-user expenditures on security products and services will grow 78 percent between 2003 and 2007, from US\$4.5 billion to US\$8 billion, according to Infonetics Research.

Sizing Up the Risk

“In terms of threats, it's not a safer world than it was five or ten years ago,” observes Lance Hayden, business development manager for Cisco's Advanced Services for Network Security (ASNS) practice. The number of security incidents reported to the CERT Coordination Center, a US government-funded research and development center operated by Carnegie Mellon University, rocketed from 9859 in 1999 to 114,855 in the first three calendar quarters of 2003.

Not only is the volume of threats rising, so is the damage potential. Today, most attacks involve denial of service (DoS). If Trojan Horse or worm exploits start harnessing distributed computing power, or establishing peer-to-peer file sharing, risk will soar. “The current worm scenarios are akin to someone firing a million missiles, but with no sophisticated warheads,” says Hayden. “It's dangerous of course, because you have debris falling out of the sky and smashing down on the

infrastructures, but nothing is really blowing up. Imagine if a hacker coded sophisticated, malicious payloads into a worm, taking advantage of distributed computing power or peer-to-peer disk storage. By harnessing the true power of the network, controlling tens of thousands of compromised hosts, the hacker could do far more damage than the worms we've seen.”



NETWORK

NETWORK THREATS DEMAND NEW SECURITY STRATEGIES AND TECHNOLOGIES.

To defend their networks, IT professionals need to be aware of the new nature of security threats, which includes the following:

Shift from internal to external attacks. Before 1999, when key applications ran on minicomputers and mainframes, threats typically were perpetrated by internal users with privileges. Between 1999 and 2002, reports of external events rose 250 percent, according to CERT.

Shorter windows to react. When attacks homed in on individual computers or networks, companies had more time to understand the threat. Now that viruses can propagate worldwide in 10 minutes, that “luxury” is largely gone. Antivirus solutions are still essential but are not enough: by the time the signature has been identified, it is too late. “With self-propagation, companies need network technology that can autonomously take action against threats,” says Platon.

More difficult threat detection. “Attackers are getting smarter,” notes Jeff Wilson, principal analyst for virtual private networks (VPNs) and security at Infonetics Research. “They used to attack the

network, and now they attack the application or embed the attack in the data itself, which makes detection more difficult.” An attack at the network layer, for example, can be detected by looking at the header information. But an attack embedded in a text file or attachment can only be detected by looking at the actual payload of the packet—something a typical firewall doesn’t do. “The burden of threat detection is shifting from the firewall to the access control server and intrusion detection system,” says Hayden. “Rather than single-point solutions, companies need holistic solutions.”

A lowered bar for hackers. Finally, a proliferation of easy-to-use hackers’ tools and scripts has made hacking available to the less technically-literate. “The advent of ‘point-and-click’ hacking means the attacker doesn’t have to know what’s going on under the hood in order to do damage,” says Hayden.

Toward the Self-Defending Network

As security threats have shifted from individual networks to the infrastructure, so too has Cisco’s approach to security. “We’ve evolved from providing point security solutions to integrated systems for secure connectivity,” explains Platon. These integrated systems defend the network from several positions. For example, Cisco offers three defense systems for networks: firewalls, intrusion detection system (IDS), and behavioral anomaly software, which employs sophisticated mathematical algorithms to judge what is “normal” network activity and what is anomalous—and therefore a potential threat. Platon likens this three-pronged approach to the physical security measures used at a bank. The firewall is the guard standing outside the bank. The IDS is a 24-x-7 video monitoring system. Behavioral anomaly software is someone standing outside the vault who is alert for suspicious behavior and prepared to take immediate action. Identity management is analogous to asking customers for a photo ID or to enter a personal identification number (PIN) before receiving bank services. With a variety of ways to detect and prevent attacks, the organization strengthens its defenses.

And yet, though a broad range of effective security technology is readily available, a significant number of



JIB HUNT

TRENDS GLEANED FROM

CISCO SECURITY POSTURE ASSESSMENTS

Industry has been sluggish about taking preventive action against threats it has already experienced, even though the technology is readily available. For instance, it has been two decades since the movie *War Games*, in which a teenage hacker used a backdoor unauthenticated dialup modem to inflict damage. Still, the Cisco Advanced Services for Network Security team regularly sees these same unsecured modems in large enterprise networks. Here are some other key trends Lance Hayden and his Cisco team have noted over the last several months.

- *More compromised hosts.* This is probably due to increased worm activity, says Hayden. Next-generation worms act like little automated hackers.
- *Propagation code tied tightly to the vulnerability.* Two tracks for building code have emerged. In one, the self-propagation code is developed independently of the code used as the vector. In the other, the propagation and vulnerability are enmeshed.
- *New places to look for bugs.* In the mid- to late 1990s, hackers focused on NetBIOS. In early 2000, they shifted to Microsoft Internet Information Services (IIS), and soon thereafter to Microsoft SQL Server. The latest focal point is underlying communications protocols such as distributed component object model (DCOM) and remote procedure call (RPC). Hackers are moving away from searching for bugs in specific applications and toward the underlying software infrastructure.
- *Bad passwords.* The Cisco Advanced Services team can usually crack 70 to 85 percent of most organizations' passwords, often because password policy, if one exists, is not written well or enforced. If security administrators take just one piece of advice, says Hayden, it should be to disable Windows LANMAN authentication, which has become trivial to crack. Next, they should get serious about password security. "Take whatever measures you need to force people to pick secure passwords. Then periodically run a password cracker program to identify passwords that are too easy to crack and ask the owners to change them."
- *Clear text protocols with UNIX systems.* Many companies that use Telnet and Remote Shell (RSH) are still sending passwords and usernames completely in the clear. There's no excuse for not transitioning to safer protocols such as Secure Shell (SSH) and IP Security (IPSec), both of which are supported in Cisco routers.
- *Insecure wireless networks.* Wireless security technology is easy to use. Companies should take advantage of Extensible Authentication Protocol-Cisco Wireless (Cisco LEAP), and backend authentication servers available in the Cisco Aironet® product line.
- *More traffic on port 80.* Now, with so much traffic traveling through port 80, shutting down this port is no longer an option. It is becoming more difficult to monitor exploits because they are wrapped in ubiquitous Web protocols.

companies do not take advantage of it. "No matter how bad attacks become, some companies hesitate to invest in an IDS, conduct vulnerability assessments, hire and train people to monitor the network, and take other action," says Wilson. "What restrains them is the complexity. But if the network becomes self-defending—that is, it blocks attacks effectively and doesn't require highly trained security professionals—the equation changes and they'll make the investment."

To help businesses and government organizations identify, prevent, and adapt to new IT infrastructure threats, Cisco has created the Self-Defending Network Initiative. The first solution, called *Network Admission Control* (NAC), tackles trust and identity

management—which historically has referred to requiring a username and password. While this authenticates the user, it ignores the security credentials of the computer. To add this essential dimension of trust, Cisco recently announced a collaborative effort for NAC with antivirus vendors Network Associates, Symantec, and Trend Micro.

Here's how it works: The *Cisco Trust Agent*, which is integrated into the *Cisco Security Agent* (CSA), collects security state information from PCs and hosts, such as the version of antivirus software and operating system patches. When the node attempts to connect to the VPN, the Cisco Trust Agent transmits the security state information to Cisco network access devices, such as

routers, switches, wireless access points, and security appliances, which enforce admission control. These devices relay the security credentials to the *Cisco Secure Access Control Server* (ACS), which makes the decision to permit, deny, quarantine, or restrict based on customer-defined policy. NAC can even enforce differentiated access policy for hosts with and without the Cisco Trust Agent. What will make NAC effective on a wide scale is the ubiquity of the client agent. "By forming a partnership with leading antivirus vendors, Cisco was able to ensure that the Cisco Trust Agent will reside on nearly 95 percent of all computers," says Platon. "By coming together to solve the critical problem of security that faces the

industry today, we are uniting behind the power of the network to defend IT infrastructure against attack.”

Used with IDS and behavioral anomaly software, Cisco NAC becomes even more powerful. For example, when IDS is embedded in Cisco routers and detects the signature of a DoS attack, within a few seconds the Cisco ACS can command all routers in the network to deny that traffic. And if traffic patterns change in an alarming manner but the signature isn’t known, behavioral anomaly software can instantly shut down the suspicious traffic according to the company’s own business rules. “NAC complements, rather than replaces, classic security technologies that are widely used, including gateway firewall, intrusion protection systems, user authentication, and communications security,” says Platon.

First Steps

How does a company get started with an effective security approach? As a first step, the Cisco Advanced Services team provides organizations with a comprehensive evaluation of their network security called a Security Posture Assessment (SPA). SPAs are conducted by Cisco ASNS experts with leading credentials, including planning, designing, implementing, and optimizing security operations for global Fortune 500 companies and government agencies. Many hold CCIE® and Certified Information Systems Security Professional (CISSP) certifications.

The main goals of a SPA include quantifying the vulnerability state of networked systems to minimize serious or long-term exploitation, determining the ability of current personnel to detect and respond to security incidents, and providing recommendations to help systems and network administrators improve their security. The ASNS team uses sophisticated tools and a methodology to perform extensive, nondestructive probing and exploitation of network vulnerabilities, mimicking the actions and attacks of an intruder attempting to gain unauthorized access.

A variety of network components are assessed, including the network perimeter firewalls, routers, interior firewalls, switches, and virtual LANs. The SPA also looks at hosts, servers, user workstations, and messaging as security policies are tested for operational effectiveness.

Security is evaluated from an internal perspective to review the potential exposures that might allow sensitive information to be breached by unauthorized or unintended actions, and from an external point of view to identify vulnerabilities that might allow internal data to be compromised or resources denied. SPAs are conducted using a suite of software developed by Cisco engineers and tools available on the Internet including full open source tools, modified freely available tools, and custom proprietary tools.

During the internal SPA, Cisco engineers simulate an intruder’s attack in a controlled, safe manner to manually confirm vulnerabilities, including the potential for exploitation of trusted relationships between hosts, password weakness, or administrative access to systems. The consultants approach the network as if they were a disgruntled employee, rogue contractor, or other trusted inside user. During the external SPA, Cisco engineers can simulate hostile activities typical of malicious attackers attempting to compromise perimeter devices and Internet security controls. This might include a “war-dialing” analysis of an organization’s telephone numbers, which typically provide easy backdoors into the network through unknown or unsecured remote-access servers. The Advanced Services team provides similar services for wireless networks as well.

When the assessments are complete, the Cisco ASNS team assembles a report in which, among other technical details, discovered vulnerabilities are prioritized and specific actions are recommended to improve network security. Most often, recommendations focus on technologies and policies, and, sometimes, even a network redesign. In any case, the Cisco Advanced Services group can design and implement defense-in-depth security throughout an enterprise (for more on security best practices across the enterprise, see “A Winning Game Plan,” page 29).

IT Governance and Communication

Security also requires IT governance. That is, companies are finding they need to think about security not only laterally, from the desktop to the network core, but also vertically, from the chief technology officer or board-of-directors level down through the policies implemented in the firewall. “It’s essential to consider both the technological strategies and their business implications,

including the processes that need to be in place,” says Hayden.

Wilson of Infonetics agrees. “Any time a company looks to become more secure, they have to balance the hit to business operations. A security policy that restricts connectivity—for example, by requiring passwords so long and complicated that people can’t remember them—restricts productivity and won’t last long. The need for security is never stronger than the mandate for increased productivity and connectivity. Connectivity always trumps security.”

Arguably the most powerful defense against attacks is open communication and collaboration across businesses and governments. The importance of communication as a weapon against network attacks was underscored in December 2003, when Tom Ridge, US secretary of homeland security, urged technology companies to share information to protect networks from exploitation by terrorists.

Ultimately, Cisco’s approach to security is not about protecting systems *per se*, but protecting the productivity of users—their ability to do their work. “There are no silver bullets in security,” says Hayden. “New technologies such as Cisco NAC and the Cisco Trust Agent are very powerful, but still need to be considered as part of an overall solution set. An effective security strategy requires both technology and attention to governance issues around business process and procedures.”▲▲

FURTHER READING

- Cisco Advanced Services for Network Security: cisco.com/packet/161_6a1
- Cisco Self-Defending Network Initiative: cisco.com/packet/161_6a2
- Cisco wireless LAN security: cisco.com/packet/161_6a3
- Cisco IOS security: cisco.com/packet/161_6a4
- Cisco Secure Access Control Server: cisco.com/go/acs
- CiscoWorks VPN/Security Management Solution: cisco.com/go/vms



A WINNING GAME PLAN

BEST PRACTICES FOR END-TO-END DEFENSE IN DEPTH

BY
GAIL
MEREDITH
OTTESON

SECURING A NETWORK IS LIKE WINNING THE BIG GAME—BOTH sides have talented players and smart coaches, but the victor has the most effective game plan and tightest teamwork. Likewise, a winning security strategy combines multilayer design, robust products, and best practices to create a synergy of expertise and technologies that protects the business. Security operators and the network infrastructure must interact to rapidly identify, prevent, and adapt to security threats.


The dramatic spread of worms such as Slammer and Blaster highlights the need for multiple levels of defense, where many security functions interoperate to mitigate threats throughout the network. This defense-in-depth strategy is the essence of SAFE—the Cisco blueprint for combining security design, technologies, and best practices to protect technology assets and the businesses that depend upon them.

The Best Defense? A Strong Offense

Even the most advanced security products have limited effectiveness if the people operating them are careless or inconsistent. Andrew Peters, manager of VPN and security marketing at Cisco, talks about the interplay of the four P's—policies, processes, people, and products—and how they determine strong or weak security. “Technology alone will not solve all your security problems,” says Peters. “Best practices are truly operations oriented and if they are not complete, it doesn't matter what defense technologies you're using.”

Securing the network is fundamental, and diverse network environments—the campus, the data center, full-service branches, and teleworkers—have distinct security challenges. Best practices start with a proactive security posture, not a reactive one. With strong defenses in place, many events can be detected and mitigated automatically, leaving fewer for security operators to react to, and those reactions can be managed more effectively.

For example, an excellent practice is patching server operating systems for known vulnerabilities; however, it is less stressful to



update tested, validated patches on a scheduled basis rather than hastily upload unproven fixes after a successful attack. This practice requires an interplay of strong design, consistent processes, and effective people. The design must shield servers from attack, perhaps by installing Cisco Security Agent on each server to thwart any known or unknown malicious behavior. (To find out how to get a 30-day trial of the Cisco Security Agent, contact your Cisco account representative.) The consistent process is scheduled patching. Effective people perform thorough patch testing with affected applications wherever possible prior to uploading scheduled fixes, leading to more stable, secure application environments.

Security Policies

The security policy defines who can do what, where, and when. “It should tie the business strategies of the Board of Directors down to the folks in the trenches configuring boxes,” says Lance Hayden, business development manager in the Advanced Services for Network Security practice at Cisco. “You have to consider top-to-bottom usefulness, so it’s really an issue of governance.”

A security posture assessment can identify vulnerabilities and suggest fixes. In performing security posture assessments for Cisco customers, Hayden reports that he rarely finds a customer with a completely viable security policy in place. As a result, organizations are often ill prepared to handle next-generation security threats. Assembling a strong security policy is a huge effort, he admits, but worth it. “If you are IT dependent, your IT should align with your business, and that includes security. If security and business don’t map, then you’ve got a problem.” The organization also needs to learn from its mistakes and fix what’s broken, and an agreed-upon security policy makes it easier to do that.

Hayden attributes the lack of effective security to poor communication both within and between organizations. He warns that industries don’t communicate nearly as well as the hacker community does, putting the “good guys” at a perpetual disadvantage. He identifies three pitfalls in security policy development and suggests ways to overcome them.

The first pitfall is a policy that does not map to the business environment. As a

result, users find ways to get around policies, or policies are not enforced. This adversely impacts an organization’s ability to do business by creating vulnerabilities and requires “knee-jerk” reactions to security issues. Security operators can avoid this problem by learning about their business or industry, then writing policies within the context of the business environment. Regulated industries such as financial services or healthcare require strict compliance to rules, and the policy must reflect them. Sharing information among hospitals or banks may help everyone within those industries to attain tighter security.

The second pitfall is a policy that does not conform to the operational environment, so it is poorly enforced. Security staff should test the network for policy effectiveness. One example is a lax password policy. The Cisco Security Posture Assessment team can usually grab and crack up to 90 percent of user passwords within minutes. Other instances are when firewall rules don’t comply with defined access privileges, or companies fail to secure wireless LANs.

“Wireless technology is not inherently insecure,” says Hayden. “It’s people that make it insecure. You have to go out and test the security of your wireless LAN, looking for vulnerabilities or rogue access points. Cisco solutions such as Cisco WLSE [Wireless LAN Solution Engine], which can not only detect rogue access points, but can also help you pinpoint their locations in the building, can help provide this proactive security testing and monitoring.”

The third pitfall is when policies clash with corporate culture. People with the best insights are not asked for input, and policies are written in a vacuum. For example, a security administrator does not consult the customer service manager before writing a policy about securing the customer relationship management (CRM) application. The manager is the expert on how her department uses CRM. She believes the policy hinders her department’s ability to do business, loses respect for the administrator, and finds ways to work around the policy.

Fostering a strong sense of commitment to security throughout the organization overcomes this problem. Business groups must contribute to policy development, because security affects how they work. IT administrators must actively market the

value of security to everyone in the organization and encourage participation from all business groups. Says Hayden, “Without training and awareness and the buy-in of everyone from the top to the bottom of the organization, no security policy can be effective. People need to understand that the security policy is not an annoyance but a critical component of both corporate and personal success.”

Security Management

Best practices for security management must incorporate tools that can digest the massive amount of data generated by multiple security devices on the network, and provide uniform configurations and changes to devices in a timely manner.

“Our customers see that managing security through the command-line interface [CLI] leads to vulnerabilities,” says Bob Dimicco, director of marketing in the Secure Managed Networks group at Cisco. “Day-zero mitigation using CLI is difficult where there are many security devices. People need management tools as a way to enforce policies everywhere.”

CiscoWorks VPN/Security Management System (VMS) defines role-based management among multiple teams responsible for network security. For example, the security operations team might designate a policy, the central network management team implements and enforces it, and administrators at remote locations can localize it, all through the same application.

CiscoWorks VMS enables the following best management practices:

- Hierarchical rules definition—translates high-level rules into configurations that are pushed to all relevant devices such as firewalls and intrusion detection systems (IDS); this provides consistent policy enforcement
- Change management—using the CLI are frequently unlogged and lost, so tracking actual device configurations can be tricky. CiscoWorks VMS tracks who did what to which devices and when, leaving a clear record.
- Localization management—some policies are mandatory and global, others are suggested and can be localized. CiscoWorks VMS defines roles that limit what local administrators can do and tracks their actions.

- Regulatory audits—complete reporting simplifies regulatory compliance.

Cross-business unit collaboration to develop and maintain accurate user permissions for the Cisco Secure Access Control Server (ACS) database paves the way for greater acceptance of security policies in the user community. Each business unit can help define how its users can access its applications and network services.

Security managers are often overwhelmed with the massive amounts of data generated by multiple devices distributed throughout the enterprise, and important events are missed. CiscoWorks Security Information Management System (SIMS) is a scalable way to normalize, aggregate, correlate, and visualize security data in a manner useful for threat mitigation and rapid response to attacks.

A comprehensive security solution requires management of network devices as well as security devices—and Cisco solutions span both. Innovative Cisco IOS® Software features such as AutoSecure help ensure network availability, and integrate with current network management and security tools, including features such as Embedded Syslog Manager, Simple

Network Management Protocol version 3 (SNMPv3), and Secure Shell (SSH).

Embedded Syslog Manager provides a customizable framework for correlating, augmenting, filtering, and routing syslog messages generated by the IOS logger. SNMPv3 provides secure access to devices by authenticating and encrypting packets over the network. SNMPv3 uses a Data Encryption Standard (DES) encryption of the SNMP packets, as well as R/W passwords and host-based access to prohibit misuse.

With SSH, your network supports secure, remote connections to routers through encryption across a Telnet session. SSH encrypts all traffic, including passwords, between a remote console and a network router across a Telnet session. Because SSH sends no traffic in the clear, network administrators can conduct remote access sessions that casual observers cannot view. The SSH server in Cisco IOS Software works with publicly and commercially available SSH clients. This can be an essential solution for integrated branch office security solutions.

Campus Networks

The campus network is not homogeneous—it is actually divided into several security

“modules” in the SAFE blueprints from Cisco, some facing outward to public networks and customers, the rest serving internal corporate users.

Most organizations are alert for external threats such as worms, viruses, and Distributed Denial of Service (DDoS) attacks, but internal hackers can do far more damage, in part because they are harder to detect and people aren’t looking for them. This occurs when “the network is crunchy on the outside and soft in the middle,” says Russell Rice, product marketing manager in the New Security System Technologies business unit at Cisco. “It’s important to compartmentalize the network to minimize damage. You need to build the network to protect key assets and services.”

An emerging campus security threat is grabbing, listening to, and spoofing voice over IP (VoIP) traffic (see “Securing Voice in an IP Environment,” page 39).

Cisco has advocated a defense-in-depth security strategy since introducing SAFE in 2001. These security solution “layers” are embedded in Cisco IOS Software, dedicated appliances, and integrated modules for Cisco routers and switches. The

CISCO NETWORK ADMISSION CONTROL

Cisco dramatically enhances the day-zero antivirus and antiworm functionality of Cisco Security Agent with Network Admission Control (NAC), the first step in the multiphase Cisco Self-Defending Network initiative. Available in mid-2004, Cisco NAC allows organizations to discover operating system patch, antivirus, and hot fix status of devices requesting network access. It can then relegate noncompliant and potentially vulnerable systems to environments with limited or no network access. Noncompliant endpoints can be denied access, placed in a quarantined area, or given restricted access to computing resources—perhaps to allow upgrades and patches to attain policy conformance.

Cisco NAC has four components:

- Cisco Trust Agent—software that resides on an endpoint system, collects security state information about the system, and communicates it to Cisco network access

devices. Cisco has licensed this technology to its antivirus co-sponsors—Network Associates, Symantec, and Trend Micro—for integration with their antivirus clients. Cisco Trust Agent will be integrated with Cisco Security Agent.

- Network access devices—enforce admission control policy by demanding host security “credentials” from endpoints requesting access and relaying information to the policy servers; access devices then permit, deny, quarantine, or restrict access according to policy.
- Policy server—evaluates endpoint security information and determines the appropriate access policy to apply.
- Management system—Cisco NAC management will be incorporated into existing security management tools.

For more information about Cisco NAC, visit cisco.com/packet/161_6b2 (Cisco.com login ID required).

Cisco security solution comprises the following three categories:

- Threat defense—watches for improper behavior in the network; examples include firewalls and network- and host-based intrusion detection/prevention systems (IDS/IPS)
- Trust and identity management—permits

or denies services to devices and users based on policies; an example is a Remote Access Dial-In User Services (RADIUS) access control server

- Secure connectivity—provides confidentiality across public links such as the Internet; for example, a virtual private network (VPN) with encryption

Detailed best practices for deployment, configuration, and management of campus networks are described in the SAFE white papers. The more important practices (many of which also apply to the data center and full-service branch office) are the following:

- Deploy behavior-based host IPS software in addition to signature-based solutions for

SECURE BANKING OVER A FLEXIBLE WAN

Chevy Chase Bank wanted secure data communications between its 220 branches in the Washington, D.C. area and its central data center. Security is always a concern for banks, and the push by regulatory agencies for continual improvements heightens that concern. Chevy Chase Bank, headquartered in Bethesda, Maryland, also wanted more communications bandwidth. It needed to revamp its local and wide-area networks to accommodate rapid growth—23 new branches in 2003 alone—and change how it handles and uses data. For example, the bank is preparing to deploy significantly more browser-based applications throughout the branch network.

Chevy Chase Bank chose to implement an IP VPN to secure communications between its branches and data center. The secure connectivity solution, which takes advantage of the existing infrastructure, will initially be used to secure the transport of information across its current Frame Relay network while positioning the bank to migrate to other transport media in the future.

The new, secure WAN will consist of integrated VPN acceleration modules in Cisco 2691 routers in the branch and 7200 routers in the data center. Encrypting in hardware saves the router processing power, helping the bank to meet its cost goals. Installation of the 2691 routers began in four pilot branches in late 2003; the remainder of the secure WAN will be installed by June 2004.

“We were looking for one single integrated platform that could handle broadband communications over dedicated links, and increase the security of those links by encrypting the data they transport,” says Bob Spicer, executive vice president and chief information officer at Chevy Chase Bank. “The Cisco routers and VPN cards provide an integrated platform that will be centrally managed and will work seamlessly with our legacy LAN.”

Installation, he adds, “is as easy as Cisco claims it is. While careful planning and design is critical to any network plan, the physical deployment has been greatly simplified by built-in Cisco features.”

The new, secure WAN, which provides IPSec protection for the VPN links with 3DES data encryption, will ensure that all data going out from the dozen or so devices in each branch will be protected. In the future, it will support further changes the bank plans. By the end of 2005, the bank intends to upgrade all branch LANs to Ethernet. That will now be easier since “you don’t have to touch the router,” Spicer says. “You just plug the LAN into a different port. That’s a real advantage.” He’s also planning to upgrade all the PCs in the offices to handle more browser-based tasks. As the bank employs more Internet-based applications, the secure WAN router network can prioritize different types of traffic with the integrated QoS available today.

The Cisco 2691 and 7200 routers, which make up the secure WAN solution, will also be able to handle the increased bandwidth needed for the new applications. Currently, most branches use 56 kbit/s links—fast enough for subsecond financial transactions. As the Internet-based applications are implemented, those speeds could increase to 768 kbit/s, and upwards to 1.5 Mbit/s as needed.

In the future, if the bank decides to employ the recently developed Advanced Encryption Standard (AES), the integrated VPN acceleration modules can perform AES as well as 3DES, leveraging today’s secure connectivity investment for tomorrow’s needs.

“The nice thing about these routers,” Spicer says, “is that they can handle both our legacy and our new networks. Installation goes very quickly. We can get all the networking in place and then follow up with our PC and LAN and capacity upgrades.”

day-zero worm attack mitigation on key hosts—systems protected by Cisco Security Agent successfully thwarted Slammer and Blaster worms (see sidebar, “Cisco Network Admission Control,” page 31).

- Disallow outbound session initiation from Web servers to prevent virus and worm propagation
- Define unique hostnames and passwords for routers and switches; change passwords periodically
- Differentiate groups using filters and virtual LANs (VLANs) for containment
- Use an out-of-band management network and encrypt management traffic; SNMP version 3 supports encryption
- Limit data flows with Cisco IOS quality of service (QoS) features such as Committed Access Rate (CAR) to minimize the impact of DDoS attacks (see “Mitigating Network DoS Attacks,” page 21)
- Configure Network-Based Application Recognition (NBAR) to filter worms and unpermitted HTML traffic
- Use Layer 2 security features such as Dynamic ARP Inspection (DIA), and Dynamic Host Control Protocol (DHCP) snooping to prevent spoofing. Shut down unnecessary firewall and router services, which hackers can exploit to gain entry; for example, AutoSecure, a “one-click” feature in the Cisco IOS Software, disables router services that hackers commonly exploit, such as Finger, BOOTP, and Proxy-ARP; enables protection through enabling services such as password encryption and TCP keepalives; and secures the forwarding plane
- Keep server operating system software up to date with the latest patches. The Slammer worm exploited a known vulnerability patched by Microsoft six months prior to the attack
- Use 802.1X port-based authentication
- Close unused ports on all devices
- Expire user passwords after a specified period; enforce strong password rules
- Go beyond WEP when securing the wireless LAN, and regularly look for rogue access points
- Conduct periodic security audits of the network, preferably by a reputable third party

The Data Center

The heart of the IT infrastructure is the data center, which houses the applications that keep organizations in business. Best practices here start with an integrated security-in-depth design. The data center has the highest density of business-critical resources, including applications, servers, storage, and networks that need to be protected, so it's important to consolidate security services into a dedicated services tier placed between the aggregation and core network tiers. Larger data centers can fuse security services with the network infrastructure using integrated security modules for the Cisco Catalyst® 6500 and 7600 series platforms. “A dedicated security services tier with integrated modules is easier to manage and delivers the higher performance levels required to simultaneously protect multiple application, server, and network environments within a data center,” says Greg Mayfield, senior security product marketing manager at Cisco.

The most critical data center security best practices include the following:

- Compartmentalize the data center into security “zones” and define policies for each one, including access rules and rules for how zones interact. Use of VLANs and virtual SANs (VSANs) enables zones to be created for the different tiers (for example, Web, application, and database) within a multi-tier applications and between different applications. Placing integrated firewall modules between zones efficiently enhances threat defense for the different application and server environments against directed and indiscriminate security attacks.
- The data center design should include both network- and host-based IPS that watches every server, switch, and router for suspicious activity, then configure IPS to automatically reconfigure firewalls to block packets from identified malicious sources
- Cisco Security Agent on each server can mitigate virus and worm attacks, as well as ease the patching and maintenance burden
- Where users require remote access to data center applications, use IP Security (IPSec) VPNs or, alternatively, Secure Sockets Layer (SSL) VPNs for partners. Also use VPNs for interconnecting remote data centers for backup and replication. Managers must also protect the storage-area network (SAN), which is often located off-

site, by isolating it from outside user access and using the security features of Cisco MDS 9000 Series switches. For example, storage managers should lock down ports with Port Security, secure Fibre Channel and Internet Small Systems Computer Interface (iSCSI) traffic, and establish VSANs for absolute partitioning between virtual fabrics. Integrating security into the IP and storage network foundation of the data center substantially enhances service availability and reduces costs associated with recovering servers and applications, while protecting confidential business-critical information.

Full-Service Branch Offices

Full-service branch offices are a logical extension of the campus LAN, but physical distances limit the ability of headquarters to protect the branch. As with the data center, best practices start with good security design. A key design consideration at the branch is expense—that's why Cisco recommends using a single IOS-based router with integrated services at the WAN edge to

Continued on page 35

FURTHER READING

- **SAFE Blueprints from Cisco:**
cisco.com/packet/161_6b4
- **Advanced Services for Network Security:**
cisco.com/packet/161_6b5
- **CiscoWorks VPN Management:**
cisco.com/packet/161_6b6
- **Cisco Secure ACS:**
cisco.com/packet/161_6b7
- **CiscoWorks SIMS version 3.1:**
cisco.com/packet/161_6b8
- **Improving Security on Cisco Routers:**
cisco.com/packet/161_6b9
- **Cisco AutoSecure:**
cisco.com/packet/161_6b10
- **Cisco Aironet Wireless LAN Security Overview:**
cisco.com/packet/161_6b11
- **Cisco WebVPN:**
cisco.com/packet/161_6b12
- **Integrated Security Flash Demo:**
cisco.com/packet/161_6b13

Defense in Depth, Continued from page 33

protect voice, video, and data traffic. These services (many of which were traditionally confined to the campus) should include VPN, firewall, IDS, content filtering, and advanced authentication including AAA and 802.1X support. Cisco IOS Software supports all these services—offering an integrated branch office solution that is easy to configure and manage. For increased performance, an alternative is to use hardware modules for VPN acceleration and IDS services. Another option is a security appliance installed just behind the router to provide firewall and IDS services.

Other best practices for full-service branch offices can include:

- Configure VPNs to prevent “back-door” hacker access to the campus network and data center
- Where VoIP traffic traverses the WAN, configure a voice- and video-enabled VPN (V³PN)
- Configure access control lists (ACLs) to conform to security policies
- Install Cisco Security Agent on critical local hosts and servers

For more information, see “Security for the Branch or Small Office” in the Fourth Quarter 2003 issue of *Packet*® at cisco.com/packet/161_6b1.

Teleworkers

Most organizations that depend upon IT for their business provide remote access services for home-based and mobile employees, or teleworkers, to increase productivity. However, recent advances now allow for smooth transition from the corporate office to the home office, greatly improving a business’s ability to continue operations in the face of disruptions such as inclement weather (shutdowns due to Hurricane Isabel cost the US government \$60 million in lost productivity). The Cisco Business Ready Teleworker solution provides this improved business continuance posture by giving employees “the same access privileges and applications at home that they can get at the office,” according to Pete Davis, product line manager for remote access VPNs in the VPN and Security group at Cisco.

An issue with traditional teleworking schemes is that teleworkers add significant uncertainty to an enterprise’s security profile. They might connect directly to other networks—such as the Internet—pick up a

virus or worm, and unknowingly transmit it to the enterprise network the next time they log in. Hackers can disguise themselves as remote or mobile employees and gain entry. And there can be hundreds, even thousands, of end-user systems to manage.

Cisco remote access VPN solutions allow companies to extend internal resources securely over the Internet to day-extenders, teleworkers, and business partners. By using Cisco VPN 3000 Series concentrators, or many other devices at Cisco that also support remote access VPN termination, a corporation can reduce expensive long distance charges and enable higher-speed broadband access to internal resources by teleworkers. Cisco remote access VPN solutions are available through installed software clients, clientless technology (Cisco VPN 3000 Series only), and hardware client solutions.

Cisco Business Ready Teleworker provides an always-on, secure and centrally managed teleworker solution by placing a Cisco 800 Series Router behind the teleworker’s broadband (cable or DSL) modem. This secure solution provides the same user applications (including IP telephony) at home as those available in the corporate office, improving business continuance and productivity without compromising network security.

With Cisco Business Ready Teleworker, the headquarters network can verify the configuration of end-user systems connected to the corporate network. Using Cisco NAC (see sidebar, “Cisco Network Admission Control,” page 33). Network administrators can push automatic security updates to users at any time. Cisco NAC can also check for required active software before permitting user connections from hotspots or dialup.

To protect the headquarters network and its resources, Cisco recommends the following best practices for teleworker security:

- Require a corporate software image on identical laptop PCs—including standard operating system configuration, applications, and utilities. Some corporations may allow departments to overlay specialized software atop the standard image.
- Use a uniform VPN access method for home workers—such as a Cisco 831 IOS-based router for critical staff and “power teleworkers,” and appliances such as a Cisco PIX® 501/506 Firewall or Cisco VPN 3002 Hardware Client for occa-

sional teleworkers. Uniform systems are far easier to manage than multiple connection options.

- Enforce strict password rules with frequent password rotation—while the most secure options are one-time password or digital certificate systems, not all enterprises can afford them. Password enforcement and a strong AAA system at headquarters can mitigate the risks associated with reusable passwords.
- At headquarters, deploy remote access gear behind the WAN edge router and firewall. In larger deployments, an IDS appliance can monitor traffic.
- Protect the system from new attacks, viruses, and worms with security and antivirus software—including Cisco Security Agent—on teleworker PCs.

Cisco eases the management cost of maintaining teleworker networks with Cisco IP Solutions Center. Using this solution, IT staff can see, troubleshoot, and manage all router-based teleworker endpoints from a central location.

Goal: The Self-Defending Network

Cisco is building intelligent security systems that ease operational headaches for its customers. “The network needs system-level defenses that can detect suspicious behavior and respond immediately,” says Steve Collen, director of security in the VPN and Security Services group at Cisco. “The goal of the Cisco security strategy is to enable a self-defending network that can identify an attack, close it down, and reconfigure the network to prevent recurrence.”

However, the self-defending network will always need operators who consistently apply best practices for effective policy enforcement. ▲▲

GET MORE SECURITY CONTENT. NOW.

For white papers, best practices, ROI calculators, and more information on integrated security, visit cisco.com/powernow/packet.



SECURITY

INCIDENT RESPONSE TEAM REACTS QUICKLY TO PRODUCT SECURITY ISSUES AND WORKS HARD TO KEEP CUSTOMERS' TRUST.

AS THE INTERNET EVOLVES TO BECOME AN INCREASINGLY CRITICAL part of business-as-usual for enterprises, so evolves the sophistication of network malefactors who make it their business to seek out and exploit untended vulnerabilities. Fortunately, when trouble strikes, Cisco customers can call on a trusted and expert resource to quickly mitigate the problem and prevent recurrence.

Formed in 1995, the Cisco Product Security Incident Response Team (PSIRT) provides a global, 24x7 body of experts for handling customer security incidents that involve Cisco products, as well as the handling of product vulnerabilities (cisco.com/go/psirt). The team rapidly gathers required resources from throughout Cisco to analyze the problem, develop a solution, and communicate the problem resolution to the appropriate parties. PSIRT leverages resources from Cisco's global Technical Assistance Center (TAC) and various business units to provide swift, efficient, and accurate response for customers. As Cisco's most customer-facing security team, PSIRT takes charge of generating Security Advisories, driving the resolution of security-related bugs, interfacing with external response teams (through such organizations as the Forum for Incident Response and Security Teams, or FIRST), and assisting customers in responding to real-time network security threats such as denial-of-service (DoS) attacks. The majority of the calls and e-mail PSIRT members receive in response to real-time incidents deals with some form of DoS attack.

The team provides incident response assistance to customers with basic service contracts for any incident in which a Cisco product plays a significant role, regardless of whether the problem involves a Cisco product. "When Cisco customers call us, we help them understand the issues, get enough information to mitigate the problem, and either find software to fix the problem or a workaround if a fix is not immediately available," explains Richard Aceves, manager of the PSIRT.

PSIRT members are highly technical engineers, all of whom share a strong interest in security. Most have worked in the TAC prior to joining PSIRT. As one team member puts it, "We are an overly caffeinated group of type-A personalities." Given the nature of their work, which requires them to be electronically reachable nearly every moment of their lives, the high-energy personality trait serves the team well.

The team acts in two modes—primarily they are reactive, responding to whatever security issues arise. But when time allows, they also operate in a proactive capacity, participating in activities

geared toward education and prevention. Operating in reactive mode, PSIRT represents a collection and reporting point for any potential security vulnerabilities in Cisco products. "The first thing we always need to do in that role," says Aceves, "is decide whether we're truly seeing a vulnerability." The team uses the "CIA" qualification

method: if the issue poses a threat to *confidentiality*, *integrity*, or *availability*, it constitutes a security defect.

"One of the best compliments our customers have paid us is asking us to help with security issues not only on Cisco products, but on end-user workstations," Aceves says. Internet worms have been a major concern for Cisco customers, particularly this past year. And although worms don't typically target network equipment, they do propagate on the Net. "We've had customers call us to help get their networks healthy again, but once we've done that, then they ask for help finding all the infected devices on their networks—not Cisco devices, but end-user workstations. We help them to the best of our abilities, but first and foremost we want to help people with their Cisco issues." In the future, Aceves projects that Cisco will extend its purview to help customers mitigate security issues right out to the end users.

The PSIRT Process

Vulnerabilities are reported to the PSIRT through external channels, either Cisco customers or third-party researchers, or through internal channels—for example, other groups at Cisco. Once it's clear that the team is working with a bona fide vulnerability or security defect, the next step is to characterize the problem and assess how to deal with it going forward. This is potentially one of the most time-consuming parts of the process. "Customers look to us to be the experts," says Aceves, "so I hold my team to extraordinarily high standards and make sure we've investigated every detail before moving on." This commitment to thoroughness helps PSIRT provide Cisco customers with as much information as possible in terms of problem scope, fixes, mitigation techniques, and so on.

In the case of vulnerabilities being reported through external sources, a major challenge facing PSIRT is how to communicate the problem to Cisco customers without jeopardizing their security in any way. Timing and level of detail are the two major factors here; they must be finely balanced to make sure customers know what they need to know, when they need to know it—without opening up the vulnerability to exploitation. Announce a



ADVOCATES

BY JOANNA HOLMES

vulnerability too early, and all the pieces may not yet be in place, leaving loopholes open to predators. Announce it too late, and someone might exploit it in the interim. And if you provide too much detail on the nature of the problem, says Aceves, you may well be giving the bad guys the recipe to exploit it.

If the vulnerability extends to multiple vendors' products, a third-party organization such as the CERT Coordination Center (CERT/CC) might be asked to coordinate between vendors so that everyone announces the problem and the fix simultaneously, with the goal being that no networks are compromised by someone leaking the information prematurely.

When it comes to informing the network community of security vulnerabilities, there are generally two schools of thought: *responsible disclosure* and *full disclosure*. Cisco stands firmly in support of responsible disclosure on the premise that full disclosure, which calls for immediate and general announcement of vulnerabilities, creates more risks to Cisco customers than it eliminates. "While the 'responsible disclosure' issue is still being worked on," notes Aceves, "Cisco is actively involved to help lead the discussion and educate both our customers and vendor peers."

Aceves stresses that at all times his team's highest priority is to protect Cisco customers' networks. "Even when there's an issue that is already being worked through our system, we continue to monitor what's going on out there, and we do that with extreme care," he says. "If we see anything going awry, we will immediately change our plans and accelerate, at a minimum, publication of the problem to customers, or we'll push hard on the development organization to accelerate release of a fix." But this, too, is a balancing act. If there is no evidence of exploitation, the team's credo is that Cisco customers are better

PSIRT Manager
Richard Aceves

CORNERSTONES OF A SUCCESSFUL INCIDENT RESPONSE TEAM

Many of Cisco's enterprise customers have or plan to form a security incident response team to support their network-based business functions. For any such team, Aceves identifies three fundamental elements for success.

Make allies of legal and public relations staffs. An incident response team needs to forge strong relationships with its company's legal and public relations departments. These alliances can prove invaluable in the event of a security incident that affects the company's customers.

Get buy-in from the top. Make sure your team has buy-in at the executive level, whether it's the president of a university or the senior staff of a corporation. Keep these

vital supporters well informed during any significant incident—but equally important, try to educate them in advance as to the importance of your team's function within the company.

Assemble a stellar team. Whether the team you put together is two people or two dozen, it should comprise the most talented individuals you can assemble. A security incident response team plays a vital role in a company, and its members must be technically sophisticated and articulate enough to present to audiences of executives, as well as legal and public relations staff—while potentially operating on very little sleep.

served by integrating the fix into the next release and not creating unnecessary fire drills for them.

"We want to keep our customers' trust," says Aceves. "We believe our customers have historically trusted the PSIRT team—and Cisco in general—because we've been very open about any defects in our products. I want to make sure our customers keep trusting us to disseminate the information they need, when they need it, but also to avoid alarming them unnecessarily."

Proactive Measures

In proactive mode, PSIRT members maintain involvements in activities geared toward prevention and education. The team stresses the importance of network security as a deployment issue, and works continuously to educate customers on best practices for secure deployments.

PSIRT members have participated in Cisco's multipronged program that focuses on delivering innovative, integrated security features in Cisco IOS® Software and educating Cisco customers on secure deployment practices, such as the Cisco SAFE architecture. One example of how PSIRT members are contributing to the IOS project is *Cisco AutoSecure*, a new, "one-touch" device lockdown process incorporated into Cisco IOS Software that enables network operators to quickly secure a device under attack. AutoSecure applies best-practice

recommendations from Cisco experts and the US National Security Agency. For more on IOS enhancements that deliver innovative, integrated security throughout the network, see "Locking Down IOS," page 17.

In another example of his team's contributions, Aceves recalls how a PSIRT member took an interest in random sequence numbering for packet-level communications. At the time, the algorithm for producing random sequence numbers constituted a vulnerability that, fortunately, had not been exploited. The PSIRT member's efforts led to changes in the Cisco IOS Software. "Some time after that," Aceves notes, "an independent organization ran tests on sequencing of numbers for various network vendors' products, and Cisco came out extraordinarily well."

"We all need to acknowledge that *nobody* writes bug-free code," Aceves says. "With that in mind, vendors—Cisco and others—must make sure that we find any problems that have been introduced, and educate our development community on how to avoid those pitfalls in the future."

PSIRT members also leverage their day-to-day customer interaction by representing customer concerns to Cisco's product development organizations. "We constantly hear

about challenges our customers are facing with respect to securing their networks," says Aceves. "We pass that information on to our development organization, advocating solutions for our customers to solve those problems."

In addition, the team stays actively involved with a variety of security-related organizations, such as FIRST and the Infraguard program sponsored by the US Federal Bureau of Investigation (FBI). PSIRT members believe that as advocates for the security of Cisco customers' networks, they must be security experts in every sense—not strictly for Cisco products, but for the end-to-end scenario. "Our customers have asked for this, and we're responding," says Aceves. "We're continuously increasing our TAC services and publications, looking at it from the whole perspective of how to put together an end-to-end solution that's as secure as possible."

It's exciting work that Aceves acknowledges can be stressful at times. "But there's also a great deal of satisfaction in contributing to the safety of the Internet, and making sure that our products are deployed in a way that delivers their full functionality with a high enough level of security that our customers don't need to worry about it." ▲▲





SECURING VOICE

IN AN IP ENVIRONMENT

DEFENSE-IN-DEPTH STRATEGIES FOR MAKING VOICE
AS SECURE AS ANY OTHER MISSION-CRITICAL APPLICATION



iNTegrating voice services and systems with data IP networks has clearly won the hearts, minds, and wallets of enterprises. But as voice technologies meet the world of IP, many IT, telephony, and business managers are concerned about its safety from attack. Voice requires high security standards, equal to those of high-security data applications.

The information in voice calls—strategic, personal, or financial—can be just as proprietary or damaging if intercepted as that in data. Moreover, no enterprise or service provider can afford a denial-of-service (DoS) attack that shuts down voice communications. Not only is voice still the lifeblood of most businesses, users absolutely must be able to get through to emergency services.

BY JANET KREILING

SPENCER TOY

Many enterprises have implemented programs to ensure data security—which, of course, also protect voice traveling as data—based on Cisco SAFE guidelines. To focus specific attention on voice security, Cisco has further developed comprehensive SAFE guidelines for IP voice that build on those for data networks and focus on the same three areas: secure connectivity, managing trust and identity, and threat defense.

Secure connectivity encompasses not only the underlying data infrastructure that carries voice—there are specific provisions that augment safety for voice. The SAFE blueprints for IP networks carrying voice pay special attention to protecting the four main voice systems: the IP phone, the call-processing manager, the voice-mail system, and the voice gateway. Similarly, there are specific provisions within trust and identity management and threat defense for voice. For example, there are measures to segment and separate voice calls from data, to secure IP phones, to prevent infiltration of voice platforms by viruses that have penetrated the data network, and

to broker connections between the voice and data segments of the network. (In this article, the term “IP voice” includes both voice over IP or VoIP calls that may begin in analog form and are converted to digital for transit, and IP telephony—voice calls that travel in IP form from start to finish.)

Points out Roger Farnsworth, senior marketing manager and long-time security specialist at Cisco, “A broad and integrated approach is essential to securing both data and voice. While no security program is 100 percent foolproof, careful adherence to the SAFE guidelines will make your Cisco IP voice communications as safe as they can be. Cisco’s leadership and expertise in data security solutions make its solutions unique in their ability to protect IP telephony and voice over IP.”

Secure Connectivity

The first step should be developing a security policy for voice, paralleling the one that you should already have developed for data.


Only after this policy has been created should you give your attention to protecting IP voice. Jason Halpern, manager of technical marketing for security at Cisco, notes that many existing data security steps also protect this emerging technology, and should be implemented before addressing voice security specifically. First, says Halpern, “You need secure connectivity, which requires a secure underlying data network—that is, one hardened wherever possible at Layer 2 and Layer 3.” Among hardening measures, he says, are “increasing security on routers and switches by taking advantage of security features built into Cisco IOS® Software such as stateful firewalls and intrusion detection, enabling features that promote security, disabling those that can expose the network, and hardening device configurations.”

Hardening a router includes actions such as locking down Simple Network Management Protocol (SNMP) access, turning off unneeded services, using secure management methods such as Secure Shell (SSH), and authenticating routing updates. Hardening a switch includes Layer 2 hardening through features such as Address Resolution Protocol (ARP) inspection or private virtual LANs (VLANs), using IP permit lists to restrict access to management ports such as SNMP, using a dedicated VLAN ID for all trunk ports, and disabling unused ports.


Similar hardening should be performed on call-processing managers, voice mail systems, and voice gateways. “It’s crucial that these appli-

THE THREATS

Many threats to IP voice are, like DoS, familiar from data communications: viruses, worms, Trojan horses, man-in-the-middle attacks, packet sniffing, IP spoofing, password attacks, trust exploitation, and the like. Then there are threats such as toll fraud explicitly directed at voice. Actually, attacks against voice calls and voice-related systems are not new. (In fact, hacking got its start in the exploitation of telephony systems as “phone phreaks” sent dual-tone multifrequency [DTMF] signals over telephone lines to simulate call control signaling and steal phone calls.) Traditional voice communications can be vulnerable to DoS attacks, fraud, eavesdropping, and other security breaches; these can also threaten IP voice. In an IP network, a hacker might attempt to get into a voice gateway for “free” calls. To violate privacy, he or she might employ a bit of software nastily acronymed VOMIT (voice over misconfigured Internet telephones), designed as a proof-of-concept VoIP eavesdropping tool, to reassemble an IP voice trace captured by the UNIX tool tcpdump into a listenable wave file. Many commercially supported diagnostic tools do the same thing more legitimately for troubleshooting voice quality and other service parameters. Application-layer attacks can specifically target call management, voice mail, and unified messaging systems.



SAFE BLUEPRINT FOR MEDIUM-SIZED ENTERPRISE NETWORK



cations are configured properly according to the vendor's best practices," Halpern adds—a step missed by administrators who assume the systems are shipped with all security aspects enabled. Given the diverse ways they are deployed, these systems must be tailored to each installation.

Hardening the network helps secure it against all the ills data can fall prey to—DoS, spoofing, packet sniffing, viruses, worms, and the like. The next step, segmenting network functions, helps make sure that even if an invader gets into the data network, it won't affect voice traffic, by providing more effective access control and successful attack mitigation. As a first step, the network should be segmented into function modules. A campus network, for example, might be segmented into a management module, building module (users), server module, lab module, building distribution, core, and edge distribution.

Once that's done, Halpern says, segment voice from data and segment voice users into coherent groups by grouping them into VLANs set up on your Ethernet LAN switches. Connections between the

voice and data segments of the network (which is still converged, as these are all logical, not physical, separations) should be limited to specific points, such as the call processing and voice mail systems.

"Segmentation is especially valuable in today's age where worms are becoming more prevalent," Halpern points out. "Sequestering data from voice renders it less likely that attacks in the data segment will affect the voice segment. For example, voice streaming occurs over the Real-Time Transfer Protocol [RTP], which uses a very large range of ports—anywhere from 16,384 to 32,768. Without proper filtering and segmentation, an attack could traverse from data into voice through one of them."

You will also want to turn off system features that can automatically allow would-be users onto the network or restricted segments. For example, many call-processing managers provide an automatic phone registration feature that bootstraps an unknown phone with a temporary configuration and then allows it onto the network. The danger is obvious: It's possible that that phone, or a device appearing to

FIGURE 1: SAFE divides the network into modules for security and manageability. It has been designed to support IP phones, PC-based IP phones, voice services, proxy services, PSTN for WAN backup and local calls, and VLANs for segmentation.




FIGURE 2: The SAFE blueprint details the security elements deployed and the systems involved in a secure IP voice network.

be a phone, may exploit this to gain additional privileges or access to proprietary data. Turn this feature off except during initial mass provisioning, advised Halpern. Another feature allows the PCs plugged into IP phones access to carry out trunking, a feature not necessary in most deployments. A third allows the IP phone to copy all voice packets to the phone's data port for local voice troubleshooting or other applications. Again, this is not necessary in most cases and could lead to snooping should the PC plugged in the data port be remotely compromised. Turn these two off as well.

Finally, access control lists (ACLs) also help secure connectivity by enforcing the separation between Layers 2 and 3. Embedded in the software of all Cisco routers and firewalls and many switches, ACLs allow or deny access between voice and data VLANs based on the requestor's IP address and protocol/port information. If you're not on the list, you don't get in. They help to prevent unauthorized access from data to voice VLANs, among voice VLANs, or among modules—particularly useful against worms or viruses that automatically propagate themselves across the network. Halpern suggests that ACLs be activated in every capable device to enforce the segmentation.

Managing Trust and Identity

Once connectivity is as secure as you can make it, you should focus on identity and trust management: How do you know that your caller—or the person you're calling—is who he or she claims to be? New technologies and features from Cisco such as the combination of Dynamic Host Configuration Protocol (DHCP) snooping, IP Source Guard, and Dynamic Address Resolution Protocol Inspection (DAI) help you be sure. Mediating access through Cisco routers and switches, these technologies specifically protect against spoofing attacks by authenticating devices requesting access to the network.

DHCP snooping intercepts untrusted DHCP messages—which originate outside the network or beyond the firewall—to thwart basic and DoS starvation attacks against DHCP servers. With DHCP, you can statically assign an IP phone's IP addresses to a known MAC address so the IP phone always has the same MAC address; it's very difficult for a hacker to coordinate both addresses.

IP Source Guard is similar in that it blocks and filters all IP packets going through a specified port except for DHCP packets (which are vetted by the DHCP snooper). It passes only those with a DHCP-assigned address, preventing, for example, a malicious host from attacking a network by hijacking a neighboring host's IP address. DAI limits MAC addresses to one per port, mitigating sniffing attacks of various kinds.

For protection at the end user level, Cisco's newest IP phones can validate the veracity of downloaded software images through the use of digital signatures, a very important new feature. "A hacker could attempt to load invalid images on an IP phone in an attempt to make it inoperable or change its mode of operation altogether," Halpern says. "The latter is more unlikely, but it is a concern we addressed." With the digital signature, the IP phone is able to verify the origin of the image.

Threat Defense

The third focus is the management and prevention of threats; the key technologies are stateful firewalls, intrusion detection systems (IDSs), and intrusion prevention systems (IPSs). Halpern recommends that stateful firewalls be deployed primarily in two places—wherever voice and data segments meet, and wherever a stateful monitor is needed to protect voice services. "Protecting every system and every voice segment in the network with a stateful firewall isn't

manageable,” he says. “It’s better to use Cisco PIX® appliances to broker connections between, for example, a voice mail system in the voice segment and an email server in the data segment; IP phones on one voice segment connecting to a call processor in another; PC-based IP phones linking to the call processor; and where the voice gateway meets converged traffic.”

The PIX appliance’s role as stateful monitor comes into play when ACLs are charged with filtering dynamic port addresses rather than static ones. ACLs are not stateful, so when dynamic ports are used, as

Cisco offers IPS functionality via Cisco Security Agent (CSA), which is now supported on many Cisco voice products as well as on the Cisco VPN Client, to provide an additional layer of “day-zero” protection. Available in two versions, desktop and server, CSA provides intrusion prevention, distributed firewall, malicious mobile code protection, operating system integrity assurance, and audit log consolidation. Wherever Cisco SoftPhone is deployed, Halpern recommends deploying CSA to provide security in the data to voice segment interaction.

CISCO’S LEADERSHIP AND EXPERTISE IN DATA SECURITY SOLUTIONS MAKE ITS SOLUTIONS UNIQUE IN THEIR ABILITY TO PROTECT IP TELEPHONY AND VOICE OVER IP.

with the RTP protocol, large ranges must be opened to allow connectivity. Using protocol awareness capabilities, a stateful firewall provides single-port accuracy so large ranges of ports need not be opened. Cisco firewalls, both IOS and PIX-based, Halpern says, have extensive enhancements to address specifically the issues presented by IP voice.

IDSs are another must-have—both network-based IDSs (NIDSs) in the network and host-based IPSs (HIPSs) on hosts. Each provides a different type of protection. NIDSs watch packet streams for signatures, or sequences of bits, of known attackers, and like antivirus systems—also recommended—protect a network from known attacks whose signatures have already been mapped. They also detect protocol and other anomalies. “IDSs should be deployed to protect all key systems in the network, including voice mail and call-processing systems,” Halpern says.

HIPSs complement NIDSs. They take a behavioral, non-signature approach, looking for extraordinary events on hosts, and offer “day-zero protection from attacks,” Halpern says. “Even if the vulnerability patch isn’t applied to the host or an attack signature doesn’t even exist yet for the new attack, IPSs still mitigate the attack.” They work by monitoring behavior and detecting anomalies, and can actually prevent a variety of attacks before they occur, as in these examples, according to Halpern: “Why would a Web server’s port 80 process modify its configuration settings in the registry? Why would a Web server copy the shell executable cmd.exe into its Web script directory? Would a Web server ever need to run shell commands? These things should never happen, but we’ve seen both recently in worms that infected systems. IPSs stop these types of events from occurring.”

Filter, Filter, Filter

Once you’ve made provisions for secure connectivity, identity and trust management, and threat management and prevention, your overarching strategy should be “defense in depth,” Cisco’s mantra. Apply any or all security technologies across the network in any and all systems where they’re appropriate. Never rely on a sole mechanism for security. Thus, IDSs should be activated on routers, switches, servers, even on client systems such as PCs. ACLs, of course, should be employed on most of the same systems; Cisco IOS-based and appliance firewalls, where they’re of value—multiple technologies on multiple systems layered throughout the network.

Building a secure IP voice network is clearly possible, Halpern says. Already-available security technologies and products, careful design and implementation, and attention to detail will help to guarantee success. However, he points out, Cisco is the only vendor offering a comprehensive set of security products, strategies, and blueprints. (See “Further Reading” box for related URL links.) Finally, he notes, your byword should be “filter, filter, filter—whenever and wherever possible in the network.” ▲▲

FURTHER READING

- **SAFE: IP Telephony Security in Depth:**
cisco.com/packet/161_6c1
- **SAFE: A Security Blueprint for Enterprise Networks:**
cisco.com/packet/161_6c2

Enterprise

SOLUTIONS

Security Advances in Mobile IP

Cisco introduces a “Zero Configuration Client” and single login for Mobile IP users. **BY GENE KNAUER**


USING AND MANAGING CISCO'S Mobile IP offerings has become much simpler for both end users and administrators with the addition of Dynamic Security Association and Key Distribution features in Cisco IOS® Software Release 12.3(4)T. The features allow mobile users to log on securely in one step instead of several using their Microsoft Windows login. For network administrators, the new features are a big time saver because it is no longer necessary to separately configure Mobile IP client software in individual mobile clients and in the network (home agent) with mobility security credentials.

“The Dynamic Security Association and Key Distribution features build on the existing authentication infrastructure of the wired enterprise and extend it to the mobile environment,” says Richard Shao, a technical marketing engineer in Cisco's Internet Technologies Division. “To the end user, the main benefit is that the client (PDA, laptop, etc.) authentication and Mobile IP authentication processes do not require separate logins. For administrators, Mobile IP login and security features are set up automatically by the software. Security is actually enhanced because keys are generated dynamically.”

“Plug and Play” Comes to Mobile IP

Mobile IP, the Internet Engineering Task Force (IETF) standard and network architecture, gives a Mobile IP device the ability to retain the same IP address and maintain continuous network connection while traveling from one network to another. The Mobile IP functionality has been available with Cisco IOS Software since 1999.

Routing in wired IP networks is based on station-



ary IP addresses. When an IP node is mobile and roams away from its home network, normal IP routing is not possible. Active sessions are terminated.

Mobile IP is based on communication between a *mobile node* (for example, a laptop, cell phone, or PDA), a router on the home network called a *home agent* and optionally a router called a *foreign agent* that functions as the point of attachment for the mobile node when it roams to a foreign network. The foreign agent delivers packets from the home agent to the mobile node.

Features in Mobile IP allow for the mobile node to find the foreign and home agents while roaming and register its location. The home agent can then set up

Cisco Structured Wireless-Aware Network Solution

In November 2003, Cisco Structured Wireless-Aware Network (SWAN) rolled out new radio management features for Cisco Aironet® access points and version 2.5 of CiscoWorks Wireless LAN Solution Engine (WLSE). A secure, integrated WLAN solution of Cisco “wireless-aware” infrastructure products, Cisco SWAN minimizes WLAN total cost of ownership through optimized deployment and management of market-proven, high-performance, multi-function Cisco Aironet access points.

“Most of our competitors have dedicated WLAN appliances that are laid over the wired networks,” says Cisco product manager Anu Gade. “As these WLANs have grown, managing and securing them has become challenging and labor intensive. Cisco SWAN greatly strengthens and simplifies the deployment, operation, management, and security of enterprise WLANs. This solution supports flexible management of a few, to hundreds, to thousands of central or remotely located Cisco

Aironet Series access points from a single CiscoWorks WLSE management console.”

Florida Hospital, an acute-care health care system with 1,800 beds in central Florida, has already seen clear benefits from managing their Cisco Aironet WLAN using Cisco SWAN. “The decision to install CiscoWorks WLSE and Cisco SWAN was driven by our need to simplify deployment, decrease management and troubleshooting complexity, and quickly implement changes,” says Gil Sturgis, network services manager at Florida Hospital.

Cisco SWAN gives Florida Hospital administrators complete visibility into their WLAN from a single management console and a full set of management and security tools. It greatly enhances WLAN deployment, administration, and security for the more than 8,000 doctors, nurses, and other staff who use wireless laptops on mobile carts at seven hospitals.

Cisco SWAN enhances network

security by scanning and monitoring the RF environment to detect rogue (unauthorized) access points. Advanced security policy monitoring, alerts, and diagnostic tools, access point auto-configuration, and centralized security settings optimize the WLAN deployment. Cisco SWAN’s RF interference detection and assisted site survey capabilities maximize wireless network uptime, saving company time and money.

Cisco SWAN includes four core components: Cisco Aironet access points; CiscoWorks WLSE; an IEEE 802.1X AAA authentication server like the Cisco Secure Access Control Server (ACS); and Wi-Fi Certified WLAN client adapters. Functionality can be extended through Cisco Aironet and Cisco Compatible client devices and, in the future, through wireless-aware Cisco wired infrastructure products for integrated wired and wireless LAN capabilities.

For more information, visit cisco.com/go/swan.

a tunnel connection to send packets to the mobile node based on the notified location.

“When we originally released Mobile IP, you had to first configure each individual device and the AAA [authentication, authorization, and accounting] server with the security associations, including security parameter index, authentication algorithm, and pre-shared key,” recalls Cisco technical lead Alpesh Patel. “With the new Security Association and Key Distribution features, administrators don’t have to provision any client software or any configuration on the home agent the cost of ownership for Mobile IP goes down. Smaller environments will also be attracted to this simpler deployment that requires less staff and expertise.”


By using an existing authentication infrastructure, the Security Association and Key Distribution features require only the configuration of the AAA server so it can use the Windows Domain Controller database or Active Directory to authenticate the Mobile IP client

upon receiving a RADIUS request from a home agent (see figure).

“The enhancements also make Mobile IP more secure,” says Richard Shao. “Before, the security was static; unless you changed the security associations, they remained the same for the clients. Now, the keys are dynamically generated and different for each session and possibly during the session.”

Another feature of the Dynamic Security Association and Key Distribution enhancements to Mobile IP includes the ability to dynamically assign an IP address to the mobile device after user authentication. A user can log in to the Windows domain using multiple mobile devices. Each device would obtain a unique IP address for communications.

In response to a successful registration, basic configuration parameters such as the Dynamic Host Configuration Protocol (DHCP) server address, home address prefix length, and domain name system (DNS) address are also passed on to the mobile node



STREAMLINED DEPLOYMENT: The Dynamic Security Associations and Key Distribution feature in the Cisco IOS Software simplifies Mobile IP deployment by streamlining login, configuration, and provisioning. With this feature, the security associations do not have to be configured manually in advance; the Mobile IP client can derive the security associations from a user's Windows login name and password. The home agent authenticates the user from a Windows Domain Controller or Windows Active Directory. Once the user is authenticated, the home agent generates the user's security associations dynamically to perform authentication of Mobile IP registration.

in the form of extensions to the registration reply message that is sent by the home agent.

Ongoing Enhancements to Mobile IP

“IP was developed for wired networks, and Cisco sees Mobile IP as the means to translate all of the innovations from Cisco’s scalable, feature-rich, and manageable wired networking products and technologies to the mobile or wireless enterprise,” says Kent Leung, technical lead and Cisco Distinguished Engineer in the IP Routing Technologies Division. “These new security and login features for Mobile IP are just the beginning of many new enhancements to make mobile computing easier, more secure, and more functional.”

Just as an array of proprietary telephone systems was developed over time, a variety of wireless networks exists and pose challenges for engineers seeking to streamline Mobile IP integration.

“Legacy wireless networks were designed differently than wired networks, even within the same company,” says Leung. “Wired networks are logically and administratively partitioned, whereas most wireless networks are one big separate entity and mobility is managed at the physical or application layers. Mobile IP manages mobility at the network layer. The challenge is to enhance Mobile IP in successive releases of

Cisco IOS [Software] to interoperate with dozens of different types of wireless and wired networks and wireless devices.”

The goal of Mobile IP development at Cisco is to make the wireless portions of networks look and act the same as wired networks and to allow users to traverse multiple network segments while staying connected and using the same IP address. Growing use of Mobile IP for data roaming services supporting a broad array of business users, the military, researchers, and consumers is finally bringing the “anytime, anywhere” connectivity promise closer to reality.

Client software for Mobile IP is expected to greatly enhance the utility and appeal of Mobile IP, allowing users and applications to seamlessly connect and reconnect across different types of infrastructures without user intervention or down time. Cisco is at the forefront of these developments, supporting an open, standards-based approach that allows users to turn virtual private networks (VPNs) on and off; to support Network Address Translation (NAT), leading wireless standards such as Code Division Multiple Access (CDMA), General Packet Radio Service (GPRS), and IEEE 802.11b; the zero configuration client; and various operating platforms such as Windows 2000, Windows XP, and Windows Pocket PC. ▲▲

Read a white paper on Cisco Mobile IP Dynamic Security and Key Distribution at cisco.com/packet/161_7a1.

Deployment Diary

New York University's "Tetrahedron Core" goes live, provides unsurpassed capacity and resilience.

WHEN NEW YORK University (NYU) began implementing a novel IP architecture with unprecedented redundancy, resilience, and performance, the university's network architect, Jimmy Kyriannis, maintained a record of the process. In recounting the birth of the Tetrahedron Core, Kyriannis provides a rare glimpse of the gritty details behind a major network upgrade, including plans, progress, setbacks, and successes.

The Second Quarter 2003 issue of *Packet*® presented part 1 of Kyriannis' diary (cisco.com/packet/161_7b1), an account of how Kyriannis and his IT team conceived, planned, and began testing the Tetrahedron Core. Part 2 of the diary picks up with intensive testing in the testbed area, and continues through deployment and production.

April 22, 2003

We'll be running some intensive tests in the testbed area over the next two weeks. First we'll confirm that load balancing based on Cisco Express Forwarding over multiple equal-cost parallel links equally distributes traffic over all paths. If the test succeeds, we'll know that one path in the Tetrahedron won't be overloaded even while another is underutilized. We'll also want to check for proper routing behavior despite failure of any network component: individual fiber links, line cards, whole Gigabit EtherChannel® links, and even an entire Cisco Catalyst® 6513 Switch, which is used as a core router in the Tetrahedron Core. Finally, we'll test some representative IT applications to confirm that they behave exactly as expected when run on this unique topology.

If the tests succeed, we'll remove the Cisco Catalyst 6500 Series switches from our testbed and begin deploying them in their permanent locations across campus. At that point, we'll be ready to interconnect the switches with 64 pairs of fiber-optic cabling, just as they currently are interconnected in the testbed. But this time the fiber will stretch across the NYU campus; the total length of fiber-optic cabling used to implement the core network is approximately 39 miles.



DAILY
JOURNAL

May 19, 2003

The final tests were an unqualified success. We'll want to preserve information about the Tetrahedron Core's operational state at this point for future analysis and diagnostic purposes. Therefore, I took a snapshot of the Tetrahedron Core's AppleTalk, IPX, and IP/OSPF [Open Shortest Path First]

routing tables; Cisco Discovery Protocol neighbor information; Layer 2 and Layer 3 adjacency tables; and Gigabit EtherChannel status. This will serve as master reference data that will document the proper operational state of the core network.

May 26, 2003

We're moving along in terms of dismantling the test environment. A few things are difficult to control, such as scheduling contractors to physically relocate our equipment. We're also setting up times for fiber-optic technicians to establish the 128 fiber-optic connections in the Tetrahedron Core and confirm that they're all operating without unwanted loss. When those steps are complete, we'll run through some final tests before going into production service. With staff out tomorrow and Monday for the Memorial Day holiday, it's created some additional delays.

June 27, 2003

The four Cisco Catalyst 6500 Series core switches are now rack-mounted in their permanent locations, and we've begun establishing the 128 constituent fiber-optic connections (64 pairs). This will be a slow process because our group's top priority is ongoing network support activities, which take up most of our day.

July 21, 2003

Only about 15 percent of the cabling remains to complete. We ran out of necessary fiber-optic patch cables, of all things, and are awaiting a shipment to replenish our stock.

When we finish the cabling work, the university's IT department will begin testing. To help with initial testing, we're establishing a redundant connection between the Tetrahedron Core and NYU-NET—that is NYU's campus network—through a single Cisco 7500 Series Router acting as an access router. We're also attaching a second access router that we'll use



JIMMY KYRIANNIS,
network architect and
manager at NYU

to establish a test subnet. The university's IT department will use this subnet to run final tests of our applications before we begin migrating NYU-NET to the Tetrahedron Core for production network services.

We're anticipating two scenarios for migrating access routers from our existing backbone to the Tetrahedron. For some access routers, we can make the transition without impacting active client network traffic, by activating new links into the Tetrahedron and allowing the OSPF routing protocol to prefer the Tetrahedron over our FDDI [Fiber Distributed Data Interface] backbone. Other access routers have fewer available free slots or interfaces, and these we will need to take down to migrate onto the Tetrahedron Core. In these cases, we will schedule the transition during our regular Friday night maintenance window, when network usage is relatively light.

August 11, 2003

We've reached a milestone! Last Monday, Manny Laqui, our lead data technician, confirmed that the fiber interconnections among all four core routers are complete. Because of the highly critical nature of this network, we opted to certify the fiber connections end to end with a fiber analyzer, checking for signal loss and errors. We discovered a couple of bad fiber-optic patch cables, but nothing serious.

After replacing the cables, Manny established Tetrahedron Core connectivity to the first access router, which we selected by virtue of its placement on the current NYU-NET backbone network. This particular access router is a focal point for Internet connectivity at NYU. By placing it on the Tetrahedron Core first, we'll ensure that all routers we transition subsequently will have direct and non-blocking access—not impeded or constrained by bottlenecks or network congestion problems—to our high-speed Internet and Internet2 connections. These routers also will have access to the rest of NYU-NET through the first router's connections to the existing NYU backbone network.

As we connected the reassembled Tetrahedron Core to our first access router, we validated the connections and routing information. To do this, I took a snapshot of operational data from all four of the core routers and compared it to the reference data I had collected before the testbed environment was dismantled. The data matched perfectly! Later in the week I ran approximately 10 GB of trial data through the Tetrahedron Core links to double check for any errors on any of the component core links. Once again, our routers and links passed their tests with flying colors.

We're now ready for the final phase of testing before beginning production service. To prepare for the final tests, we're attaching a second access router to the Tetrahedron Core. We will place this test

router on the two tetrahedron faces opposite the two to which the current access router is attached. When those fiber-optic connections are complete, we'll migrate our entire NOC [network operations center] network and staff onto this test access router. At that point, all network management and monitoring systems, as well as production business traffic from our department, will run exclusively through the Tetrahedron Core.

Following phase one of this testing, the next step will be to migrate two additional IT departments onto the test access router to give our staff an opportunity to test drive the Tetrahedron Core. We'll want them to confirm that their applications and services behave as expected—and preferably with better performance characteristics. We'll continue to operate in test mode for a while so that our IT staff can run tests to validate our network's operation. I'll be on vacation, so two members of my staff—Keith Malvetti and Yoni Radzin—will conduct the trials and collect feedback from our test participants.

October 6, 2003

During testing we experienced a glitch relating to Gigabit EtherChannel. Two weeks ago, I visited Cisco headquarters in San Jose, California, to participate in the Cisco Technical Advisory Board. While I was there, Cisco development engineers working with the Cisco TAC [Technical Assistance Center] in Research Triangle Park, North Carolina, provided us with a copy of the Cisco IOS® Software Release 12.1(13)E10, which overcomes the glitch and will provide the greatest stability in our mission-critical production environment.

We were very pleasantly surprised that when we installed the upgrade to the Cisco IOS Software and rebooted according to Cisco recommendations, the network continued operating without any interruption or impact to users. And what's more, initial tests with the software ran perfectly.

Since then we have expanded our testing of the Tetrahedron Core to include production services for a good portion of our IT department. Absolutely no complaints, so it appears we have a very stable and reliable Tetrahedron Core in service!

We'll continue internal testing for a few more weeks, and then begin migrating some small departments to the Tetrahedron Core for production service on the NYU campus network. As our confidence builds, we'll begin using our weekly Friday night maintenance windows to migrate Layer 3 distribution routers that support thousands of clients onto the core.

November 3, 2003

For a couple of weeks our NOC network has been running entirely on the new Tetrahedron Core. So far, things are proceeding smoothly, and even simple

data transfer tests using FTP [File Transfer Protocol] have exhibited data rates in excess of 250 Mbit/s.

Having successfully exposed our own production networks to this new networking environment, we're now ready to transition other departments. In fact, this evening we successfully cut over our first two departmental production networks onto the Tetrahedron Core.

Our initial foray into migrating access routers will be to take a single departmental network off of the Cisco 7513 Router currently on our FDDI backbone network, and then attach that subnet to a Catalyst 6500 Series Switch (acting as an access router) attached to the Tetrahedron Core. Our plan is to move cautiously but aggressively, migrating all of the subnets off of the Cisco 7500 Series Router and onto the Cisco Catalyst 6500 Series within the next two weeks. When success is confirmed, we'll migrate entire Cisco 7500 Series and Catalyst 6500 Series appliances to the Tetrahedron Core—an approach that will significantly speed up the process. All told, we'll migrate approximately five hundred /24-equivalent IP subnets. Each /24 subnet can support up to 256 IP addresses.

November 14, 2003

The migration of our first access router onto the core was an unqualified success, with every department reporting reliable network and application performance. Already, after just two weeks of use, we're seeing increased network utilization levels on the subnet uplinks that were migrated—a result directly attributable to the dramatically higher backbone data rates available within the Tetrahedron. The net result is that university users can transfer data more quickly and in greater volumes.

During our maintenance window this evening, we'll migrate a complete Catalyst 6500 Series Switch from our FDDI backbone onto the core.

November 18, 2003

Our thorough testing over the summer continues to pay dividends in the form of a smooth migration. The migration of our second access router onto the Tetrahedron Core proceeded without a hitch. Data rates continue to climb on the core as traffic enters and leaves these access routers in a nonblocking fashion. Our users are noticing better-than-ever performance with high-bandwidth applications such as nightly network-level backups, streaming media, and distance learning.

December 1, 2003

As we transition the few remaining routers to the Tetrahedron Core, we're declaring success. We've met our goals: redundancy, resilience, performance,

and creating a next-generation campus backbone that we won't outgrow anytime soon. The tetrahedron topology, with the Cisco Catalyst 6500 Series Switch at its core, can scale to hundreds of gigabits per second.

We're anticipating three ways we can leverage the Tetrahedron Core to differentiate NYU in terms of academic research and academic activities. One is positioning us to roll out a campus-wide IEEE 802.11 wireless network, which students and faculty will be able to access from the library, lounges, study halls, and classrooms. That's an amazing feat for NYU, which, as the nation's largest private university, has tens of thousands of students and faculty and 130 buildings. We now have one of very few networks with the capacity for a campus-wide wireless rollout of this scale.

Another near-term plan is to facilitate MPLS [Multiprotocol Label Switching], which we use to support campus-wide services with unique routing requirements and to improve support for security domains. A particular subnet's traffic can be tagged and then forwarded based on the criteria that are appropriate for that network's needs. The high performance of the Cisco Catalyst 6500 Series Switch makes this application practical, while offering services such as encryption and firewalling.

Possibly the biggest benefit of the Tetrahedron Core to the university is that we can now take full advantage of Internet2, which enables much higher bandwidth applications for academic research and student projects. For example, recently our performing arts students used Internet2 to perform collaboratively and in real-time with students from another university: a flat-panel video screen behind our performers provided real-time video and audio of the other school's performance, giving the perception that the two groups were acting as one. The Tetrahedron Core delivers the extremely high bandwidth that makes this and other high-bandwidth Internet2 applications possible.

Finally, we are confident that our network can remain functioning despite many types of concurrent failures. During our testing, we've simulated fiber-optic failures, board failures, and many others. Through it all, the Tetrahedron Core adapted so that traffic continued to flow without any interruption to users.

Not only do we now have a next-generation campus backbone with the capacity for future high-bandwidth applications, we made the transition without sacrificing any of the hallmark reliability of our old FDDI network.

The best of both worlds! ▲▲

High Availability for Campus Networks

Why and Where to Deploy the Gateway Load Balancing Protocol

BY RICK WILLIAMS

PLANNERS AND MANAGERS OF ENTERPRISE networks build redundant links and devices into their topologies to ensure a very high degree of network uptime. However, most network operators would also like the ability to squeeze maximum utility out of those backup resources. In addition, as with most network designs, alleviating as much administrative complexity as possible from the configuration is an operational goal.

There are several design options using Cisco technologies for achieving these high availability and resource-efficiency objectives. From a redundancy standpoint, for example, automatic failover protocols have long been available to negate the impact of an outage at a first-hop access point—the device that aggregates traffic from multiple clients or other edge devices for forwarding. These protocols include Cisco's Hot Standby Routing Protocol (HSRP) and the Virtual Router Redundancy Protocol (VRRP), an HSRP-like protocol specified in Internet Engineering Task Force (IETF) RFC 2338. Both are supported in Cisco IOS® Software.

These protocols play an important role in increasing availability by providing dynamic switchover of the forwarding responsibility from one network device to another should an active switch or router—from here on called a *gateway*—become unavailable. The term “gateway” describes a forwarding device whose job is to link two dissimilar network segments, usually, the campus Layer 2 access network and Layer 3 distribution or core network.

However, HSRP and VRRP do not allow for the total use of network resources when multiple paths between the switched, access-layer network and the upper-layer IP network exist unless additional configuration steps are taken. HSRP, for example, allows a virtual IP address and virtual MAC address to be applied to a cluster of gateways, forming an “HSRP group.” But all clients associating with the cluster connect to the “active” HSRP device—the one that controls the virtual IP address—until a failure occurs. This scenario leaves the standby device idle most of the time.

To gain load-balancing benefits with HSRP, network operators must configure additional virtual LANs (VLANs) and subnets. As such, HSRP (or VRRP) is generally recommended in cases when just two devices and two upstream paths are being used.

Configurations for sharing traffic among redundant gateways using HSRP and other methods are described later in the section titled “Other Load-Balancing Alternatives.”

GLBP Design Scenarios

Enterprises aggregating traffic from many edge devices to several upstream paths have another option. The Cisco Gateway Load Balancing Protocol (GLBP), initially introduced in Cisco IOS Release

GLBP is an extension to HSRP designed to achieve the multiple goals of hot standby, load balancing, and configuration simplicity by segmenting traffic automatically across resources.

12.2(14)S and 12.2(15)T, automatically provides for load balancing across the multiple redundant gateways and uplinks installed for backup purposes in a common subnet.

Load sharing allows enterprises to take full advantage of all available resources and thus get more out



RICK WILLIAMS

RICK WILLIAMS is a technical marketing engineer in the Internet Technologies Division at Cisco. He can be reached at rwill@cisco.com.

of their capital investments. GLBP is an extension to HSRP designed to achieve the multiple goals of hot standby, load balancing, and configuration simplicity by segmenting traffic automatically across resources.

GLBP, which has heretofore been used primarily in enterprise WAN access routers and at the edge of service provider networks, has recently also become available for Cisco Catalyst® 6500 Series switches, when the Supervisor Engine 720 is in use.

The purpose of GLBP is to allow traffic load sharing across uplinks that connect access-layer endpoints and Layer 3 gateways, which support device failover protection.

There are two spots in the network where the use of GLBP is highly recommended: in the campus network, in an aggregation gateway where the Layer 2 switched access layer meets the Layer 3 routed layer (see figure); and at the WAN edge of remote sites, where dual- or multi-homed connections are in place for high availability networking.

GLBP can also be used to share outbound network traffic from server farms. However, in this scenario, GLBP is recommended for environments that do not require additional server load-sharing features to balance the resource consumption of the hosts themselves—a capability that is best solved using Cisco content switching software, modules, and appliances.

How Does GLBP Work?


Multiple gateways in a “GLBP redundancy group” respond to client Address Resolution Protocol (ARP) requests in a shared and ordered fashion, each with their own unique virtual MAC addresses. As such, workstation traffic is divided across all possible gateways.

A controller called the *active virtual gateway* (AVG) assigns a unique virtual MAC address to each device in the group and responds to ARP requests on behalf of the group. The other members of the group are backups to the AVG. They also serve as *active virtual forwarders* (AVFs), in that they are responsible for forwarding traffic that the AVG matches to their virtual MAC address using ARP.

ARP is a protocol for mapping an IP address to a physical machine (MAC) address that is recognized in the local network. A table, usually called the *ARP cache*, maintains a correlation between each MAC address and its corresponding IP address. ARP provides the protocol rules for making this correlation and providing address conversion in both directions.

When an incoming packet destined for a host on a particular LAN arrives at a gateway, the gateway asks the ARP program to find a MAC address that matches the IP address. The ARP program looks in the ARP

GLBP CAMPUS DESIGN



LOAD EFFICIENCY: When deployed at the campus distribution layer, GLBP enables enterprises to make better use of their upstream Fast Ethernet or Gigabit Ethernet links without having to manage and segment clients among different Layer 3 default gateway addresses and VLANs.

cache and, if it finds the address, provides it so that the packet can be converted to the right packet length and format and sent to the LAN-attached device.

If no entry is found for the IP address, ARP broadcasts a request packet in a special format to all the devices on the LAN to see if one knows that it has that IP address associated with it. A device that recognizes the IP address as its own replies. ARP updates the ARP cache for future reference and then sends the packet to the MAC address that replied.

There is a Reverse ARP (RARP) for hosts that do not know their IP address. RARP enables them to request their IP address from the gateway’s ARP cache.

Other Load-Balancing Alternatives

There are three primary ways to eliminate a single point of failure in the forwarding path while also allowing traffic loads to be shared among gateways. The simplest is GLBP, in that GLBP has both capabilities inherently built into the protocol. As such, it also achieves the third desired benefit, administrative simplicity, as GLBP-enabled devices require no special tuning to split and share traffic.

The other two network load-sharing options are *Layer 2 Multi-VLAN* and *Layer 3 Multi-Group HSRP* (MHSRP), which each require a bit of tweaking.

- **Layer 2 Multi-VLAN.** To achieve both redundancy and load sharing at Layer 2, network operators can use a Cisco IOS Software feature called *Per-VLAN Spanning Tree* (PVST). In this design, two or more VLANs are created on the access-layer switch; half

Continued on page 56

Deploying GLBP, Continued from page 53

CAMPUS ALTERNATIVES FOR REDUNDANCY AND LOAD BALANCING		
Multi-VLAN (Layer 2)	Multi-Group HSRP (Layer 3)	GLBP (Layer 3)
Uses Per-VLAN Spanning Tree, whereby two or more VLANs are created on the access-layer gateway. Half the VLANs use the forwarding path through one uplink, and the remaining VLANs use the other uplink.	Two or more Layer 3 gateways on a common subnet share a virtual IP address and a virtual MAC address to form an HSRP group. A gateway that is "active" for one group is a "standby" for another group. Clients are configured to forward traffic to different default upstream devices, which thus share the load.	Multiple gateways in a GLBP redundancy group respond to client ARP requests in a shared and ordered fashion, each with their own unique virtual MAC addresses. Thus, workstation traffic is divided across all possible gateways. A controller called the active virtual gateway assigns the virtual MAC addresses and responses to ARP requests on behalf of the group.

the VLANs are forwarded across one uplink and the second half use the other.

One distribution-layer gateway serves as the "root bridge" for the first VLAN and another serves as the root bridge for the other. Both distribution-layer gateways are linked with an additional trunk that carries both VLANs. Each access-layer switch, then, forms the point of a triangular loop with two distribution-layer gateways; however, the Spanning Tree Protocol removes the loop, leaving the desired forwarding arrangement and retaining redundancy.

- **Layer 3 MHSRP.** As mentioned, basic HSRP allows for failover among devices within an HSRP virtual router group, but does not inherently support general load sharing. However, multiple HSRP groups can be configured on a common subnet to provide this capability.

Use of network protocols that provide for hot failover and load sharing enable enterprises to eliminate single points of failure and achieve high levels of network availability during the long periods of time when there are no failures.

In this configuration, two or more Layer 3 gateways on a common subnet share a virtual IP address and a virtual MAC address to form an HSRP group. Each gateway's LAN interface is configured with two HSRP groups. A gateway that is "active" for one group is a "standby" for another group. Clients are configured to forward traffic to different default upstream devices, which thus share the load.

There is some administrative overhead associated

with dividing client devices on a common subnet among multiple gateways. Another option is to combine Layer 2 Multi-VLAN load balancing with MHSRP, in which each LAN interface on the upstream gateways are configured as trunks, and each trunk can carry two VLANs. Multiple HSRP groups are defined so that one Layer 3 gateway will be active for odd-numbered IP subnets while the other will be active for even-numbered IP subnets.

Why Use Network Load Sharing?

Use of network protocols that provide for both hot failover and load sharing enable enterprises to eliminate single points of failure and achieve high levels of network availability without having backup resources go unused during the long periods of time when there are no failures.

Using GLBP for this purpose can reduce congestion in output queues during peak times, because multiple paths are automatically used for upstream forwarding.

Because GLBP has been designed to provide first-hop, Layer 3 redundancy plus load sharing over multiple devices and paths, enterprises can gain better performance and greater utility out of their capital investments with less of an administrative challenge than has been previously possible. ▲▲

FURTHER READING

- **Gateway Load Balancing Protocol:**
cisco.com/packet/161_7c1
- **Cisco Catalyst 6500 Series switches (supporting GLBP):**
cisco.com/packet/161_7c2
- **High availability white papers and best practices:**
cisco.com/packet/161_7c3

Extending SANs


Cisco DWDM platforms and storage fabric switches meet demand for business continuance. **BY DAVID BARRY**

WITH WIDESPREAD POWER OUTAGES, potential threats of terrorism, and natural disasters continuing to make headlines, business continuance and disaster recovery planning has moved to the top of the agenda at many organizations. The cost of downtime can be immense—almost US\$1 million of potential lost revenue per outage-hour for synchronous mirroring applications, according to Gartner Group. But this cost pales in comparison with the longer-term damage to a company caused by customer dissatisfaction, brand dilution, and legal liabilities due to an inability to deliver services or products for which customers contracted. In fact, according to a recent Gartner study published in *Disaster Recovery Journal*, seven out of 10 companies that encountered a disaster and had no disaster recovery plan in place were out of business within two years.

Today, companies with a business continuance plan are addressing the issue with optical metro dense wavelength-division multiplexing (DWDM) combined with storage-area network (SAN) extension solutions. Together, these two technologies provide disaster recovery for high-end synchronous mirroring applications across distances where no downtime can be tolerated. In these designs, synchronous replication is used to ensure that multiple synchronized copies of data exist at multiple sites or data centers. Every write-to-disk operation is synchronously replicated across the network to a storage array in an alternative data center. The synchronous application that resides on the intelligent controller waits for both disk drives to complete writing data before it returns an acknowledgment to the input/output (I/O) requestor or initiator.

For synchronous mirroring to be successful, it requires that synchronous replication of data occur in real time and that accurate up-to-the-second data must exist in the standby data center in the event of a fault in the production data center. Synchronous replication is optimized for local high-speed connections with low latency and, when properly configured, operates effectively across a metro optical network that provides the low latency performance required.

The demand for this type of service is growing rapidly. Market researcher IDC projects that the DWDM metro market will grow from US\$600 million in 2003 to US\$1.5 billion in 2007—and that 30 percent of this growth will be for SAN extension technologies.



Cisco Metro Optical DWDM Platforms

The two Cisco metro optical DWDM platforms that are optimized for extended storage applications are the Cisco ONS 15540 Extended Services Platform (ESPx) and the Cisco ONS 15530 Multiservice Aggregation Platform. The Cisco ONS 15540 supports industry-leading wavelength aggregation of 32 wavelengths per fiber, and offers 1+1 protection for a total of 64 wavelengths. It can carry 10 Gigabit Ethernet, Gigabit Ethernet, ESCON, FICON, Fibre Channel, and SONET/SDH services. Scalability is built in with the ability to grow from 2.5 Gbit/s to 10-Gbit/s wavelengths.

The Cisco ONS 15530 DWDM platform excels at providing multiservice aggregation onto a single wavelength. It can transmit up to eight Gigabit Ethernet, Fibre Channel or FICON services, or 40 ESCON services per wavelength. This high density translates into capital and operational savings for customers who can utilize each wavelength more efficiently and can operate with fewer platforms. The Cisco ONS 15530 also supports non-storage traffic including SONET/SDH and ATM at OC-3/STM-1, OC-12/STM-4, and OC-48/STM-16. All of these disparate services can be multiplexed together onto a single wavelength providing multiservice transport.

Figure 2 illustrates a mixture of both platforms in a data center replication application. The Cisco

FIGURE 1: Synchronous Replication ensures that multiple synchronized copies of data exist at multiple sites or data centers.

Cisco and IBM Collaborate to Reduce Complexity and Administration of Storage Networks

In October 2003, Cisco and IBM announced a jointly developed solution that tailors IBM's TotalStorage SAN Volume Controller storage software to be hosted on the Cisco MDS 9000 Multilayer Intelligent Family of SAN directors and fabric switches.

The new solution allows customers to manage SAN volumes, handle data replication, and create point-in-time copies directly from the network—giving them a single point of control and management across multiple storage subsystems.

The solution will allow IT administrators to consolidate, or "virtualize," multiple and disparate storage subsystems into logical pools. By dynamically provisioning storage from the SAN switches, users can avoid taking servers or storage subsystems offline as changes are required. This capability is intended to reduce the

amount of planned downtime normally associated with storage infrastructure changes and helps simplify administrative tasks. The ability to manage heterogeneous storage environments also allows customers to use lower-cost storage devices for their applications or business functions.

Cisco has developed a new Caching Services Module (CSM) for the Cisco MDS 9000, and it will host the IBM TotalStorage Volume Controller Storage Software. The CSM, which includes two independent nodes for IBM's software, offers eight gigabytes of cached memory and is hot-swappable, with redundantly configured components for high-performance, low-latency, highly available SAN transactions.

The IBM SAN Volume Controller for Cisco MDS 9000 is available through IBM and IBM Business Partners.




FIGURE 2: The Cisco ONS 15540 offers industry-leading wave-length aggregation per chassis. When combined with the Cisco ONS 15530's unmatched services aggregation capacity, enterprises can deliver rapidly scalable ESCON, Fibre Channel, and FICON storage services across a metro optical DWDM network.

ONS 15530 aggregates multiservice traffic, including multiple Gigabit Ethernet data services from a Cisco Catalyst switch, ESCON services from a storage array, and Gigabit Fibre Channel from Cisco MDS 9000 multilayer storage switches (that aggregate traffic from storage-area networks). In turn, the Cisco ONS 15530 transmits this aggregated multiservice traffic onto two 10 Gigabit Ethernet wavelengths to the Cisco ONS 15540. The Cisco ONS 15540 also aggregates wavelengths from IBM zSeries Servers using new technology developed by IBM and Cisco that extends the distance of the IBM 9037 Sysplex Timer to Server link (ETR). IBM initially designed the ETR to transfer data across local distances of up to 40 kms.

Through joint technology development, Cisco and IBM developed software that extends ETR to 100 kms—an important metric that is stipulated by the US government for certain financial transactions as outlined in the government document, *Intra-agency Paper of Sound Practices to Strengthen Resilience of US Financial System*.

Business Continuance Application Providers

Architecting and deploying business continuance and disaster recovery solutions demands specialized expertise and highly trained resources. Cisco has certified the leading business continuance application service providers to deploy the Cisco ONS 15540 and Cisco ONS 15530 optical DWDM platforms—IBM Global Services, AT&T Business

Solutions, British Telecom, Qwest, Verizon, SBC, and Giant Loop.

Cisco has also achieved certification status for the Cisco ONS 15540 and Cisco ONS 15530 platforms to operate with the market-leading storage applications. These include IBM Geographically Dispersed Parallel Sysplex (GDPS), EMC Synchronous Remote Data Facility (SRDF), HP StorageWorks Continuous Access EVA, and Hitachi Data Systems (HDS) True Copy.

The certifications are as follows:


- Cisco ONS 15540 SAN Qualification
- ✓ IBM GDPS—1G and 2G FC/FICON (100kms) and ESCON (40kms)
- ✓ EMC SRDF/Mirrorview—1G FC, 2G FC and ESCON (200kms)
- ✓ HP CA-EVA / DRM—Fibre Channel (100kms)
- ✓ HDS TrueCopy—Fibre Channel and ESCON (100kms)
- Cisco ONS 15530 SAN Qualification
- ✓ IBM GDPS—1G and 2G FC/FICON (100kms) and ESCON (40kms)
- ✓ EMC SRDF/Mirrorview—1G FC and ESCON (200kms) [2G FC testing in progress]
- ✓ HP CA-EVA/DRM—Fibre Channel (100kms)
- ✓ HDS TrueCopy—Testing in progress

For more information about Cisco's business continuance networking solutions, visit cisco.com/packet/161_7d1.▲▲

Learn more!
Read the white paper, *Understanding Alternatives for Extending Storage- Area Networks*, at cisco.com/packet/161_7d2.

Service Provider

SOLUTIONS



Triple-Play Cable

Time Warner Cable rolls out Cisco enabled voice over IP phone service to millions of residential customers.

THE ANNOUNCEMENT LAST YEAR THAT Time Warner Cable (timewarnercable.com) will offer voice over IP (VoIP) telephone service to a huge customer base, allowing those customers to bypass traditional phone companies, signaled the start of a technological shift that could change the cable industry—one of the biggest and most important industries in the US economy.

Service providers delivering broadband access are expanding their service portfolios to include voice, as well as broadcast video and high-speed Internet access—a “triple play” of services all in a single package. While many cable operators are still in the trial

stages with their voice offerings, Time Warner Cable was one of the first to enter the emerging market, in May 2003 introducing its Digital Phone residential voice service to customers in Portland, Maine. At the end of the year, in a move that further demonstrated its commitment to service rollout, Time Warner Cable announced partnerships with both MCI and Sprint to transport its calls over their long-distance networks.

An important milestone, both for broadband service providers and residential telephone customers, the Time Warner joint venture with MCI and Sprint to deliver VoIP is as significant as the first rollout of high-speed Internet access broadband services, or the

RON CHAN

launch of HBO, and represents a huge opportunity for Time Warner to expand its service offering to customers in some 31 markets. That expansion will almost certainly result in a prospective revenue stream with strong growth potential.

By adding the new service to its existing advanced multiservice network infrastructure, Time Warner Cable gains a new revenue stream—with minimal investment in new technology—while reducing customer churn and laying the foundation for continued innovation of its communications service portfolio. The bundle also enables the cable company to maximize the return on its IP network investment and increase the average revenue per user.

Digital Phone

For customers, the Time Warner Cable VoIP service affords more choice in residential voice service providers. The cable company’s Digital Phone service offering provides all the same features as local carriers, including important capabilities like emergency 911 services and provides a single monthly phone bill for all its triple-play services.

Subscribers to the new service get unlimited calling anywhere in the continental US as well as standard features such as 411 directory assistance, 611 service calls, enhanced 911 service, operator-assisted calls, and the ability to view call detail records online. The package also includes caller ID, call waiting, and call waiting ID. The cost: a flat fee of US\$39.95 per month.

The initial network deployment is being rolled out to more than 100,000 Time Warner Cable customers. As the service expands, the network can easily be upgraded to provide capacity for millions of users across Time Warner’s network. In the near future Time Warner Cable expects to begin taking advantage of the unique capabilities of a multiservice infrastructure to offer integrated services such as unified messaging, and caller ID notification on the TV set, as well as other features like voice recognition and variable dial tone (sports, weather, or news instead of standard dial tone).

The Triple Play Win

The VoIP milestone is also great news for consumers because they can have the convenience of receiving all bundled services from one company, and they’re being offered competitive rates.

Gerry Campbell, senior vice president of Voice for Time Warner Cable, says the company—along with other cable operators that are rapidly following suit—sees a clear competitive differentiator to being able to offer the service bundle in such a highly competitive market. “As cable operators, we’re able to offer a combination of services that potentially no one else can. We can offer a full suite of services, whereas the

CABLE SERVICES COMPARISON

Service	Monthly Churn
Cable TV only	1.5%
Cable TV and high-speed data	1.1%
Cable TV, high-speed data, voice	0.7%

consumer would have to go to two companies to get the same range of services from our competitors.”

In an In-Stat MDR analyst report, Cox Communications was reported as saying that monthly churn is demonstrably lower for consumers who subscribe to cable TV, high-speed data, and voice from a single service provider. Cox estimates the cost per churned customer to be US\$50 for cable TV and US\$150 for high-speed data. It concludes that voice adds about US\$1 of free cash flow per subscriber per month from reduced churn on cable TV and high-speed data—figures that can have a huge impact on the profitability of a service offering.

The Right Network for the Job

To add telephony to its bundle, Time Warner Cable needed a communications-grade infrastructure for voice, video, and data. For cable operators, preparing their infrastructures to support voice introduces some challenges. For example, providing toll-quality telephony requires a cable network infrastructure capable of delivering the quality of service (QoS) required for voice communications.


Time Warner’s Digital Phone primary-line service is built on top of the company’s existing IP infrastructure for high-speed Internet access. The enhanced network enables reliable voice services that can scale as the cable company’s Digital Phone service grows.

The standards-based IP technology provides many advantages. For example, a standards-based packet network makes it much easier for Time Warner Cable to integrate new multimedia and interactive capabilities. Further, standards-based equipment will endure in the operator’s network for many years without becoming obsolete.

Open-platform, flexible architectures are important because they help enable interoperability among different vendors. “As the cable industry rolls out IP communications, by using common platforms we’ll be able to interconnect our systems to one another, greatly expanding the reach of our services. Also important, because this solution is standards-based, we’ll have a choice of different vendors that support the same technology,” says Campbell. “And the cost to deliver the service will continue to go down.”

A standards-based infrastructure also provides the foundation for future services. “Because cable is all one common network platform, companies that

SERVICES ON THE CUTTING EDGE: Capitalizing on IXC local points of interconnection, its own IP-based broadband cable system, Time Warner Cable can deliver local and long-distance telephony services more efficiently, at a lower cost, and with the reliability and quality of service that customers require.



develop innovative services will reap the reward of being able to deploy these services across the cable industry as a whole,” Campbell predicts.

PSTN Interconnection

Time Warner Cable has been able to address the requirements for PSTN interconnection by partnering with two established interexchange carriers (IXCs): MCI and Sprint. Through the partnering arrangement, Time Warner Cable’s regional IP networks can interconnect to the nationwide networks of these established IXCs to transport calls outside of its service area. This handoff occurs directly to the carrier’s existent time-division multiplexing (TDM) switches, but in the future the handoff will occur as an IP call, enabling the carriers to cost-effectively transport the call through their network before egressing to the PSTN.

“Through these relationships, Time Warner Cable can immediately take advantage of a nationwide network and offload complex PSTN interconnection requirements, facilitating the rate at which they can turn up service in new markets,” says Rob Kissell, product manager for the Cisco Broadband Local Integrated Services Solution (BLISS) for cable, the solution that Time Warner is using to deliver IP services.

Cisco Voice Solution

Cisco BLISS is a fully integrated, tested, and Cisco supported residential voice solution consisting of

Cisco and partner products that make up a PacketCable-based architecture. Cisco BLISS enables cable operators to deliver IP-based voice services over their CableLabs DOCSIS IP network infrastructure. The Cisco BTS 10200 Softswitch, the Cisco uBR7246VXR Universal Broadband Router, and the Cisco MGX™ 8850 Voice Gateway products are key elements of the network infrastructure.

Softswitches provide the call-control plane and service intelligence for delivering telephony over packet networks. Because softswitches appear as an application server on IP networks they enable service providers to support new, enhanced communications services and features by melding voice communications, data communications, and entertainment

Continued on page 88

FURTHER READING

- Cisco cable product information: cisco.com/packet/161_8a1
- Cisco cable-ready solutions for VoIP: cisco.com/packet/161_8a2
- Cisco BLISS for cable solution: cisco.com/packet/161_8a3
- Article on the Cisco BTS Softswitch 10200 for triple-play cable: cisco.com/packet/161_8a4

Proactive Protection

New techniques and best practices help service providers counter increase in cyber attacks.

BY DAVID BARRY

WHEN THE SQL SLAMMER WORM struck on January 25, 2003, it infected 150,000 to 200,000 servers worldwide within hours of its first appearance, according to Network Associates' AntiVirus Emergency Response Team (AVERT). Customers of the Canadian Imperial Bank of Commerce in Toronto, Canada, and the Bank of America in the US were unable to withdraw funds using automated teller machines during part of the day following the attack and into the next day. And in South Korea, most of Korea Telecom Freetel's and SK Telecom's fixed-line and mobile Internet services failed, stranding millions of the country's Internet users.

Cyber attacks show no signs of abating. Carnegie Mellon University's CERT Coordination Center for reporting Internet security problems states that through the end of September 2003, users and ISPs reported 114,855 security breaches—32,761 more than during all of 2002. These reports include all types of security policy violations, from distributed denial-of-service (DDoS) to hacker attacks.

As the level and sophistication of these attacks increase, they are no longer focused solely on disrupting customer networks and data—many now target the Internet infrastructure itself. As more service providers offer converged services, including voice, the importance of network defense escalates. This migration of the core network is critical to their financial success and competitive differentiation but cannot come at the cost of reliability.

Service providers are quickly realizing that failure to respond to this new threat can be costly in terms of lost revenue caused by network downtime and disruption to customer service. At one time content to provide only “transport,” many now realize they must be more proactive in combating these attacks.

Sprint, for example, deploys new tools in its network to better defend itself against attacks. Sprint is specifically focusing on distributed DoS mitigation and intrusion-detection products that it plans to deploy in its backbone within the next year, says John Pardun, senior product manager of network-based IP VPN and security services at the carrier.

Sprint plans to offer customers an “additional level of monitoring and mitigation” as an add-on service that it will charge for, Pardun says. Both MCI and AT&T

also say they will charge customers for their planned distributed DoS mitigation and reaction services.

Security as a Process, Not a Point Product

Because service provider infrastructures are so large and the potential for attacks exists at multiple points in the network, Cisco recommends a multifaceted defense approach.

Barry Greene, a consulting engineer at Cisco, is a leading security thinker within the service provider community and a major force in the founding of NSP-SEC (puck.nether.net/mailman/listinfo/nsp-security), an online forum for security engineers at service providers to exchange information and coordinate response activities during malicious attacks.

“Service provider security is a life-cycle process that must encompass continual preparation, communications with colleagues, and reassessments to keep pace with the changing nature of security attacks,” says Greene. “Service providers must continually prepare, test, and deploy new approaches and continually reassess what worked and what didn't after each attack.”

Cisco, in collaboration with its service provider customers, has developed a six-stage security framework for service providers that combines industry best practices with specialized routing techniques and tools:

- Preparation
- Identification
- Classification
- Traceback
- Reaction
- Postmortem

Top Priority: Preparation

Service providers have little chance of successful defense against a malicious worm or virus attack unless they have laid the groundwork before an attack. “Trying to respond to an attack without preparation is like trying to line up a militia to combat an invasion when the invaders are already crawling up the beach,” says Roland Dobbins, a network engineer in the Cisco IT Internet Services Group.

Aggressive preparedness by Cisco IT enabled the company's entire network to escape unscathed during the Slammer attack that devastated other companies. According to Dobbins, “We did not lose a single packet on the production network due to the Slammer

SECURING POP AND CORE INFRASTRUCTURE




FIGURE 1: Anomaly-detection systems, such as those from Arbor Networks, use raw packet data from NetFlow to detect and characterize undesirable traffic, such as DoS attacks. In addition to detecting suspicious traffic, they can also recommend where filters should be applied to limit attacks.

worm.” As a major content provider, e-commerce site, and ISP for 32,000 people worldwide (and with 93 percent of its revenue booked online), Cisco takes security of its worldwide network seriously (see sidebar, “Cisco Security Best Practices,” page 67).

“Our motto is that ‘proactive work buys us time to be reactive,’” says Dobbins. “We owe our success during Slammer to the fact that we had the systems, communications path, escalation path, and processes already in place. Our rules state that if the attack is rated a P1—the highest level—management must be updated every 15 minutes, with additional network engineers brought on continually until the attack is under control.”

A critical part of preparation is the creation of a security or incident response team, which must be given the necessary resources and authority to combat attacks, and a team member must be available 24 x 7. As part of a response plan, team members should also establish relations with security team members at other peering partners and IXCs so that security engineers can be reached quickly at any hour. This enables teams to rapidly share information in the early stages of an attack and can help to determine the real source of the attack and possible remediation.


Relationships with key infrastructure vendors should also be established. Cisco maintains the Product Security Incident Response Team (see “Security Advocates,” page 36) as a key resource for both reporting security incidents and for obtaining information

about newly discovered vulnerabilities (security-alert@cisco.com for emergencies or psirt@cisco.com for nonemergencies).

One of the problems with communication among top engineers at various service providers is that often these engineers are not easily reachable—and they might not answer phone calls during an emergency. To help ensure that top security engineers at service providers can reach one another during an attack, a dedicated Session Initiation Protocol (SIP)-based network operations center (NOC) hotline system was created. This voice over IP (VoIP) system, managed by the nonprofit research institute Packet Clearing House (pch.net) and sponsored by Cisco, provides a direct connection to the engineers on the incident response teams at each service provider. When the NOC hotline rings, the NOC engineer knows it is one of the NOC engineer’s peers.

The online security forum NSP-SEC maintains a closed security operations alias and provides a vehicle for security engineers within NSPs and service providers to communicate and share information before or during attacks. During the Slammer attack, many of the participants in what is known as the “Skitter Group,”—the top Tier 1 and Tier 2 service providers—credited the NSP-SEC with helping to mitigate the impact of the worm. NSP-SEC was the first to report the worm to the general public and, in fact, CERT/FIRST teams received their alerts from NSP-SEC.

FIGURE 2: Backscatter traceback uses a network's own routing protocols to identify where an attack enters the network when the attack is using "spoofed" source addresses. When the source of the attack is identified, service providers can apply mitigation techniques. This technique was pioneered by UUNET Operational Security (OPSEC) engineers Chris Morrow and Brian Gemberling, and is offered to the greater Internet community as a tool to trace back an attack to the edge of an organization's network.



Detect the Threat

The ability to quickly identify an attack is critical to minimizing the damage that the virus or worm can ultimately cause. But aside from waiting for customers to deluge your NOC with complaints or your network management alarms to begin ringing in the NOC, how can you detect an attack?

Cisco and several service providers extensively deploy a NetFlow-based anomaly detection system that uses NetFlow data exported from Cisco routers to Arbor Networks PeakFlow Traffic and PeakFlow DoS systems. The system allows organizations to detect and characterize undesirable traffic such as DoS attacks (see Figure 1).

NetFlow Data Export must be enabled and the Arbor Peakflow system must be deployed before an attack. The first step is to use Arbor Peakflow Traffic to characterize network traffic; that information is used to set initial thresholds for the Arbor Peakflow DoS system, which dynamically updates traffic profiles and alerts designated personnel when anomalous traffic patterns such as SYN-floods, fragmented traffic streams, and the like, are detected. Based on information from the Arbor system

(anomaly type, sources, destinations, protocols, ports, packet sizes, packets per second, specific routers and interfaces involved in the attack, for example) operational personnel instantly receive the information they need to respond effectively.

“Our Arbor-based anomaly detection system is a tremendous force multiplier,” says Dobbins. “During the Slammer attack, it took us straight to the router and interfaces where we could see an abnormally high amount of UDP traffic flows coming in through port 1434. We could characterize the type of attack and where it was coming from within minutes. This gave us a tremendous time advantage to apply ACLs [access control lists] to our border routers at our POPs [points of presence] worldwide—while many organizations were still working to characterize the attack and understand its ramifications.”

Traceback

Once a service provider has detected an attack, often the next step is traceback—trying to determine the source of the attack so that the service provider can apply mitigation techniques or, if the source of the

Cisco Security Best Practices

When the SQL Slammer virus hit last year, Cisco used the following six-phase implementation to prevent damage to its own network.

Preparation. People, processes, procedures, lines of communication, architecture, and automation tools all need to be in place *before* an event occurs. Security and networking teams must work together smoothly, with no contention over authority. A duty manager is available around the clock to make business decisions as needed, and engineers are empowered to take decisive actions. Professional relationships were established with Cisco's Product Security Incident Response Team; the Cisco Technical Assistance Center (TAC); Cisco Advanced Services; and ISPs, peers, and customers that run other large networks. Cisco participates in FIRST. A detailed communications and escalation plan, facilitated by the operations group and used daily, is in place and well understood.

Identification. Use of Cisco and Cisco partner products and technologies (NetFlow on routers and switches exporting to Arbor Peakflow Traffic and Peakflow DoS anomaly-detection system) allows Cisco to know what is normal for its network, and what is abnormal (for example, unusually high numbers of UDP/1434 traffic flows) and potentially hostile.

Classification. Knowledge of Cisco's own network architecture, network traffic patterns, systems and input from Arbor/NetFlow instrumentation allows Cisco to quickly classify and scope threats.

Traceback. Instrumentation plus knowledge of the Cisco network allows Cisco to identify all presently visible and potential sources and vectors of the attack. This proves to be critical; in the case of Slammer, many organizations failed to account for indirect vectors, such as virtual private networks (VPNs) and laptops, that carried the virus into companies on Monday morning.

Reaction. Cisco immediately "dropped the shutters" through the use of ACLs at all Internet POPs worldwide. A well-designed, "bulkheaded" network allowed a pause so that Cisco could determine its follow-on actions. Knowledge of the virulence and threat level caused Cisco to push ACLs down to the desktop level in every Cisco facility worldwide, along with strategically placed ACLs in the Cisco WAN backbone. This ensured that Cisco wasn't affected on the following Monday. Operations teams provided focal point and bridges for all intergroup communications—network engineers didn't have to search for telephone numbers while dealing with mitigation. Thorough, complete, draconian, and pervasive ACLs were essential.

Postmortem. Cisco conducted daily followup sessions for two weeks to ensure that the threat was eradicated and to discuss lessons learned. Management issued the directive to prioritize and implement lessons learned in concrete, measurable ways, with meaningful followup to ensure future success.

attack is from another network, inform the respective peer.

Attacking sources fall into two groups: valid source addresses and spoofed source addresses.

With valid sources, traceback is not required. Instead, service providers can query online databases such as ARIN's WHOIS database (www.whois.net) to determine address ownership. Additionally, they can use network utilities such as Domain Name Service and traceroute. But these techniques take time and might not be able to determine the source address if the attack is using

spoofed addresses. *Backscatter traceback* is the commonly used technique for rapid traceback (see Figure 2). Backscatter takes advantage of Border Gateway Protocol (BGP), the routing protocol pervasively deployed in service provider networks, to drop traffic originally destined for the victim and enables the creation of unreachable Internet Control Message Protocol (ICMP) messages to identify routers that are transmitting data intended for the victim. Once ingress routers have been identified, upstream peers can be contacted to continue traceback on their networks.

The key for backscatter traceback is the existence of a *sink hole*, which is a dedicated portion of the network that attracts traffic. The sink hole advertises large quantities of unused address space, typically referred to as *_BOGON* or *_DarkIP* (that is, unused or “unlit” portions of the overall IP address space that have not been allocated by the regional registries and therefore are not valid sources or destinations).

When traffic is dropped at the edge of the network, and if the source of this dropped traffic corresponds to advertised space from the sink hole, the ICMP unreachable generated from the router performing the drops are routed to the sink hole. The sink hole can be monitored to record these ICMP messages and identify the ingress routers, tracing the attack to the edge of an autonomous system.

Reaction or Containment

When an organization knows where an attack is coming from it can apply containment mechanisms such as ACLs. When attack traffic has been detected and classified, appropriate ACLs can be created and deployed to the necessary routers. Because this manual process can be time consuming and complex, many service providers use BGP to propagate drop information to all routers quickly and efficiently. This technique, *remote triggered drop*, sets the next hop of the victim’s IP address to the null interface. Traffic destined to the victim is dropped on ingress into the network.

Another option is to drop traffic from a particular source. This is similar to the drop described above but relies on the pre-existing deployment of unicast RPF (uRPF), which drops a packet if its source is “invalid”; invalid includes routes to null0. Using the same mechanism of the destination-based drop, a BGP update is sent, and this update sets the next hop for a SOURCE to null0. Now all traffic entering an interface with uRPF enabled drops traffic from that source.

Although scalable, the BGP-triggered drops limit the level of granularity available when reacting to attack: they drop all traffic to the black-holed destination or source, as described above. In many cases this is an effective reaction to a large attack.

Depending on the type of attack, service providers can use additional reaction methods, other than simply dropping traffic. They can rate-limit using committed access rate (CAR), which limits the rate of administratively identified traffic. For instance, a provider can, upon detecting a high rate of ICMP traffic, rate-limit that traffic to alleviate the effects of the attack. As with traffic filtering, a provider can use QoS Policy Propagation via BGP (QPPB) to remotely trigger CAR configurations.

Postmortem

The final phase of security best practices is post-mortem—reviewing what was most effective during an

attack and what could be improved. Postmortems should be conducted not only internally, but with other providers as well.

The Skitter Group—a group of the top 10 service providers within the larger NSP-SEC group—came out of an industry-wide postmortem after Cisco issued a major Product Security Incident Response Team advisory. When Cisco issued the advisory, it wasn’t communicated consistently to all of the top-tier service providers.

“As a result, some service providers implemented the changes while others didn’t, and this led to problems,” says Greene. “My contacts at these top-tier service providers later told me that they needed more direct, encrypted communications among themselves. The resulting Skitter Group has further improved coordination against future attacks.”

As organizations implement these best practices, the industry is gaining a new edge in the security battle. While attacks are destined to continue unabated, service providers stand a much greater chance of mitigating their most damaging effects through close collaboration and ongoing internal due diligence. ▲▲

◆ ◆ ◆

Joe Dallatore (jdallato@cisco.com) and Paul Quinn (paquinn@cisco.com) of Cisco’s Advanced Services engineering group provide security support to service providers and contributed to this article.

FURTHER READING

- **NSP-SEC SP security organization:**
puck.nether.net/mailman/listinfo/nsp-security
- **Cisco NetFlow product information:**
cisco.com/packet/161_8b1
- **Arbor Networks anomaly detection:**
arbornetworks.com/products_platform.php
- **Cisco Catalyst® Network Analysis Module:**
cisco.com/packet/161_8b2
- **SQL Slammer Mitigation white paper:**
cisco.com/packet/161_8b3
- **Internet security site:**
www.cymru.com
- **ISP Best Practices Tutorial:**
www.getitmm.com/bootcampflash/launch.html
- **NetWorm.org information center:**
www.networm.org
- **Renesis Internet monitoring:**
renesis.com/
- **ISP Resource Center:**
ispbook.com/


Troubleshooting MPLS VPN Networks

Getting the network humming is easier than you might think.

BY RAJIV ASATI

BEING RELATIVELY NEW, MULTI-protocol Label Switching virtual private networks (MPLS VPNs) might seem mysterious to troubleshoot. How do we tell whether a problem is occurring within the core or at the edge? Is it a problem in the MPLS provisioning or the Border Gateway Protocol (BGP)? Which database do we look in to make sure that labels and prefixes are correct?

In fact, however, troubleshooting an MPLS network relies on black magic no more than troubleshooting any other kind of network does. Understanding how MPLS labels and prefixes are attached, stored, and advertised is crucial to successful troubleshooting. The key protocol in an MPLS VPN network is Multi-Protocol BGP (MP-BGP), which provides the framework to exchange reachability information about many protocols such as



IN THE CONTROL PLANE: The key protocol in an MPLS VPN network, MP-BGP allocates labels for VPN prefixes in the advertising provider edge router (PE1) and accepts them in the receiving provider router (PE2). In this control plane example, the sending customer edge router (CE1) advertises an IP prefix to PE1, which converts the IP prefix into a VPN prefix. PE1 then allocates the label, installs it in both BGP and LFIB, and advertises the VPN prefix and associated label to PE2 via MP-BGP.

This article is based on a presentation at the Cisco Networkers 2003 User Conference. To learn more about Networkers, visit cisco.com/networkers.

IPv4, VPNv4, IPv6, multicast, and others. MP-BGP allocates labels for VPN prefixes in the advertising provider edge router (PE1) and accepts them in the receiving provider router (PE2).

Like any network, an MPLS VPN network logically comprises two planes: *control* (analogous to call setup) and *forwarding* (analogous to call transmission). The diagram on page 69 depicts the basics of the control plane. In this diagram, the sending customer edge router (CE1) advertises an IP prefix, “5.5.5.5/32,” to PE1 via a routing protocol such as BGP, Enhanced Interior Gateway Routing Protocol (EIGRP), Router Information Protocol (RIP), or the like.

After receiving the advertisement, PE1 converts the IP prefix 5.5.5.5/32 into a VPN prefix, “1:1:5.5.5.5/32.” Next, PE1 allocates the label (say, 20) to 1:1:5.5.5.5/32 and installs it in both BGP and the Label Forwarding Information Base (LFIB). PE1 then advertises the prefix to PE2 via MP-BGP. Please note that P routers, which use an Interior Gateway Protocol (IGP) such as Intermediate System-to-Intermediate System (IS-IS) or Open Shortest Path First (OSPF), do not have any VPN knowledge. P routers might optionally exchange IPv4 routes via MP-BGP.

PE2, after receiving the MP-BGP advertisement, checks whether the VPN prefix is acceptable by comparing the route target (RT) values. If acceptable, PE2 installs the prefix and label in the BGP and VPN Routing and Forwarding (VRF) FIB tables and advertises the prefix to CE2. In the meantime, FIB performs the “recursion resolution” to find a valid route and label to VPN prefix's next-hop (for example, PE1). If FIB finds so, it installs the next-hop label in the label stack that also contains the VPN label. This label stack is what FIB (at PE2) will use to forward VPN packets toward PE1.

Once we gain a sound understanding of MPLS VPN technology and how it operates, there are tips and logical series of steps that spotlight errors in configuring network components, advertising prefixes, and creating labels. For example, we can often tell if the problem is inside or outside of the core by going through the initial steps for troubleshooting the forwarding plane, outlined in the box on page 72. We check out whether packets are getting from the local PE1 through the core to the far edge PE2, and whether the interfaces between the CEs and PEs are dropping the packets. Further tests eliminate

the PEs themselves as culprits. If the packets never get from PE1 to PE2 and they're both set up correctly, the problem might be inside the core. Similarly, whether we should be checking VPN labels in the LFIB or the FIB is determined by whether we're troubleshooting for the VPN prefix at the advertising or receiving PE router. If it's the advertising PE router, check LFIB; if the receiving one, check FIB. Per our MP-BGP control plane diagram, 1:1:6.6.6/32 is advertised by PE2 (and received by PE1). Hence, we should check for 6.6.6.6/32 in LFIB at PE2, and in VRF FIB at PE1.

Fundamental Troubleshooting Tips

Before getting into heavy-duty troubleshooting, we should perform some simple configuration checks.

- First, validate the VRF configuration, because if the RTs aren't sent with VPN routes, nothing is going to work. Make sure that “export RT<X>” on the advertising PE router matches “import RT<X>” on the receiving PE router. To check, use this command on PE routers: `show ip vrf detail <vrf> | inc Export | import | RT`
- If export and/or import maps are configured in VRF, then validate the RT in the “match” and “set” clauses, using the commands `show ip vrf de <vrf> | inc route-map` and `show route-map <name>`.
- If BGP is not used as the PE-CE protocol, make sure the redistribution between BGP's VRF instance and the respective IGP VRF instance is correct.
- If route reflectors (RR) are present, make sure that the PEs are configured as route-reflector-clients in the `VPNv4 address family at RR`.
- MP-BGP neighbors must be configured to send “extended community” in the VPNv4 address family.

Troubleshooting the Control Plane

Next, we'll want to check for control plane problems in more depth. Common questions concern what's needed to configure an MPLS VPN, how to do it, which `show` commands to use, and the necessity of being a rocket scientist. (The answer to the last one is “no.”) If we find out that the customer's VPN traffic is not getting through, then we might have to ascertain, for example, that VPN routes are missing from one or more tables, that the routes are there but the labels aren't, that the labels are in BGP but not in LFIB.

Working through the following real-world problems, in sequence, should tell us why and what to do. The basic `show` commands for the control plane are in my *Networkers* presentation, “Troubleshooting MPLS VPN Networks,” at cisco.com/packet/161_8c1.

Problem #1: The prefix is in the VRF routing table on the PE1, but not in the MPLS table (LFIB). Use the following `show` commands to check whether the prefix is in the MPLS table, then the BGP VRF table, which allocates labels, and then to fix the error.

```
PE1#sh mpls forwarding vrf v1 | inc 200.1.61.4
```



RAJIV ASATI

RAJIV ASATI, senior networking consulting engineer in Cisco's High-Profile Solution Engineering team, presented “Troubleshooting MPLS VPN Networks” at *Networkers 2003*. He has also been a contributor at several national and international conferences, including APRICOT, MPLSCon, and MPLS World Congress. He can be reached at rajiva@cisco.com.

```
PE#1
PE1#sh ip route vrf v1 200.1.61.4
Routing entry for 200.1.61.4/30
  Known via "connected", distance 0, metric 0
  (connected, via interface)
  Routing Descriptor Blocks:
    * directly connected, via Serial2/0
      Route metric is 0, traffic share count is 1
PE1#PE1#sh ip bgp vpn vrf v1 200.1.61.4
%Network not in table
PE1#
```

This command sequence produces the information that the VPN prefix is not even in the BGP table. Because 200.1.61.4/30 is a connected route, it is likely not getting redistributed into BGP, which is why the label and prefix aren't entered in BGP. Configuring "redistribute connected" fixes the problem, as shown below:

```
PE1#sh ip bgp vpn vrf v1 label | inc 200.1.61.4
200.1.64.4/30 0.0.0.0 30/nolabel
PE1#sh mpls forwarding vrf v1 | inc 200.1.61.4
30 Aggregate 200.1.64.4/30(V) 0
PE1#
```

Problem #2: Rarely, the label is in the BGP table, but not in the LFIB. Use the command `clear ip route vrf<vrf><prefix>` to resynch the BGP and LFIB for the specific route. If the problem persists, clear the MP-BGP session and also inform the Cisco Technical Assistance Center (TAC).

Problem #3: PE1 advertises a VPNv4 prefix, but PE2 doesn't have it. We've already validated the BGP table and VRF configurations at PE1. The following `show` commands will indicate that PE2 doesn't have the VPNv4 prefix because there is no import RT configured on PE2.

```
PE2#sh ip bgp vpn vrf v1 200.1.61.4
%Network not in table
PE2#sh ip vrf det v1 | beg import
No import route-target communities
No import route-map
No export route-map
PE2#
```

The fix is to configure the correct import route-target and also check the import map configuration, if any.

Problem #4: PE2 still doesn't get the VPNv4 prefix from PE1. Since we've validated the configuration at PE2 previously, let's go back to PE1 and make sure it indeed is advertising the prefix to PE2 or other selected

edge routers. Using the same commands but substituting "export" for "import," check first that the export RT is correctly configured in PE1's VRF. Also, make sure that the RT is getting tagged to the VPNv4 prefix, and check the RT in the "set" and "match" clauses in the configured export map within VRF. PE1 might need this export map capability to be selective about which prefixes are advertised to which other PE routers—the essence of keeping VPNs private. But this selectivity also makes configuration more complex.

Problem #5: PE2 still doesn't get the VPNv4 prefix from PE1. Having checked Problem #4, we now know that PE1 is indeed advertising the prefix to PE2. Let's check the RR to make sure that the VPNv4 prefix is reflected appropriately—the RR should be configured with "neighbor<PE2>send-community extended" under "vpn4 address family" for all PE routers. The command to check route reflector #1 is:

```
RR1#sh run | community
neighbor 10.13.1.61 send-community extended
RR1#
```

As evident, PE2 (10.13.1.62) is missing above. PE1 and PE2 must both be configured as RR-clients under vpnv4 address family on RRs.

Bear in mind with all of these steps that even if we've configured everything correctly, there may still be lapses in packets getting from PE1 to PE2—perhaps a corrupted database. Packets might still not get through from PE1 to PE2, but there are further steps to follow, and others to troubleshoot problems such as label mismatches or packets not getting to the customer edge router. The point is that the process is logical.

Troubleshooting the Forwarding Plane

The forwarding plane deals with transit of packets across the MPLS network. Problems might involve how routers use header information, how packets are load shared, and other functions. (As with the control plane, the basic `show` commands are available in the Networkers presentation at cisco.com/packet/161_8c1.) Our first task, when traffic is not going through, is to isolate the problem into one of two categories: a problem inside the MPLS core or outside (for example, the customer-facing interfacing). Here's how:

- Issue "VRF ping" from an ingress PE to egress PE to verify that the PE-to-PE labeled switched paths (LSPs) are set up correctly and MPLS backbone is functioning properly: `ping vrf <vrf> <prefix>`. Pick up a connected VPN prefix as the <prefix> for the simplicity.
- Next, ping from one CE to another CE to verify that the PEs are forwarding the VPN traffic correctly. This again confirms that the LSP is established between PEs or not. If not, then enable "debug ip icmp" on

Find out more about MPLS VPNs from our peers and Cisco experts at the Cisco Networking Professionals Connection "Virtual Private Networks" forum: cisco.com/discuss/vpn.

MPLS VPN FORWARDING—BASIC TROUBLESHOOTING STEPS

So, you have received report of a VPN traffic outage:

1. First, verify “VRF ping” from PE1 to PE2.
2. If passed, then either CE->PE or PE->CE might be the problem=>not a MPLS core problem. Stop and check whether the packets are getting dropped by ingress LC on PE.
3. If failed, then MPLS core might be the problem. Proceed.
4. “Ping” ingress PE to egress PE to verify the IP reachability.
5. If failed, STOP and verify egress PE’s route hop by hop.
6. If passed, traceroute PE1->PE2 and PE2->PE1 to ensure the PE-to-PE LSP setup.
7. Check for the labels in each line of the traceroute output (watch out for the PHP).
8. If traceroute fails for some reason, stop and verify the label on every hop.
9. If good, the problem might be very specific to the HW on either PE or P routers. Find out if HW is dropping the packets.

CEs to further narrow down either the broken LSP or the erroneous PE-CE interface.

Specific Forwarding Plane Problems

The following steps also build on each other, as with the examples for troubleshooting the control plane.

Problem #1: VPN connectivity is broken between two CEs (CE1 = 5.5.5.5 and CE2 = 6.6.6.6). First, check the control plane information, making sure that the labeling is correct by using these commands:

```
PE1#sh ip cef vrf v1 6.6.6.6
PE2#sh mpls forward vrf v1 | inc 6.6.6.6
PE1#sh mpls forward vrf v1 | inc 5.5.5.5
PE2#sh ip cef vrf v1 5.5.5.5
```

Remember that we verify FIB information at the advertising PE router and in FIB at the receiving PE router for a given VPN prefix. The labels must match.

Having verified the control plan correctness using the commands above, enable “debug ip icmp” on both PEs, and then “VRF ping”—that is, `ping vrf v1 <remote PE-CE address>` from PE1 to PE2 by picking up a PE-CE prefix such as 200.1.62.5.

If ping fails and the “debug ip icmp” output doesn't show any “ICMP echo” received at PE2, then we know that the PE1->PE2 LSP is broken. If debug at PE2 does show the received “ICMP echo,” but debug at PE1 doesn't, then PE2->PE1 LSP is broken. Broken LSP typically occurs because of incorrect labels in LFIB.

Problem #2: VPN connectivity is broken between CEs. If the problem is not incorrect labels, it might be an IGP label mismatch on a P or PE. Work down through the steps shown in “MPLS VPN Forwarding,” page 71. First, try VRF pinging from PE1 to PE2. If the ping fails, go to Step 3 and verify the label information in the FIB for the VPN prefix. If that seems correct,

move to Step 4 to verify the PE->PE IP connectivity to the egress PE by pinging from an ingress PE. If the ping doesn't get through, Step 5 suggests next verifying the egress PE's route and label on the path to egress PE hop by hop by Telnet and also checking the LFIB entry for the egress PE prefix. This step should catch any label mismatch that might have occurred between the LFIB of two adjacent routers. This would typically be an LDP problem.

To fix the problem, validate the LIB bindings for the egress PE prefix with the command `clear ip route<prefix>`. If this doesn't fix, then flap the LDP neighbors to get them in sync.

As with the control plane sequence, additional steps narrow down the cause so it can be corrected.

This is by no means an exhaustive walk through of all the actions to be taken in troubleshooting a problem in an MPLS VPN network. However, it should demonstrate that the process is not mysterious or difficult, as many perceive—but that it is doable. There have been many enhancements in Cisco IOS® Software to further ease the troubleshooting: worth mentioning here are the “lsp ping” that can find the broken LSP and a command to flap an LDP neighbor. Lastly, remember that the Cisco TAC is always available to help. ▲▲

FURTHER READING

- **Networkers 2003—“Troubleshooting MPLS VPN Networks,”** Rajiv Asati:
cisco.com/packet/161_8c1
- **Networkers 2003—“Deploying MPLS VPNs,”** Zahir Aziz:
cisco.com/packet/161_8c2

The IP Revolution in Mobile Messaging

IP-enabled SS7/C7 applications yield significant gains in bandwidth and greatly reduced operations costs.

BY GREG WALKER

IN THE MOBILE WIRELESS ENVIRONMENT, THE phenomenon of mobile text messaging can easily be characterized as one of the most intriguing success stories of the past decade. Short message service (SMS) was never formally planned, designed, or marketed, yet millions of people all over the world adopted SMS—and changed the mobile world in the process. Today, many carriers worldwide recognize between 5 to 20 percent of their total revenue from person-to-person SMS text messages. This contribution, in turn, has initiated a drive to introduce additional, more advanced revenue-generating services. The common denominator for these advanced services is IP.


While the de facto bearer for SMS has traditionally been Signaling System 7 (SS7)/C7, the ongoing need to grow average revenue per user (ARPU) and average revenue per message (ARPM), while maintaining an optimal quality of service (QoS) and significantly reducing infrastructure costs, has pointed to IP as the logical next step. Taking their cues from customers and the marketplace, Cisco and LogicaCMG (logiacmg.com) recognized this need early on and have jointly developed an IP-based next-generation messaging architecture.

Meanwhile, with Cisco's participation, the Internet Engineering Task Force (IETF) has defined a set of standards for transporting SS7 traffic over IP networks (SS7overIP). As part of this effort, specialized protocols were developed to IP-enable traditional SS7 end-node applications, generically referred to as signaling control points (SCPs). Typical SS7 end-node applications include network elements such as home location registers (HLRs), and billing applications and service delivery platforms such as short message service centers (SMSCs). By IP-enabling these traditionally circuit-based nodes, tremendous gains in bandwidth and capacity are achieved along with greatly reduced infrastructure capital and operational costs.

Benefits Realized

Replacing rigidly defined, costly low-speed SS7 link limitations with cost-effective IP bandwidth—SS7overIP—will yield gains in end-node processing and transaction capacity. Consider an HLR, which is essentially a large database and computing platform. With limited SS7 connectivity options, using traditional SS7 design, one could connect up to 16 low-speed SS7 links at

For greater detail on SS7overIP, the SIGTRAN protocols, and Cisco ITP, check out the white paper, *Next-Generation SS7 Networks with the Cisco IP Transfer Point*, at cisco.com/packet/161_8d2.



HERCULEAN BENEFITS: The IP-enabled LogicaCMG SMSC connects to a mated pair of Cisco ITPs using the SIGTRAN SUA protocol. These ITPs, which can also perform MLR, transport SS7overIP using the M2PA protocol to distributed ITPs, which function in the network as STPs and SMS FDA nodes. Connectivity to the legacy MSCs is accomplished via traditional TDM low-speed links or high-speed links.



GREG WALKER

GREG WALKER, a product manager in Cisco's Mobile Wireless Business Unit, is focused on Cisco ITP. Working with mobile wireless and wireline customers worldwide, he is responsible for disseminating ITP product knowledge and defining the product roadmap. He can be reached at gwalker@cisco.com.

56/64K, for a total bandwidth to the end node of 896/1024K. With some creativity in the application of SS7 point codes, one might even deploy up to 32 low-speed SS7 links to the end node (for a total bandwidth of 1792/2048K). However, this multiple point-code approach quickly becomes fraught with network complexity. What's more, with the SS7-defined requirement for complete redundancy, half of this limited bandwidth must be set aside and reserved for failover, thereby effectively cutting the available bandwidth in half. In the case of large, centrally located SS7 end nodes, which must process high volumes of transactions, this limited bandwidth constrains the node and results in an inability to effectively use the inherent processing/transaction capacity of modern computing platforms.

Compare this to the bandwidth available via Fast Ethernet and Gigabit Ethernet, and the possibilities are enticing. An IP-enabled node supporting 10, 100, or even 1000 MB of bandwidth can fully realize its inherent processing/transaction capacity, often increasing existing transaction capacity by a factor of 3 to 5-plus times. Investment in existing SS7 end-node applications can thus be preserved and extended simply by implementing a client-stack IETF Signaling Transport (SIGTRAN) protocol and a signaling gateway function.

Equally important, this IP bandwidth is achieved at a fraction of the cost of dedicated leased lines, both from a capital and operational cost perspective. For instance, consider a traditional end node with SS7 connectivity. This node would require both expensive SS7 hardware and expensive, typically third-party, SS7 protocol stacks. But using IP and widely available SIGTRAN client stacks, these nodes can be IP-enabled for a fraction of the cost. These reduced product costs, in turn, allow for a more competitive product offering in the marketplace.

Additionally, from an operating cost perspective, expensive leased-line time-division multiplexing (TDM) connectivity is removed and replaced with cost-effective IP bandwidth. When the benefits are analyzed, this reduction in TDM leased lines *alone* provides a significant, compelling cost benefit.

SIGTRAN Standards

The IETF SIGTRAN standards were developed to preserve the reliability and redundancy characteristics demanded by an SS7 network, and benefit from the sim-

licity and cost effectiveness of IP. Each standard has been designed for specific functionality. Stream Control Transmission Protocol (SCTP), for instance, serves as the foundation reliable transport layer for SIGTRAN "adaptation-layer" protocols, such as MTP2-User Peer-to-Peer Adaptation Layer (M2PA), Message Transfer Part Layer 3 User Adaptation (M3UA), and Signaling Connection Control Part (SCCP) User Adaptation (SUA).

Using SCTP's reliable transport, the M3UA and SUA protocols provide an internetworking function between IP and SS7 networks while preserving the redundancy and reliability features associated with SS7 (see cisco.com/packet/161_8d2 for a graphic representation of the SS7 protocol stack). In a client-server architecture, M3UA provides an internetworking function at the MTP3 layer, and SUA provides this function at the SCCP layer. Using SS7overIP internetworking/adaptation layer protocols, higher-layer SS7 protocols are transported unaltered as payload across the IP network. IP-enabled end nodes and applications can thus seamlessly interwork with existing SS7 network elements. Examples of higher-layer SS7 protocols include ISDN User Part (ISUP), Transaction Capabilities Application Part (TCAP), Mobile Application Part (MAP), and IS-41.

The Cisco IP Transfer Point

The Cisco IP Transfer Point (ITP) is a truly flexible platform for SS7 transport. For legacy SS7 routing and transport, Cisco ITP incorporates the functionality of a TDM-based signaling transfer point (STP). Leveraging the SIGTRAN protocols, the ITP also functions as a next-generation STP, supporting SS7overIP and legacy SS7overTDM simultaneously. Using the SIGTRAN M3UA and SUA protocols, Cisco ITP performs as a classic SS7/IP signaling gateway with advanced QoS and load-balancing features. Additionally, GSM MAP and IS-41 upper-layer services have been implemented to enable wireless LAN subscriber identity mobile (WLAN SIM) authentication, SMS Multilayer Routing (MLR), spam filtering, and distributed SMS first delivery attempt (FDA) functionality.

With a flexible approach to implementing SS7overIP, mobile operators can migrate to next-generation signaling networks at a manner and pace consistent with their business requirements and goals. For instance, some operators might choose to begin migrating to SS7overIP by IP-enabling legacy SS7 nodes in their network. Another less invasive step for many operators involves offloading SMS traffic to IP. Having gained sufficient comfort and experience with the technology,

Continued on page 89

FURTHER READING

- **Cisco ITP home page:**
cisco.com/packet/161_8d1

Ringin Up Profits

Cogeco Cable sets sights on profitable business market.

BY DAVID BARRY

COGECO CABLE (COGECO.COM), CANADA'S fourth largest cable operator, is seeking to enlarge its market share by targeting the lucrative business market. Until now, Cogeco has primarily served residential customers with analog and digital video and high-speed Internet services through its two-way broadband network. But in a move that opens the door into Canada's business market, Cogeco recently added the Cisco ONS 15454 Multiservice Provisioning Platform (MSPP), equipped with the ML-Series packet-switching line cards, to its Cisco end-to-end network.

The ML-Series adds two new key capabilities to the company's existing SONET/SDH infrastructure that enable the efficient delivery of Ethernet: *packet multiplexing* and quality of service (QoS) with committed information rate (CIR) and peak information rate (PIR). Also available on the same platform is Resilient Packet Ring (RPR) technology—providing cable operators and service providers with the ability to support multiple applications over various topologies (for example, point-to-point, hub and spoke, and RPR) using a single multiservice optical platform.

“The most impressive capabilities of the new ML-Series card are the efficiency and device consolidation they provide.”

—CHRIS MACFARLANE, VICE PRESIDENT OF IP TRANSPORT AND ENGINEERING, COGECO CABLE

Packet multiplexing allows the intelligent and flexible use of fixed amounts of bandwidth and enables much greater efficiency when delivering Ethernet services across the SONET/SDH backbone. Unlike legacy techniques that reserve fixed bandwidth—regardless of whether it is actually used—packet multiplexing allows the bandwidth to be dynamically reallocated, thus delivering more efficiency from the same amount of bandwidth.

Packet multiplexing allows Cogeco to oversubscribe the allocated network bandwidth. But to accomplish this, Cogeco must leverage QoS solutions that

guarantee that business customers get the bandwidth they require. The QoS capabilities in the new ML-Series hardware and software release 4.1 allow Cogeco to offer customized services that will attract and retain more customers and increase the company's profitability. Cogeco will be able to provide its business customers with Ethernet CIR, which ensures a guaranteed transmission rate regardless of congestion. PIR, a peak rate, will be available if the network is not congested. These QoS capabilities enable Cogeco to support more customers by leveraging packet multiplexing to add network efficiency and scalability.

QoS and packet multiplexing will allow Cogeco to provide customers the service-level agreements (SLAs) necessary for voice over IP (VoIP) and multimedia applications. Ethernet circuit provisioning and monitoring of SLAs is accomplished through the GUI-based interface offered by Cisco Transport Manager—the integrated optical element management system for the entire Cisco COMET (Complete Optical Multiservice Edge and Transport) portfolio of optical networking products.

“The most impressive capabilities of the new ML-Series card are the efficiency and device consolidation they provide,” says Chris MacFarlane, vice president of IP Transport and Engineering, Cogeco Cable. He continues, “We'll be able to use a resilient packet ring topology to offer Ethernet services and reduce the number of STS-Nc circuits we currently use. Instead of only transporting one 100-Mbit/s Ethernet service over an STS-Nc, like we do today, we'll be able to offer *multiple* Ethernet services, to several different customers, all over that same STS-Nc. That really adds efficiency to our network without an expensive buildout.”

ML-Series interface enhancement enables Layer 2 and Layer 3 switching from the Cisco ONS 15454 MSPP, which already supports TDM and optical services. It also interoperates with Cisco routers and switches to provide consistent end-to-end SLAs through Cisco IOS® Software. This integration of multiple functions—IP, Ethernet, and RPR—into a single platform helps Cogeco reduce the type and number of network elements required for Ethernet and IP transport.

“The ML-Series card will simplify our network by integrating IP, Ethernet, and RPR features for multiple

Continued on page 88

Small AND Midsized BUSINESSES

Paradise Reconnected

Cutting Costs and Expanding Communications Horizons at the Eden Project BY GRANT ELLIS

AN ABANDONED CORNWALL CLAY mine seems an unlikely site for a major tourist destination. But the Eden Project (edenproject.com), which opened in 2001, appeals to a basic human need: to understand, enjoy, and preserve the natural world. It also appeals to our love of the exotic. What really grows in a rain forest? What does a banana blossom look like?

That appeal is strong enough that nearly two million visitors a year come to Bodelva, near St. Austell in Britain's Cornish peninsula, to look at carefully tended plantings that accurately reflect local vegetation at locations around the earth. There, visitors can admire 5000 plant species. They can experience exotic habitats in the largest conservatories on the planet—glazed freeform bubbles called *biomes*. A lush, steamy rain forest fills the 240-foot Humid Tropics Biome, and the Temperate Biome is a pleasant stroll through the Mediterranean climate zones of the world.

While the enormous public interest in the project has been exciting, says Howard Jones, head of organizational development for the Eden Project, the resulting communications traffic eventually strained telephone and IT infrastructures to the breaking point. Cisco partner BT has installed Cisco converged network technology that not only solves these problems but also reduces costs, improves productivity, and enables the project to extend e-learning opportunities to Web audiences worldwide.

First Priority: IP Telephony

Some 600 people work in a startling variety of environments within the 35-acre landscape of the Eden



Project. They use 288 phones and 250 PCs, and by 2001 both the analog phone network and the IT network were failing often from the sheer burden of activity.

The Eden Project had other communications concerns. Although it operates as a commercial business, it is also a visionary facility dedicated to research and education about human interaction with the planet. There was a strong desire to use Web-enabled technology and a variety of digital media for e-learning and for data-sharing with other scientific organizations.

When Jones contacted BT, his first priority was to replace the ailing phone system. BT invited Jones to see a Cisco IP communications solution in action at nearby Cornwall College, which was running a converged voice and data network on Cisco AVVID (Architecture for Voice, Video and Integrated Data). It was obvious to Jones that Cisco AVVID would not

only reduce costs and improve productivity at the Eden Project, but would also create an infrastructure that could scale to support its far-reaching goals. “I thought it was tremendous,” says Jones. “This was going to be the answer, no question.”

The Eden Project management team reflected this openness to new solutions when BT presented its proposal. “They’re so different from any client I’ve ever come across,” says BT Business Manager Simon Tilden. “They’re very willing to embrace new technology.”

Wanted: Environmentally Aware Vendors

Engineers from BT and Cisco toured the Cornwall site to gather information for a proposal. “It was fascinating,” says Cisco Engineer Steve Hammond. “We began to see the challenges they faced. There were all sorts of communications environments, including low-population buildings such as snack bars and parking huts, scattered around an immense property. They also have a highly mobile staff moving around the site. They needed great flexibility and wireless IP.”

Before the proposal could be written, the Cisco UK team had to answer some searching questions. Was Cisco committed to protection of the environment? Did the commitment go beyond mere corporate window-dressing? What were its policies on recycling products? A full day of investigation at Cisco’s London office convinced the Eden Project’s management of the depth of Cisco’s commitment. It became clear that BT, Cisco, and the Eden Project shared concern for the reduction of waste and the sustainability of the environment.

The project commissioned BT to design and install a converged voice and data system based on Cisco AVVID, which underlies Cisco IP communications solutions. BT expanded the existing fiber network, installing two Cisco Catalyst® 6509 switches as the network core and eight Catalyst 2950 switches at the edge. Nineteen Catalyst 3524-PWR XL switches provide integrated voice, video, and data service to the project’s 250 desktops and to its 288 highly customizable Cisco IP 7940 phones. The 7940’s Extensible Markup Language (XML) capabilities provide phone users with access to a variety of information sources, including Web-based content. Two deployments of Cisco CallManager Version 3.3 support functions such as unified messaging, multimedia conferencing, and interactive multimedia response.

BT completed the network and IP communications installation early in 2003 as phase one of the Project’s communications upgrade program. According to Jones, “It was far more cost-effective and functionality-rich than any other phone solution available.” Customer service improved with more efficient call handling, and internal communication became more positive and immediate.” Jones estimates that choosing the Cisco IP communications solution has saved the project US\$170,000 in terms of initial capital outlay and US\$85,000 a year in operational costs. “In terms of cost, it was a complete no-brainer,” says Jones.

Interchangeability at the Desktop

With the new backbone and IP communications in place, the second phase of the upgrade was to create a converged voice and data LAN that serves the project’s 250 desktop PCs as well its phones. By October 2003, BT engineers had replaced the entire office IT system with Cisco AVVID technology running on the new backbone. Cisco Unity™ 4.0 unified communications software, deployed on every desktop, integrates with desktop applications such as Microsoft Outlook and Lotus Notes to simplify communication for Eden Project staff members.

“Cisco Unity is tremendous,” says Jones. “It enables me to deal very quickly with a variety of media using whichever piece of hardware I choose. For example, when I’m logged on it will pick up voice messages from my phone and display them or let me hear them, and I can respond by e-mail, which is really handy when I’ve got no time to talk and the message needs a quick answer. Or vice versa—it will let me listen to my e-mail over the phone so I can respond to it with a phone call. This interchangeability is tremendous; it enables me to do a variety of functions with a single tool—in my case, a laptop.”

Jones estimates that the project’s annual saving in IT costs from using the Cisco IP communications solution on the desktop is approximately US\$340,000 a year. That doesn’t include the efficiencies generated by Cisco Unity in terms of saved messaging time and faster access to data.

Despite the obvious efficiencies created by Cisco IP communications, Eden Project staff members were given plenty of time to absorb the technology before it became fully operational. “BT and the customer didn’t

The new Cisco IP communications system will enable the project to achieve one of its most important goals: the creation of a rich learning experience that can be accessed for classroom sessions in schools across Britain.

do a ‘Big-Bang’ change,” says Cisco Account Manager Nathanael Whitbread. “It was very much a phased business-as-usual change that introduced the technology in a gradual way.”

Another example of how innovative uses of Cisco AVVID are improving communication and business processes at the Eden Project is Mayfly—the name given to the front Web page of the Eden Project intranet (because it lives for just 24 hours).

Mayfly describes the day’s activity programs, including any last-minute changes in assignments. It can be accessed from all desktops, and XML applications embedded in the network can be used to transmit it instantly to the personal digital assistants (PDAs) and mobile phones used by staff members. It literally puts all staff members on the same page every day.

Concluding the Dream

The US\$136 million next construction phase of the project is now in the design stage. It includes an education center that will likely become the most exciting and important building in the project. Jones plans to apply the most advanced technology available to capture visitors’ imaginations about how humans interact with the natural world. Cisco technology will be part of the network buildout and a key player in the construction process.

“We’re aware that Cisco is very good at using AVVID for building infrastructure technology for everything you can imagine,” says Jones, “from building control systems to security and even the design and build itself, Web-enabling the process from architect to building site. Our building program will be hugely enhanced by Cisco wireless IP. The builders on site will use laptops or PDAs linking with the architects live to make minor changes or adjustments to specification.”

The vision of site-wide wireless communication will soon be realized for Eden Project staffers who roam the grounds to tend, manage, supervise, and inspect. BT is deploying Cisco Aironet® wireless technology and XML applications to provide instant communication through multiple platforms, including PDAs and cell phones. This is more difficult than it sounds. The old clay mine left a broad pit 260 feet deep, and personnel and facilities are spread over this irregular depression out to the edges of the site.

The converged Cisco solution will help simplify the problems by moving voice and data in a single digital stream.

“We need the robustness of passing things through single pipes on Cisco AVVID,” says Jones.

Jones plans to use Cisco AVVID and the wireless facility to create new teaching relationships with vis-

itors. One possibility is giving visitors PDAs that will enable them to learn through interactive media as they move through the exhibits.

The Final Network Extension: E-Learning

It is important to remember that the Eden Project was conceived and organized primarily by professionals with an academic background, outspoken concern for people and the planet, an interest in research, and a desire to educate. The new Cisco IP communications system will enable the project to achieve one of its most important goals: the creation of a rich learning experience that can be accessed for classroom sessions in schools across Britain.

“When we were talking about the architecture we wanted on this backbone, we knew we wanted something that would carry all forms of media, both internally and for our potential audiences,” says Jones. “First and foremost, we’re an educational platform and a forum for learning about the relationship of people to the planet and the natural world, using plants as an instructional tool.”

One example of the way the Eden Project intends to use the converged communications platform for classroom teaching is the addition of Web cams. Recorded virtual tours of the two biomes and outdoor plantings are already posted on the Eden Project Website. Those tours will ultimately be live. In fact, one live presentation has already appeared on the Website: a bird’s-eye video of the interior of a biome, taken from a remote-controlled dirigible.

Eden Project scientists have still another application in mind for their Web-enabled network. They will use it for real-time collaboration and knowledge sharing with corporations, research institutions, and the academic world.

Jones sums up their satisfaction: “The whole thing is working remarkably well. In the truest sense, this is converging technologies—putting it all together in a very simple way, making simplicity out of complexity.” ▲▲

FURTHER READING

- Cisco IP communications solutions:
cisco.com/packet/161_9a2
- Cisco Networking Professionals Connection
“Voice and Video” Discussion:
cisco.com/packet/161_9a3
- Cisco IP Communications Application Central:
cisco.com/packet/161_9a4

Meet your peers online at the *SMB Networking Connection* and talk about small and midsized business networking topics. Join the Discussion Boards to ask questions and share ideas, then visit the Learning Center to find out about solutions to help your business problems:
cisco.com/packet/161_9a1

A Glimpse into the World
of a CCIE Lab Proctor

The PROCTOR FACTOR



by RHONDA HELDMAN RAIDER

THE CCIE® IS THE HIGHEST CERTIFICATION THAT Cisco grants, opening new career opportunities for recipients and inspiring confidence in prospective

customers. The first stage of certification is to pass a two-hour multiple-choice exam. Those who pass qualify to take an eight-hour hands-on lab exam. It's tough: since the program was introduced more than ten years ago, less than 3 percent of Cisco certified professionals have earned CCIE certification.

In September 2003, Cisco added the CCIE Voice track, which joined the three other tracks: Routing and Switching, Security, and Service Provider. The Voice track certifies and recognizes experts in configuring and maintaining IP telephony solutions. The demand for such knowledge is significant: more than 800 professionals sat for the Voice written exam in the first month, a record for any track in the CCIE program to date.

On the day of the hands-on portion of the exam, the candidate's world inevitably narrows to the lab where his or her own knowledge will be put to task, and at the front of the room, sits that mysterious lab proctor. *Packet*® recently caught up with Voice Lab Proctor Ben Ng to find out what his job entails. Ng is intimately familiar with intricacies of the CCIE exam. He earned his CCIE in Routing and Switching in October 2000, and in Voice in August of this year to prepare him for the role of proctor.

What does a lab proctor do?

My official job title is customer support engineer. It fits because my primary role is to provide support to candidates taking the lab exam. This support takes a variety of forms. One is to clarify the lab exam questions for candidates—so long as their query isn't really an end-run to find out the answer! Also, I fix equipment should a hardware or software failure occur. Basically, I view my job as doing whatever is necessary to make the exam less stressful for candidates, short of providing beer and cigarettes.

How many people have taken the Voice lab exam?

Currently I can accommodate up to three candidates a day. That works out to about 15 a week or 60 a month. We've been very busy since the lab opened. Most candidates are from the United States, but I've tested candidates from virtually every part of the world.

How many of them passed the exam?

So far [mid-January 2004], 31 people have earned their CCIE Voice certification, which is less than 10 percent of those who passed the qualifying written exam. Some pass on



their second try, and already I am testing some who are on their third try. It's a very rigorous exam, and that is why the CCIE certification is held in such high regard.

What's so special about the CCIE?

CCIE certification recognizes network engineers for in-depth knowledge of networking protocols—specifically, the ability to configure networking equipment so that it works together, and to troubleshoot and recover if anything goes wrong. The objective isn't training; it's identifying experts who are capable of understanding the subtleties and snares of end-to-end networking. We don't have formal prerequisites to sit for the exam, but we strongly encourage candidates to have at least three to five years of hands-on job experience before they attempt the lab exam.

What type of equipment is in the lab?

We select the equipment and compose questions based on what candidates are likely to encounter in real-world scenarios. A full list of equipment that might be used is posted on the CCIE Website [cisco.com/go/ccie]. We provide basic network connectivity, and candidates are asked to build an IP tele-

phony system, including establishing parameters for quality of service, virtual LANs, gateways, gatekeepers, and so on. We present problems that people actually run into when they're building voice networks. My prior life as a technical lead in the Cisco TAC [Technical Assistance Center] prepared me well for this job, because I know the areas where people are most likely to make mistakes.

Do you change the exams?

Yes, frequently.

Describe a typical work day.

The other proctors and I arrive at the lab between 7:00 and 7:30 a.m. I make sure the equipment racks are ready for use and the equipment is working, and set out the examination books and scratch paper. When everything is ready, I greet the candidates in the lobby and bring them to the lab to give them their seat assignments. During a ten- or fifteen-minute briefing, I go over the exam do's and don'ts, when they'll break for lunch, and other matters like that. During the course of the exam, I answer questions from candidates who are unsure about something on the exam, basically to clarify

the meaning of the questions for them. Sometimes they approach me because they think the equipment isn't working correctly. In those cases, I'll check to make sure the hardware and software are free of problems.

When things are quiet, I grade exams from the previous day by verifying that the configuration is correct and that it actually works. Grading is based on results only, so if the configuration looks right but doesn't work, the candidate might not earn the point. First I troubleshoot and try to find out why the configuration didn't work. I only deduct points after I can confirm that the problem was the configuration and not a hardware or software failure. I want to make sure candidates get credit if they're entitled to it.

After the exam, I like to take some time to chat with the candidates and find out how they're feeling. I realize they've spent an intense day. When proctors aren't monitoring or grading exams, we generally spend our time learning, working with new equipment, and creating future lab content. We work with the rest of the CCIE team to make the lab a better experience for candidates.

A Candidate's Perspective

T.J. Schuler, network engineer with Cisco reseller Flair DataSystems, Inc., knew he wanted to take the CCIE Voice exam as soon as it was announced. "I'm a glutton for punishment," says Schuler, who at the ripe age of 24 already has eight years of networking experience and holds two other CCIE certifications, in Routing and Switching and Security. "Even though I've deployed a lot of IP telephone networks, I knew I

would learn while preparing for the exam, and the certification gives our customers confidence."

Schuler passed the exam on his second attempt, choosing to retake it as soon as possible so that the technology wouldn't have changed. "Ben [Ng] was great. I really get the sense that Ben and the other proctors I've met understand how difficult it

is to fly to a different location and take a difficult test."

Schuler wasn't shy to ask questions during his exam. "Every word on the exam is there for a reason, and I took the opportunity to make sure I wasn't reading too much into it. The proctor is there to guide you in the right direction. Ben was extremely helpful in clarifying whether or not I knew what the question was asking."

What's an example of the type of question candidates shouldn't ask?

Candidates are entitled and encouraged to ask any questions if they feel the wording of the exam content is not clear. However, generally I will decline to answer a question if it is formulated for me to choose or confirm a specific solution. We can't give candidates a direct reply if they're trying to narrow down the answer. But we don't mind clarifying the questions if they ask. The worst that could happen is I tell them I can't answer.

What are some of the stranger incidents you've encountered?

I take candidates out for a 30-minute lunch break. One candidate cussed me out during lunch, which was quite a surprise. Another guy attempted to simulate a successful fax transmission by physically inserting the original upside-down into the receiving fax machine, as if it had been sent. Then there was a person who tried to use a fake ID to take the exam for someone else. And once when another proctor called "time's up," a candidate fainted, and within a few minutes stood back up as if nothing had happened.

Do people have misconceptions of lab proctors?

Yes! The main one I would like to dispel is that we're unapproachable. It's true that being a proctor implies a certain degree of authority. Whatever we say in the lab goes. However, we very much welcome questions as long as they're appropriate. And we're very committed to making the

candidates comfortable. For example, recently a candidate was shivering from the cold, so I offered him a jacket from our rather large collection of jackets and sweaters left behind by preoccupied predecessors. If someone starts to feel ill during the exam, we try to offer a remedy or something to help them relax, such as hot tea. Candidates are welcome to make brief visits to the lab days before the exam to help ease their anxiety, so they can better focus on studying.

What's the hardest part of your job?

By far the biggest challenge is resisting the urge to solve problems for candidates who are struggling—a mindset that's especially strong for me because my previous job was with Cisco TAC. I also have to be very mindful of answering people's questions about the exam in a way that simply clarifies the question instead of answering it. This sometimes requires careful thought because English is not my first language.

It's also emotionally taxing to have to fail people. I often remind myself that it's a challenging examination that recognizes only the top few percent of engineers, so some excellent engineers won't pass. I also like to remember that regardless of whether a candidate passes or fails, the exam is a valuable learning experience. Candidates must solve a real-life problem under time pressure, so they're better prepared to cope when they encounter the same experience on the job. I'm convinced that going through any of the CCIE labs helps make people better engineers. So what's a little suffering!

And your favorite part of the job?

I enjoy meeting people from diverse backgrounds. Everyone who passes through my lab is in the networking industry, but they all have unique experiences and challenges. I especially enjoy our lunchtime conversations because I get fresh perspectives on what's going on in real production networks. Our customers implement and maintain some really innovative and complex networks.

Definitely the best part of the job is when candidates make positive comments about the lab. If someone says that studying and preparing for the lab and then solving problems in a stressful environment has made them a better engineer, I know my hard work is making a difference.

And every time a candidate passes, I get a boost. Some people think that proctors are mean-spirited nerds who enjoy failing people! They couldn't be more wrong. Whenever a Voice candidate passes, I immediately share the good news with the other proctors and send out an e-mail of congratulations. Really, we're just engineers who love networking technologies and are excited about how such technologies are changing our everyday lives. ▲▲

FURTHER READING

- **CCIE home page:**
cisco.com/go/ccie
- **Voice track home page:**
cisco.com/packet/161_10a1

Technically Speaking

Voice Call Transfer, Forwarding in IOS

BY JASON DACTLER

IF YOU WORK WITH VOICE OVER IP (VoIP) networks, ensuring compatibility between equipment is a constant challenge. Even basic call connections can be challenging because of the variety of standards-based signaling protocols—H.323, Session Initiation Protocol (SIP), Media Gateway Control Protocol (MGCP), H.248, and so on—and the varying vendor implementations. With supplementary services, interoperability is even more of an issue.

The ITU currently defines 12 recommendations (H.450.1–H.450.12) for supporting various supplementary services in an H.323 network. This article focuses on H.450.2 (call transfer) and H.450.3 (call diversion).

With H.450.2 call transfer, three parties exist: User A, the party initiating the transfer; User B, the party being transferred; and User C, the party receiving the transferred call. User A has an established call with User B and initiates a transfer to User C. This transfer initiation is either consultative, with User A privately conversing with User C before completing the transfer, or blind, where User A transfers User B without conversing with User C.

In H.450.3 call diversion or forwarding, three parties also exist: A, the calling endpoint; B, the diverting endpoint; and C, the diverted-to endpoint. Endpoint A calls Endpoint B, Endpoint B does not answer, and the call is sent to Endpoint C, such as a voice-mail system or a cell phone number. H.450.3 passes on the reason for the call diversion at Endpoint B, which can be ring-no-answer, busy, or unavailable. This allows Endpoint C to behave differently depending on the reason, such as playing different prompts for busy versus ring-no-answer.

In both cases, the intermediate party is removed from the call, freeing resources and bandwidth for other calls. But these recom-

mendations typically require that all three parties support H.450. While many VoIP devices support H.323 signaling, only a subset offers H.450 support, leading to interoperability problems when trying to implement supplementary services.

Until recently, support for H.450.2 and H.450.3 on a Cisco IOS® voice gateway was available only through a Tool Command Language (TCL) application and was limited to receiving requests only. The only exception was Cisco CallManager Express (formerly IOS Telephony Service). Beginning with version 2.1, Cisco CallManager Express also initiates call transfers and forwards using the H.450 recommendations.

To configure an IOS voice gateway for H.450 support, the H.450 TCL application and supporting files were loaded onto the gateway and applied to all dial-peers that would be involved in the call transfers or forwards. If other TCL-enabled services, such as prepaid calling card support, were required on the same dial-peer, a custom TCL application was required to combine the functionalities, because dial-peers support only one application. Potentially complicated and time-consuming, this led to longer deployment times for such services.

Now, with Cisco IOS Software Release 12.3(4)T, the default session application in IOS includes H.450.2 and H.450.3 support. The same limitations on initiating H.450 requests still apply, but the configuration to accept H.450 requests is eliminated.

While Release 12.3(4)T eases implementation of H.450-based supplementary services, interoperability issues with devices that do not support H.450 still exist. The most common example is Cisco CallManager, which uses an H.323 mechanism, *Empty Capability Set (ECS)*, to initiate call transfers and

forwards. So, how does an H.450-compliant platform like Cisco CallManager Express interoperate with a non-H.450-compliant platform like Cisco CallManager?

In Release 12.3(4)T, Cisco CallManager Express provides for a VoIP-VoIP hairpin for call transfers or forwards to non-H.450 devices using H.450 TCL version 2.0.0.8 or later. Cisco CallManager Express first attempts the H.450 transfer or forward, and if that fails, falls back to hairpinning the call. This procedure provides a good balance between end-user experience (call transfer or forward is completed), network resources (only non-H.450 calls are hairpinned), and service provider configuration tasks (no special dial-plan provisioning).

What's Next

Cisco IOS Software Release 12.3(6)T integrates this hairpin function into the default session application. H.450.12 capabilities exchange allows Cisco CallManager Express to auto-detect the ability to support H.450 services per call and attempt a hairpin transfer only if H.450 is unavailable. Finally, Cisco IOS Software Release 12.3(6)T allows Cisco 2600, 3600, or 3700 series routers to act as H.450 proxies for non-H.450 devices such as Cisco CallManager, the Cisco BTS 10200 Softswitch, and the Cisco PGW 2200 Softswitch. This capability eliminates the double-bandwidth limitation of these non-H.450 devices. ▲▲



JASON DACTLER, CCIE® No. 7133, is a technical marketing engineer in the Service Provider Solution Engineering Group at Cisco. He joined Cisco in 1998 and has worked on the design of both VoIP and remote-access networks. He currently specializes in Metro Ethernet solution design. Dachtler be reached at dachtler@cisco.com.

New Product Dispatches

Core Routing

Cisco 12000 Series Routers: New Models

Four new routers in the Cisco 12000 Series give service providers greater capacity and functionality in core and edge networks. The new Cisco 12816 and 12810 routers provide 40 Gbit/s per slot in a 16- or 10-slot configuration. With the industry's first field-upgradable 40-Gbit/s routing solution, service providers can double the capacity of the largest IP/Multiprotocol Label Switching (IP/MPLS) networks without a forklift upgrade. The new Cisco 12010 and 12006 routers provide 2.5 Gbit/s per slot and can be upgraded to 10 Gbit/s per slot with a simple software license key. The Cisco 12010 and 12006 routers are ideal for deploying high-performance edge services when coupled with Cisco IP Services Engine (ISE) line card technology.

cisco.com/go/12000

Cisco 12000 Series Routers: New Performance Route Processor and Line Cards

The new Cisco 12000 Series Performance Route Processor-2 (PRP-2) offers up to 100 percent performance improvement over existing PRP-1 on most system tasks, up to 4 GB of memory, a Gigabit Ethernet management port, and a hard-drive option. PRP-2 delivers the scalability needed for service and network convergence. Also new are the Cisco 12000 Series Two-Port OC-192c/STM-64c POS Line Card and the Cisco 12000 Series Eight-Port OC-48c/STM-16c POS Line Card, which cost effectively scale IP/MPLS packet infrastructures to 50 million pps per line card while providing a rich IP/MPLS feature set. The Cisco 12000 Series Four-Port OC-3c/STM-1c ATM ISE Line Card combines service-enabling edge features

with advanced ATM traffic management and sophisticated IP/MPLS-to-ATM quality-of-service (QoS) functions, without compromising line-rate performance. The Cisco 12000 Series One-Port Channelized OC-12c/STM-4c (DS1/E1) POS/ISE Line Card channelizes up to 840 channel groups of DS1 and DS0 circuits and provides an extensive set of service-enabling edge features at line-rate performance.

cisco.com/go/12000

Switching

Cisco Catalyst 4500 Series Switches: Bi-Directional Fast Ethernet Line Card

The 48-port 100BASE-BX10-D Bi-directional Fast Ethernet line card for Cisco Catalyst® 4500 Series switches provides technology compatible with the IEEE



802.3 standard to support next-generation Metro Ethernet access networks. Bi-directional Fast Ethernet interfaces operate over a single strand of fiber, enabling significant cost savings in fiber and cable management, installation, and operation. This new line card offers both high-density connectivity and long-reach Fast Ethernet over fiber distances up to 6.2 miles (10 km).

cisco.com/go/catalyst4500

Security and VPNs

Cisco VPN 3000 Series Concentrator Software Version 4.1 (WebVPN)

Cisco VPN 3000 Series Concentrator Software Version 4.1 offers WebVPN functionality that enables clientless browser



access to internal Websites, e-mail, file transfers, and TCP-based applications originating from a Web browser. Cisco VPN 3000 Series concentrators support both IP Security (IPSec) and Secure Sockets Layer (SSL) remote-access virtual private network (VPN) connectivity. The WebVPN functionality is offered with no feature licensing fees.

cisco.com/go/webvpn

Cisco VPN 3020 Series Concentrator

The Cisco VPN 3020 Series Concentrator terminates remote-access VPN connections at enterprise central sites. This new model connects up to 750 simultaneous IPSec users with 50-Mbit/s performance using the Enhanced Scalable Encryption Processor (SEP-E) hardware module that provides support for Data Encryption Standard (DES), Triple DES (3DES), and Advanced Encryption Standard (AES) encryption. The concentrator accepts VPN connections from the Cisco VPN 3002 Hardware Client, the Cisco VPN Software Client, and from Web browsers using the new clientless WebVPN connectivity offering.

cisco.com/go/vpn3000

Edge Routing, Access, and Aggregation

Cisco 7600 Series Routers and Cisco Catalyst 6500 Series Switches: Supervisor Engine 720-3BXL and Enhanced FlexWAN Module

The new Cisco Catalyst® 6500 Series/Cisco 7600 Series Supervisor Engine 720-3BXL powers high-density,

line-rate Ethernet, private line, and subscriber services. Coupled with Cisco IOS® Software Release 12.2(17b)SXA, a suite of new services are enabled on the Cisco 7600 Series via new features including multipoint Ethernet, Virtual Private LAN Services (VPLS), hardware-accelerated Multiprotocol Label Switching (MPLS) VPNs, 10 and 1 Gigabit Ethernet IPv6-enabled connectivity, and point-to-point Ethernet/Frame Relay/ATM transport over MPLS. The new Supervisor Engine 720-3BXL delivers 30 million pps performance and integrates switch fabric and route processing into a single processor card, increasing performance and simplifying operations. The Cisco 7600 Series/Catalyst 6500 Series Enhanced FlexWAN module offers Cisco IOS feature parity with the Cisco 7500 and support for existing 7200/7500 port adapters, while delivering increased performance, for a natural, cost-effective migration from the Cisco 7500 to the Cisco 7600.

cisco.com/go/7600

cisco.com/go/catalyst6500

Cisco 1700 Series Modular Access Routers: Four-Port Switch WAN Interface Card

The new Four-Port 10/100BASE-T Fast Ethernet Switch WAN Interface Card (WIC) for Cisco 1700 Series modular routers is an intelligent, managed WIC that integrates LAN switching and routing functions for small businesses and small branch offices. This new WIC (WIC-4ESW) enables the router to deliver LAN services for connecting PCs, printers, wireless access points, and other devices. Support for standards-based IEEE 802.1Q virtual LANs, 802.1p traffic prioritization, and 802.1D Spanning Tree Protocol enables a variety of network configurations, such as a demilita-

rized zone for Internet access or LAN segmentation for IP telephony and wireless access workgroup.

cisco.com/go/1700

Content Networking

Cisco IP/TV System Version 5.1

Cisco IP/TV® delivers a complete, highly scalable, bandwidth-efficient solution for high-quality video communications over enterprise networks. Version 5.1 offers tight interoperability with Cisco Application and Content Networking Software (ACNS) Version 5.1 to more effectively support unicast and multicast “island” networks by leveraging ACNS features such as stream splitting, streaming automation, and the Cisco Streaming Engine.

cisco.com/go/video

Cisco 3700, 3600, and 2600 Series Routers: New 80-GB Content Engine Network Module

A content engine network module with an 80-GB hard disk is now available for Cisco 3700, 3600, and 2600 Series routers. Also available in 40-GB and SCSI controller configurations, each module integrates a content-delivery system into the router. Running Cisco ACNS Version 5.1 Software, the Cisco 2600/3600/3700 Series content engine network modules enable companies to extend the value of their branch routers for delivering new application services such as Web-application acceleration (including software distribution), secure Web content access management, and business and point-of-sale video.

cisco.com/packet/161npld

Storage Networking

Cisco MDS 9100 Series Switches

The Cisco MDS 9100 Series multilayer intelligent fabric switches enable IT organizations to grow storage-area networks (SANs) with intelligent network services and integrated management tools. Models

in this new series are single-rack-unit, fixed-configuration switches supporting 1 Gbit/s or 2 Gbit/s autosensing Fibre Channel connectivity. The Cisco MDS 9120 provides 20 ports, and the Cisco MDS 9140 provides 40 ports for Fibre Channel connections, enabling enterprises to build and manage small and mid-sized SANs or provide edge-to-core connectivity in larger SANs.

cisco.com/packet/161npd2

Cisco MDS 9000 Family: Caching Services Module

Designed for the Cisco MDS 9500 Series directors or the Cisco MDS 9216 Fabric Switch, the new Caching Services Module (CSM) will host the IBM TotalStorage SAN Volume Controller Storage Software. This solution allows data center managers to administer volume management, data replication, and point-in-time copies directly from the network for a single point of control and management across multiple storage subsystems. The CSM includes two independent nodes for the IBM software, offers 8 GB of cached memory, and redundantly configured components for high-performance, low-latency, and highly available SAN transactions. For a related article, see "Extending SANs," page 57.

cisco.com/packet/161npd3

Wireless

Cisco Aironet 1200 and 1100 Series Access Points: IEEE 802.11g Models

New models with IEEE 802.11g-compliant radios are available for the Cisco Aironet® 1200 and 1100 Series access points, supporting data rates up to 54 Mbit/s, five times faster than the current IEEE 802.11b standard. Access points with the 802.11g radio will work with any Wi-Fi-compliant 802.11b or 802.11g client, including all Cisco Aironet 350 Series 802.11b clients, Cisco Compatible clients, the Cisco Wireless IP Phone 7920, as well as the new Cisco Aironet 802.11a/b/g CardBus and PCI wireless LAN client adapters, which are slated to be available in the first calendar quarter of 2004. For investment protection, the

802.11g radio also is available in an upgrade kit for retrofitting installed 1100 and 1200 Series access points. Complementing existing support for Wi-Fi Protected Access, Advanced Encryption Standard (AES) is also supported in this new hardware and will be enabled for all Cisco Aironet 802.11g devices in 2004 via a free software upgrade after ratification of the IEEE 802.11i standard. All of the new 802.11g products are priced the same as their Cisco Aironet 802.11b-based versions.

cisco.com/go/aironet

CiscoWorks Wireless LAN Solution Engine Version 2.5

The CiscoWorks Wireless LAN Solution Engine (WLSE) eases deployment and operation of Cisco Aironet wireless LANs. Enhancements in Version 2.5 include detection of rogue access points, radio frequency (RF) scanning and monitoring, and detection of RF interference. A location manager displays the placement and status of wireless access points and bridges within a campus or building for easier management from a network operations center. The new feature for assisted site survey facilitates access point deployment and operation by determining optimal frequency selection, transmit power, and other settings.

cisco.com/go/wlse

Network Management

Cisco Broadband Access Center for ETTx

The Cisco Broadband Access Center for Ethernet to the Home and Business (BAC for ETTx) software helps service providers simplify management of broadband subscribers and services. Cisco BAC for ETTx provides features for automated service activation, subscriber self-registration and self-care, IP address assignment and management, device inventory management, reporting, and troubleshooting. An application programming interface (API) enables integration with a service provider's traditional management systems. For service

access, BAC for ETTx supports Cisco Catalyst® 2950 EI, Catalyst 3550, Catalyst 4000, and Catalyst 3500XL Series switches.

cisco.com/packet/161npd4

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 October 2003 and January 2004. For announcements of the most recently released products, see "News Archive, News Releases by Date" at newsroom.cisco.com/dlls/index.shtml.

ABOUT SOFTWARE: For the latest updates, versions, and releases of all Cisco software products—from IOS to management to wireless—visit the Software Center at cisco.com/kobayashi/sw-center/index.shtml (access to all content requires Cisco.com registration).

Time Warner Cable, Continued from page 63 services together in unique ways. For example, soon it will be possible for customers to see not only caller name and number displayed on their TV sets when the phone rings, but they will also have the ability to retrieve their voice mail through the TV set in between innings of their favorite baseball game.

Cisco BLISS is the first carrier-class, end-to-end PacketCable-compliant service for voice over broadband. The Cisco BTS 10200 Softswitch was formally PacketCable 1.0 qualified by CableLabs in April 2003—the first softswitch to receive such certification. The Cisco uBR7246VXR product was PacketCable 1.0 qualified in 2002—the first cable modem termination system (CMTS) to achieve this distinction.

“Time Warner Cable already had a strong relationship with Cisco,” says Campbell of his company’s decision to work with the networking technology provider, “and we wanted to go with a company that was sure to move forward with the PacketCable standard so that we knew the network would adhere to future standards.”

Ahead of the Game

Even though it has been forecast for some time that cable operators will move in to take some market share from existing telecommunications operators by offering VoIP services, only recently has the technology matured and the costs decreased enough for this vision to become a reality.

Already today, more people use cable modems for their home-based broadband connections, preferring them to other broadband options and giving cable operators a distinct advantage in the residential services market. With this milestone, cable operators have an even greater advantage: the ability to offer video, data, *and* voice services. In contrast, the best that phone companies can deliver over their networks is a double play of data and voice.

Time Warner Cable has been pleased with the acceptance rate for its Digital Phone service to date. “It has proven to us that this is going to be a good service to continue to roll out across the country,” says Campbell.

In baseball, double plays are impressive—but it’s the rare triple play that really

blows the crowd away. If the initial success of Time Warner Cable is any indicator, other cable operators that bundle voice with their other offerings should find a similar positive reaction to their “triple play” of cable services. ▲▲

Coming Second Calendar Quarter 2004

IP Communications

Not a subscriber?

Sign up for your
FREE subscription!

[www.cisco.com/go/
packet/subscribe](http://www.cisco.com/go/packet/subscribe)

PACKET
cisco.com/packet

Cogeco Cable, Continued from page 75 service delivery,” adds MacFarlane. “The cost savings, efficiency, and service versatility for voice, video, and data traffic over one network provided an attractive solution.”

Software-Based Migration

Cisco IOS Software provides investment protection for cable operators and service providers, because they can implement new features and services through software upgrades, without disrupting existing customers or the network infrastructure. For example, providers can add RPR to their networks through a software upgrade to the ML-Series interface, making it cost effective to add the new functionality. RPR enables point-to-multipoint and multipoint-to-multipoint services and adds valuable band-

width efficiency features because it creates a dual-rotating ring—sending Ethernet traffic in both directions on a ring to achieve the maximum bandwidth utilization on the SONET/SDH ring. An additional benefit of RPR is its ability to provide sub-50 microsecond resilience without requiring SONET/SDH protection.

Cisco enables RPR at up to 1 gigabit on the ML-Series interface and provides interoperability with routers and switches to deliver consistent end-to-end QoS through Cisco IOS Software. This functionality enables providers to establish end-to-end SLA management across resilient packet rings and Ethernet/IP networks.

The ML-Series interface also provides direct support for SDH, allowing global

operators to easily add these capabilities to their SDH infrastructures. ▲▲

FURTHER READING

- Cisco COMET optical solutions for service providers:
cisco.com/packet/161_8e1
- ML-Series data sheet:
cisco.com/packet/161_8e2
- Cisco white paper on Metro Ethernet services:
cisco.com/packet/161_8e3
- Q&A on Cisco optical products and technology:
cisco.com/packet/161_8e4

Mobile Messaging, Continued from page 74

operators can move all SS7 traffic, including ISUP, to proven, robust SS7overIP networks. In doing so, operators can completely control their transition to next-generation networks.

LogicaCMG IP-Enabled Messaging Architecture

LogicaCMG is the worldwide leader for SMS, with more than 55 percent of the market using their systems. The close integration and cooperation of LogicaCMG and Cisco has led to a surge in uptake of SS7overIP technology, enabling increased bandwidth (transactions) to SS7 applications and dramatically reducing product (CapEx) and operational (OpEx) costs. And perhaps equally important, by moving or “offloading” this SS7 traffic to IP networks, additional benefits are realized as the traffic burden on existing, expensive legacy SS7 networks is alleviated—effectively freeing up ports and capacity on the SS7 transport infrastructure.

The growing trend toward IP-enabled SMSCs offers a perfect example of the benefits that can be realized. SMSCs are the SS7 network elements that provide SMS, also known as text messaging. On an SS7 network for transport and delivery of SMS messages, these

nodes have traditionally used SS7 connectivity.

With the Cisco ITP Signaling Gateway as a component of its messaging architecture, LogicaCMG’s SMSC connects via the SIGTRAN SUA protocol (IP) to the ITP, which connects to the SS7 network via traditional low-speed or high-speed SS7 links as well as SIGTRAN/IP (see figure, page 73). Providing not only traditional SS7 redundancy characteristics to the SMSC, the ITP also provides load-balancing services across multiple SMSCs. With its MLR capability, the ITP can also be used to customize traffic flows to and from specific SMSCs, or voting applications based on a multitude of TCAP, SCCP, and MAP/IS-41 layer parameters. Formally released to the market in mid-2003, this IP-based messaging solution has gained broad acceptance among mobile operators and is being deployed worldwide.



With ever increasing momentum, mobile wireless and wireline operators are recognizing the benefits of SS7overIP in their networks. Dramatic capital and operational costs savings, coupled with increased capacity and bandwidth, provide compelling motivation for deploying these industry-hardened solutions. ▲▲

PACKET ADVERTISER INDEX



ADVERTISER	URL	PAGE
ADC—The Broadband Company	www.adc.com/performance	D
AdTran	www.adtran.com/info/wanemulation	2
Aladdin Knowledge Systems	www.ealaddin.com/cisco	IFC
American Power Conversion (APC)	www.apc.com	4
BellSouth Business	www.bellsouth.com/business/netvpn	F
Boson Software	www.boson.com/packet	A
Cisco Networking Academy® Program	www.cisco.com/go/careerconnection	60
Cisco Press	www.ciscopress.com/ccnpgiveaway	B
Cisco Systems	www.cisco.com/securitynow	Cover Gatefold/34
Dalhousie University	www.dal.ca/internetworking	58
eiQ Networks	www.eiqnetworks.com	8
Interstar Technologies	www.faxserver.com/packet	84
netForensics	www.cisco.com/go/sims	76
NIKSUN	www.niksun.com/packet	12
NTT Communications	www.nttverio.com	54/55
OPNET Technologies	www.opnet.com	OBC
Panduit	www.panduit.com/vip04	IBC
RAD Data Communications	www.radfrisco.com	44
RedSiren	www.redsiren.com/packet.htm	22
SurfControl	www.surfcontrol.com	48
Unisys	www.unisys.com/security	16
Websense	www.websense.com	28

Cache File

SNIPPETS OF WISDOM FROM OUT ON THE NET

Top 20 Internet Security Vulnerabilities

The lion's share of worms and other successful cyber attacks are made possible by vulnerabilities in a small number of common operating system services—a conclusion underscored by the SANS Institute's most recently updated list of Top 20 Internet Security Vulnerabilities. Topping the list of vulnerabilities to Windows systems are Internet information services, Microsoft SQL Server, and Windows authentication. BIND Domain Name System, remote procedure calls, and Apache Web Server head the list of vulnerabilities to UNIX systems. Find out more at sans.org/top20/.

Text Messaging Big Among The Brits

Mobile phone users in the UK broke a text-messaging record on January 1, 2004—with 111 million new-year greetings sent via text, an 8 percent increase over the 102 million messages sent on New Year's Day 2003. The Mobile Data Association (mda-mobiledata.org) reports that UK mobile phone owners send an average of 56 million text messages in a typical day across the four UK Global System for Mobile Communications (GSM) networks, with about 2.3 million messages sent every hour. Holidays and sporting events are among the highest text messaging days for mobile users in the UK.

CYBER QUOTE

"THERE ARE two ways TO WRITE error-free programs. ONLY THE third one WORKS."

—Anonymous

DIRECTORY FOR HOTSPOTS

Looking for Wi-Fi hotspots while on the road? For your next vacation or work-related trip? Check out wi-fihotspotlist.com, one of the top resources for locating wireless access points around the world.


CANADIANS ARE CONNECTED

The Yankee Group reports that 48 percent of Canada's home Internet users access the Net at least three times per day, and 71 percent of those users spend at least 15 minutes online during each session. Most Canadian Internet users are logging on via a broadband connection, either DSL or cable modem. And according to Statistics Canada, while e-mail remains the most-used activity by the majority of surfing households, e-banking and online shopping have significantly grown in popularity among Canada's home users over the past five years.

Most Popular IM Acronyms

Among the 10 most commonly used instant messaging (IM) acronyms in the workplace are BRB (be right back), CTRN (can't talk right now), IMO (in my opinion), HTH (hope this helps), and IAM (in a meeting), according to a survey conducted by Omnipod, Inc. (omnipod.com). With currently more than 100 million IM users worldwide, Gartner Group forecasts that by 2006, IM will be used more often than e-mail as the preferred method of messaging in the enterprise.

THE 5th WAVE



"A centralized security management system sounds fine, but then what would we do with the dogs?"

©The 5th Wave, www.the5thwave.com