

PACKET

I Want
My
IPTV!

Access a World
of Content—on Any
Device, Anywhere

3 Steps to Network
Virtualization

Time to Migrate
Your Routers?



contents

COVER STORY

32

I Want My IPTV!

In fact, consumers want everything, on anything. Networked entertainment in the home is about delivering more and varied media-rich content to today's media-savvy consumers through an array of technologies more involved than just the home PC and TV.



FEATURES

38

A View from the Net

Cisco acquires SyPlix and brings IP-enabled video security to enterprise networks.

43

The Connected Home

Networked entertainment that puts you in control.

46

Time to Migrate?

Making way for the next generation of routers.

departments



1 From the Editor

I Want Everything, on Anything

3 Mail

4 Datagrams

18 Reader Tips

19 Tech Tips

79 NetPro Expert

Wireless Security

81 Advertiser Index

82 Cache File

TECH TIPS + TRAINING

7 Flexible NetFlow

New IOS feature delivers greater visibility into your network.

11 Is Your Information Security Working?

Security tools are the exceptions to the "set it and forget it" approach.

15 Meeting US IPv6 Mandates

Six steps to on-time, affordable compliance.

17 Gaining an Edge on the CCIE

New lab simulates rigorous CCIE exam experience.

21 Five Ways to Ensure VoIP Reliability

All the redundant hardware in the world won't compensate for poor design.



CHALK TALK

BEST PRACTICES

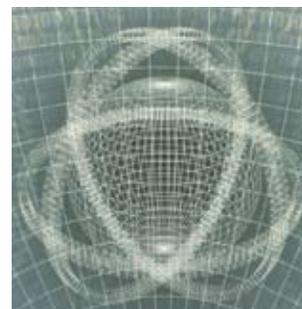
23 Securing the Edge

Solving the trust issue with BGP.

BEST PRACTICES

29 Managing MPLS

An overview of MPLS OAM tools, techniques, and standards.



INFRASTRUCTURE

DATA CENTER

51 Three Steps to Network Virtualization

Take all three with the Cisco Catalyst 6500 Series Switch.

ROUTING

55 7200 Gets a Boost

Performance, capacity gains meet rising WAN/MAN services aggregation needs.

VOICE

59 Can We Talk?

Effective call admission control for complex networks.

SERVICE PROVIDERS

63 First HDTV over IP in the US

With IP over fiber to the home, SureWest Communications delivers desired services *now*.

67 High Availability for MPLS

Increasing service availability through fast recovery from network disruptions.

BEYOND SPEEDS + FEEDS

71 Unified WLAN in Access Layer

New Cisco Catalyst 3750G Integrated WLAN Controller.

72 New Product Dispatches

77 Product Review

Cisco ASA 5500 Adaptive Security Appliance.

PACKET

DAVID BALL
Publisher and Editor in Chief

JENNIFER REDOVIAN
Executive Editor

SUSAN BORTON
Managing Editor

SUZANNE JACKSON, JOANIE WEXLER
Contributing Editors

ROBERT J. SMITH
Sunset Custom Publishing
Project Manager/Account Supervisor

NICOLE COLLINS, AMY MACKEY
Sunset Custom Publishing Production

EMILY BURCH
Art Director/Designer

ELLEN SKLAR-ABBOTT
Diagram Illustrator

BILL LITTELL
Print Production Manager

VALERIE MARLIAC
Promotions Manager

DWIGHT ESCHLIMAN
Cover Photograph

ADVERTISING INFORMATION:
Kristen Bergman, 408 525-2542
kbergman@cisco.com

PUBLISHER INFORMATION:
Packet magazine (ISSN 1535-2439)
is published quarterly by Cisco Systems.

Please send address corrections
and other correspondence direct to
packet@cabmwest.com.

Aironet, Catalyst, CCDA, CCIE, CCNA, Cisco, Cisco IOS, Cisco Networking Academy, Cisco Press, the Cisco Powered Network logo, the Cisco Systems logo, Cisco Unity, IOS, iO, Linksys, Packet, and PIX 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 © 2006 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

ROB BRODMAN

I Want Everything, on Anything

A TYPICAL HOME IN THE US HAS THREE OR MORE TVs, ONE OR MORE PCs, AND A VARIETY OF STEREO, DVD, and video-game products—along with an assortment of mobile gadgets such as PDAs, cell phones, and portable media players. While an increasing amount of digital content (video, voice, data, and music) flows among these devices, an increasing number of consumers—that means, *us*—are looking to get any content we want, when we want it, on any device we want.

In the home, this desire for “everything, on anything” is satisfied with an all-IP wired or wireless network. Our cover story, “I Want My IPTV!” (page 32), lays out the networking, multimedia, and video equipment involved in creating the “Connected Home,” and the implications to service providers of delivering the high-quality experience consumers demand. Our own rendering of the Connected Home—sporting Linksys and Scientific Atlanta gear—is on page 43.

Today, service providers and carriers are in a race to deliver broadband to the networked home. Integrating service delivery to meet consumers’ demand for information, entertainment, and communication how and when they want it depends less on the access technology than on the quality-of-experience capabilities in provider networks. Technologies that boost the quality and efficiency of multimedia service delivery are paramount. “Managing MPLS” (page 29) focuses on the service assurance functionality of MPLS OAM and the tools and features available in Cisco routers.

For enterprise networks, “Securing the Edge” (page 23) dives into the policies and protections that make BGP an ideal protocol choice when you’re connecting to networks outside your administrative control. “Three Steps to Network Virtualization” (page 51) provides information on how your company can transform physical network devices into virtual resource pools.

In the “Tech Tips + Training section,” discover the six steps for “Meeting US IPv6 Mandates” (page 15), “Five Ways to Ensure VoIP Reliability” (page 21), and how “Flexible NetFlow” (page 7) can help you increase visibility into your network. And after you find out if your information security is working (page 11), test your knowledge of security with our quiz on page 17.

In a break from the technical content, make sure to read “Time to Migrate?” (page 46) for an overview of Cisco’s retirement of the 1700, 2600, and 3700 series routers, and tips to ease your migration strategies. ■



David A. Ball

DAVID BALL
Editor in Chief
daball@cisco.com

VPN Types Versus Interfaces

ON PAGE 8 OF THE last issue of *Packet* [Second Quarter 2006], the author, Mark Lewis, seemingly neglected to mention multipoint GRE (mGRE) technology, which expands the GRE implementation area further than described in the article and makes it closer to the MPLS IP VPN area. Can you clarify?



SERGEI A. CHERNOOKI, CCNP
NPP Belsoft, Inc.,
Minsk Republic of Belarus

EDITOR'S NOTE: The following response is from author Mark Lewis.

You might have been a little confused for the following reasons:

1. Multipoint GRE is neither a VPN type nor a tunnel type; it is an *interface* type used on Cisco routers. The article focused on VPN types.
2. Your message implies that mGRE expands GRE so that GRE capability is similar to MPLS IP VPNs. Again, mGRE is actually an *interface* type, rather than a VPN or tunnel type, and can in fact be used *in support of* an MPLS VPN-over-GRE solution on Cisco routers (RFC 2547 describes MPLS IP VPNs.) For more information, refer to cisco.com/packet/183_2a1.

You might have confused this with Dynamic Multipoint VPN (DMVPN), a Cisco GRE/IPsec VPN technology that takes advantage of mGRE interfaces. When used with DMVPN, mGRE interfaces can reduce configuration complexity when compared with regular GRE/IPsec or IPsec VPNs.

Some advocates of DMVPN, with limited knowledge of each technology, often imply that DMVPN (using mGRE interfaces) competes directly with MPLS Layer 3 VPNs. In fact, they do not compete; they are often complementary technologies and are used for different purposes.

If you are interested in a more detailed discussion and comparison, my latest book, *Comparing, Designing and Deploying VPNs*, available from Cisco Press (ciscopress.com), describes these technologies in greater detail.

More About Frame Relay Reader Tip

I read an interesting reader tip about backup solutions for Frame Relay in your Second Quarter 2006 issue. An important point is that you must configure your Frame Relay end-to-end keepalive mode reply before configuring your Frame Relay end-to-end keepalive request and then check if your backup solution is working properly. If not, you can lose connectivity to the remote router. I was lucky that my backup solution (ISDN) was working fine, so I connected the remote router via a backup link.

MURILO WOZEN
SerraOn, Petrópolis, Brazil

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.

CORRECTION

The Mail page of our First Quarter 2006 issue contained a typographical error. Reader Robert McCallum's CCIE number was incorrectly printed as 875. The correct number is 8757. Our thanks to Stephen Green, who noticed the error and felt compelled to write in and point out a bit of CCIE lore—namely that CCIE numbering started at 1024, in honor of the number of bytes in a Kilobyte, of course.—*Editors*

Passing the CCIE with a Little Help from Packet

I truly value *Packet* and have greatly benefited from it. The magazine has provided me with invaluable technical assistance, and the credit goes to the people involved in publishing and distributing the magazine to all parts of the globe. Last Saturday, I passed my CCIE written exam with a high score. Now you know why I needed *Packet* so badly. Thank you to the *Packet* team for playing an important role in keeping my students and me informed of current happenings in the Cisco world...and helping me attain my CCIE.

PRADEEP VARADARAJAN
Bangalore, India

datagrams

Renting Movies with a Box and a Beam

A company backed by Disney, Intel, and Cisco has come up with yet another movie-delivery mechanism. It's a slim, silver, good-looking set-top box called MovieBeam.

The MovieBeam player connects directly to your TV set. Whenever you're in the mood for a movie, you choose from the list of 100 movies on the player's hard drive. There's no monthly fee and no minimum; you're billed only for the movies you watch. You can rewind, pause, fast-forward and replay a movie you've bought—for 24 hours from your first glimpse of the opening credits.

Each week, seven or eight new movies magically show up in the player's list, pushing an equal number of old ones off the list.

MOVIEBEAM features a slim, silver, good-looking set-top box.



This wireless movie-delivery feature gives MovieBeam its name. The company doesn't require an Internet connection or even a computer. Nor does the service depend on what cable or satellite setup you have, if any.

MovieBeam's movies are encoded in the broadcast signal of Public Broadcasting System (PBS) stations across the US. You're actually receiving MovieBeam's movies at this very moment—but they're invisible unless you have the MovieBeam box. (MovieBeam pays PBS for these piggybacking rights.)

MovieBeam is available in 29 major metropolitan areas, including Atlanta, Chicago, Houston, Los Angeles, New York, Philadelphia, and Washington. Check availability in your area at moviebeam.com. The company plans a large geographical expansion in the next year.



LAB REPLICATES CUSTOMER ENVIRONMENT

Cisco customers can actually interact and play with the technologies in the Cisco DNA lab, a new lab that showcases Cisco's data center solutions to customers. The lab replicates a typical enterprise customer's data center, comprising storage networking, high-performance compute networking, and high-density server farm networking. The DNA lab demonstrates major Cisco data center technologies, including Ethernet, Infiniband, Fibre Channel, optical metropolitan-area networking, and wide-area application acceleration services. The lab also enables the experience for remote customers, with live demos already conducted for customers as far away as India and Europe. For more information on the DNA lab, contact your Cisco account manager.

DILBERT ON THE DESKTOP

CARTOONIST SCOTT ADAMS, creator of the famous “Dilbert” cartoon strip, recently used Cisco Unified MeetingPlace to draw and collaborate with a nationwide audience in the US, highlighting the potential of Cisco Unified Communications to improve collaboration and illustrating how the network can serve as a platform for life’s experiences.

BY INTEGRATING VOICE AND WEB conferencing capabilities, the rich-media conferencing solution allowed Cisco customers and partners to brainstorm with Scott Adams and watch as he developed their ideas into a Dilbert cartoon strip.

TO SEE SCOTT ADAMS using Cisco Unified MeetingPlace to create a cartoon strip, visit newsroom.cisco.com/video/dilbert.html.



Networkers: Buzz And Blogs

Communication networks are moving beyond the office and becoming the platform for life’s experiences, based on a survey of 8,000 attendees at Cisco Networkers, Cisco’s annual users’ conference.

An overwhelming majority of attendees attributed life-changing experiences to the communications, collaborations, and success they have achieved through online experiences. Responses ranged from claims that

they met their spouses online to having used the Internet to keep in touch with family and friends.

When asked what technologies they plan to deploy soon that will give their business a competitive advantage, respondents listed voice-over-IP applications (77.64 percent), ahead of video on demand (46.23 percent) and enterprise instant messaging (36.29 percent).



New Network Lexicon

The network has even spawned its own jargon. Attendees cited personal buzzwords, including “green noise” (enthusiasm generated when working as a team) and “world without walls” (WWW). And while some may be concerned about developing a “webpendency”—a reliance on the Internet to effectively function—most reported positive experiences, even approaching “webvana.”

For more information on activities at Networkers 2006 in Las Vegas, see the blog at <http://blogs.cisco.com/networkers/>.



WORLD OF SOLUTIONS showcases Cisco and partner solutions.

Recently Announced Cisco Acquisitions

Acquired		Employees	Location
Audium	Provides VoiceXML speech self-service application development and management environments. The acquisition will enable enterprises to build automated voice response applications that are integrated with their converged IP network and can work within their Services Oriented Architecture (SOA), enabling common services across the network. Audium will become part of Cisco’s Voice Technology Group.	26	New York, New York, USA
Meetinghouse	Provides client-side 802.1X supplement security software that allows enterprises to restrict network access to only authorized users or host devices that attempt to gain access to networked resources through wired and wireless media. Meetinghouse will become part of Cisco’s Wireless Networking Business Unit.	77	Portsmouth, New Hampshire, USA
Metreos	Provides IP communication application development and management environments. Metreos technology is a platform for integrating Cisco’s Unified Communications System with enterprise business applications. Metreos will become part of Cisco’s Voice Technology Group.	19	Austin, Texas, USA

Flexible NetFlow

NEW IOS FEATURE DELIVERS GREATER VISIBILITY INTO YOUR NETWORK.

by tom zingale

V

isibility into your network is no longer a luxury. It's a necessity. In response to new security threats, business requirements, and IT demands, network operators are finding it critical to understand how the network is behaving: application and network usage; network productivity and utilization of network resources; the impact of changes to the network; network anomaly and security vulnerabilities; long-term compliance, business process, and audit trail. In short, we need a solid understanding of who, what, when, where, and how network traffic is flowing. / To help meet these new requirements and demands, Cisco is taking flow technology to a new level with *Cisco IOS Flexible NetFlow*. Flexible NetFlow promises to enhance network optimization, reduce costs, and improve capacity planning and security detection beyond what other flow-based technologies can offer today.

Application Tracking Flexible NetFlow provides the ability to track exactly the information you need for your organization. By targeting specific data, the amount of flow information and flow export is reduced, allowing for enhanced scalability.

For example, if you're interested in TCP application analysis, Flexible NetFlow can track source and destination IP addresses and TCP source and destination ports, and also examine the packets for this data. This information will effectively show who is sending and receiving the traffic per application port. In traditional NetFlow, aggregation comes with the expense of lost information; however, in Flexible NetFlow, you can actually track multiple sets of information to ensure that all flow information in the network is captured efficiently.

Security Detection Flexible NetFlow is an excellent attack detection tool with capabilities to track all parts of the IPv4 header and even packet sections, and characterize this information into flows. It is expected that security detection systems will listen to NetFlow data and, upon finding an issue in the network, create a virtual bucket or virtual cache that will be configured to track specific information and pinpoint details about the attack pattern or worm propagation. Flexible NetFlow has the capability to create *caches on the fly* that contain specific information combined with input filtering (i.e., filtering all flows to a specific destination), making it a better security detection tool than current flow technologies.



A CLOSER LOOK

Flexible NetFlow enhances your ability to detect security incidents and understand traffic behavior in your network.

Common attacks, such as port scans for worm target discovery and worm propagation, are tracked in Flexible NetFlow. So, for example, if the security detection server understands such an attack, it might program another virtual cache or bucket to export payload information or sections of packets to take a deeper look at a signature within the packet. This is just one of many examples of how Flexible NetFlow can be used to detect security incidents.

Key and Non-Key Fields

EACH PACKET THAT IS FORWARDED within a router or switch is examined for a set of IP packet attributes. These attributes are the IP packet identity or *key fields* for the flow and determine whether the information in this packet is unique or similar to other packets. For example, all packets with the same source/destination IP address, source/destination ports, and class of service are grouped into a flow and then packets and bytes are tallied. This methodology of flow characterization or determining a flow is scalable because a large amount of network information is condensed into a database called the NetFlow cache (see Figure 1).

Additional information, or *non-key fields*, can be added to the flow record. The non-key fields are not used to create or characterize the flows but are simply added to the flow. Example non-key fields might be packet counters, routing next-hop, and other fields.

Key Components of Flexible NetFlow

FLEXIBLE NETFLOW HAS THREE key components: *flow monitor*, *flow record*, and *flow exporter*.

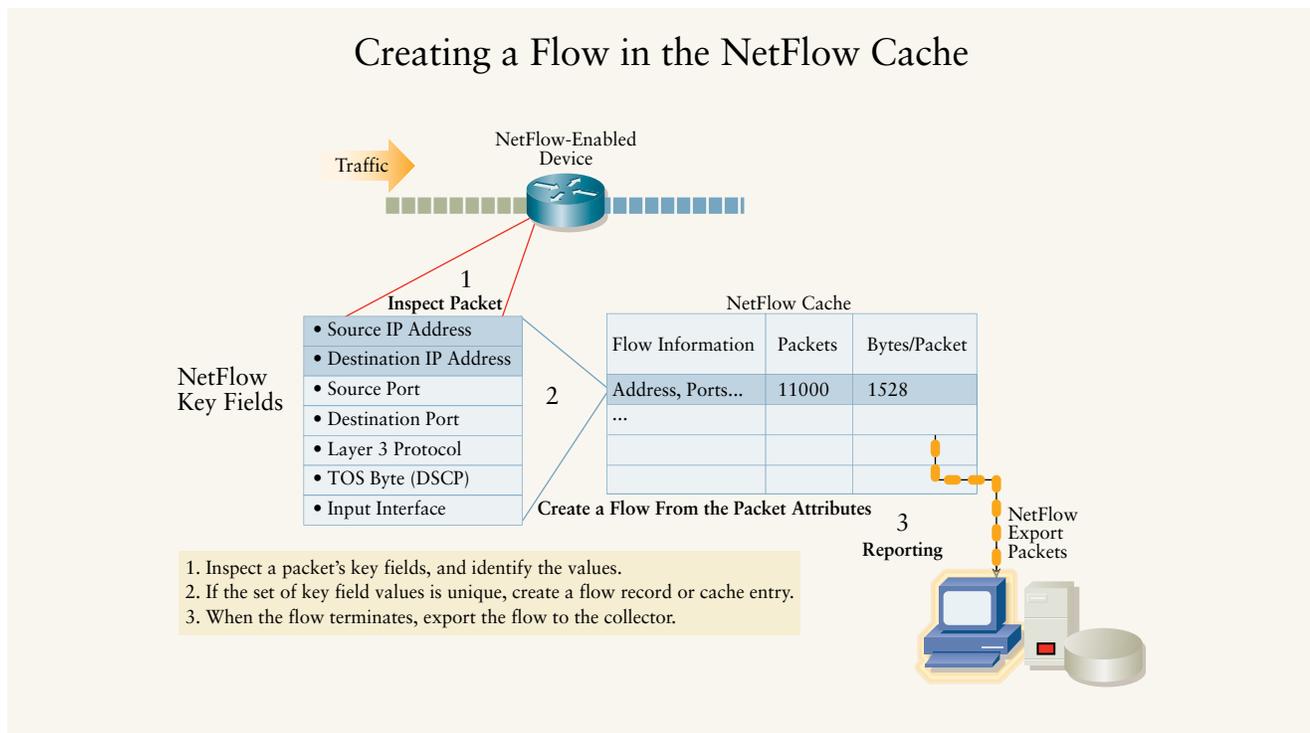
The flow monitor stores flow information and contains the flow record and flow exporter. Multiple flow monitors can be configured per interface (see Figure 2). The flow record is a pre-defined or user-defined set of packet attributes to track NetFlow information including the IPv4 header, routing, and sections of packet data. The flow records in the cache will expire or terminate and be exported to a NetFlow collector and used to create management reports.

The flow exporter allows you to define where the export can be sent, and the type of transport and properties for the export. The flow exporter supports various export formats including v5, v9, and the IETF IP Flow Information Export (IPFIX) standard. The flow exporter also supports various transport protocols including UDP and Stream Control Transmission Protocol (SCTP).

Flexible NetFlow and Version 9

THE BEST METHOD for exporting a wide range of information from the packet is to use NetFlow version 9. *Without version 9 export format, Flexible NetFlow would not be possible.* A key advantage of Flexible NetFlow is that when you configure a flow record, it is effectively converted to a version 9 template and then forwarded to the collector. NetFlow

FIGURE 1 In creating a flow based on a packet's key fields, a large amount of network information is condensed into a database called the NetFlow cache.



Sample of Flexible NetFlow Customizable Flow Monitors

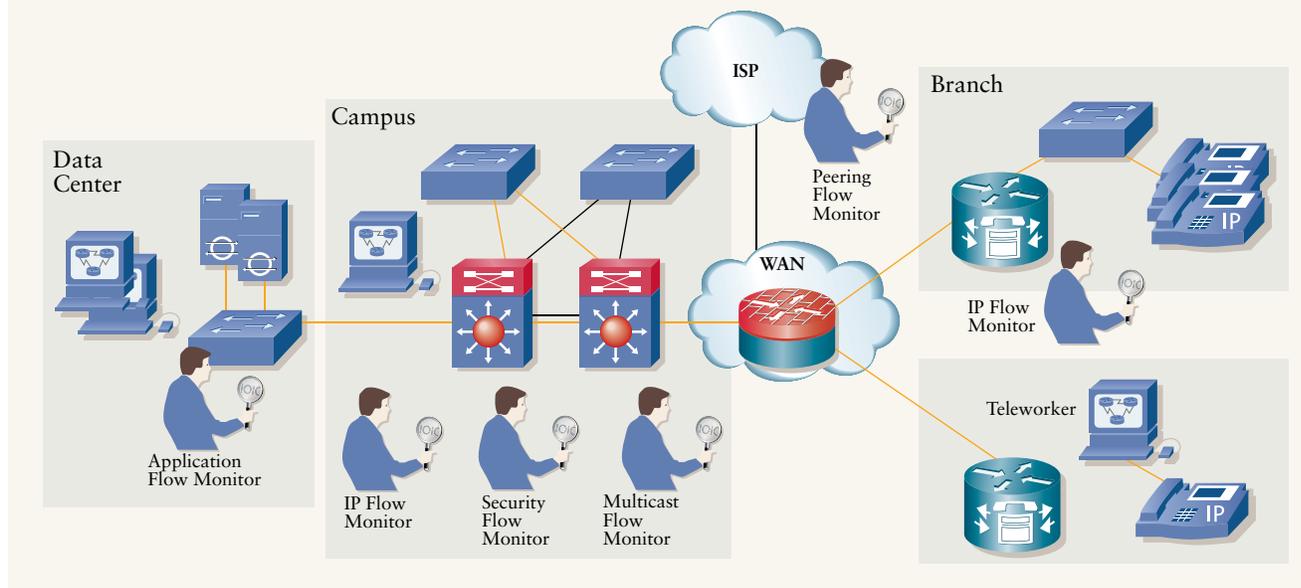


FIGURE 2 Multiple flow monitors can be configured per interface.

version 9 will periodically export the template data, so the collector understands what data is to be sent and exports the data FlowSet for the template.

NetFlow Version 9 Export Packet						
Packet	Template	Data	Data	-	Template	Data
Header	FlowSet	FlowSet	FlowSet	-	FlowSet	FlowSet

Configuration Examples

CONFIGURING FLEXIBLE NETFLOW can be quite easy with predefined or user-defined flow records.

1. Configure the exporter to send NetFlow data to a collection server.
2. Create a flow record to define flow information to capture.
3. Attach the flow record and flow exporter to the flow monitor
4. Add the flow monitor to the interface to monitor either ingress (input) or egress (output) traffic.

For a Flexible NetFlow configuration example, let's take the aforementioned task of monitoring how much traffic will be used per TCP application. In a user-defined flow record configuration, the "match" keyword is used to denote the field

as a key field and to characterize or create flows. The "collect" keyword denotes the non-key field and will be used for information that we want to add to the flow but not used when creating the flow.

flow record app-traffic-analysis

```
match transport tcp destination-port
match transport tcp source-port
match ipv4 destination address
match ipv4 source address
collect counter bytes
collect counter packets
```

flow exporter export-to-server

```
destination 172.16.1.1
```

flow monitor my-flow-monitor

```
record app-traffic-analysis
exporter export-to-server
```

interface Ethernet 1/0

```
ip flow monitor my-flow-monitor output
```

Following is a configuration example of packet section export for security monitoring. Consider a UDP port 53 anomaly spike in traffic to a Domain Name System (DNS). We will use deep packet inspection and send a section of packet to the NetFlow

collector (e.g., 1000 bytes into the payload). The NetFlow collector can then determine whether the DNS query is legitimate or an anomaly DNS attack. The flow monitor is tracking source and destination IP addresses along with UDP destination port. We can see the first packet of the flow, every packet of the flow, or use 1 in N packet sampling (e.g., sample 1:400 packets randomly). Notice that the interface has two separate flow monitors, both configured simultaneously.

flow record packet-section

```
match ipv4 section payload size 1100
match transport udp destination-port
match ipv4 destination address
match ipv4 source address
collect counter packets
```

flow monitor section-export

```
record packet-section
```

interface Ethernet 1/0

```
ip flow monitor my-flow-monitor input
ip flow monitor section-monitor input
```

Further Reading

- NetFlow
cisco.com/go/netflow
- NetFlow Technical Overview
cisco.com/packet/183_4a1

Flexible NetFlow is an important technology available in Cisco devices to enhance your ability to detect security incidents and understand the behavior of traffic in your network. An improved version of traditional NetFlow, Cisco IOS Flexible NetFlow brings greater scalability, data aggregation, and user customization to your network. **P**

TOM ZINGALE is a product manager in the Internet Technology Division at Cisco. He can be reached at tomz@cisco.com.

Is Your Information Security Working?

SECURITY TOOLS ARE THE EXCEPTIONS TO THE "SET IT AND FORGET IT" APPROACH. by bill young

No doubt, information security has grown by leaps and bounds. But as types of security devices and techniques grow in complexity and number, is your security actually improving? What about regulatory compliance? Are we on a downward-trending road paved with good intentions?

Defense in Depth

ASK SECURITY ENGINEERS what keeps them up at night. More than likely they will tell you it's a known weakness in their environment that they just don't have the time to address—something they suspect they might have overlooked, or some vulnerability they strongly suspect exists but simply don't know about. These are the kinds of thoughts that spook the information security expert.

Security professionals acknowledge these probabilities and attempt to mitigate them using a defense-in-depth strategy to reduce their company's risk from individual weaknesses.

So, what exactly is defense in depth? Put simply, it's the idea that implementing multiple layers of protection reduces your company's exposure. If one device or layer in the DMZ is compromised, you can still protect the company's critical or proprietary data. A successful defense-in-depth strategy requires trained personnel, effective tools and processes, policies and end-user training. While it's impossible to completely remove the risk of an incident, defense in depth is the most effective way to reduce its impact—and sleep better at night.

Good Intentions

REGULATORY COMPLIANCE AND AUDITS in the industry are becoming prevalent, and more companies and industries are striving to meet legal and other security requirements. The US Sarbanes-Oxley Act, HIPAA, and FISMA are just a few of the regulations

driving organizations to add increased information security controls. Successful implementation of these controls should be increasing the effectiveness of defense in depth. And, indeed, we see that these regulations are succeeding in motivating organizations to tackle security with some urgency, rather than a promise to do more next fiscal year.

Every year, the US Computer Security Institute and Federal Bureau of Investigation (FBI) release the Computer Crime and Security Survey. The most current survey is from 2004, and it shows that we might have reached a turning point. We are seeing the first signs of a decline in compromises and the dollars lost per incident. This is heartening and does show that risks can be reduced to manageable levels, but we still have a long way to go. For example, we continue to see an increase in unauthorized access, Website incidents, and viruses. A laundry list of new compromises comes out weekly.

Despite the fact that security awareness is at an all-time high, our defensive strategies have significant room for improvement.

The Devil Is in the Details

WHILE SECURITY ENGINEERS share a common goal of protecting their respective organizations, their underlying backgrounds are often very different:

- Windows administrators
- UNIX administrators
- Network/firewall administrators
- Policy developers and auditors

Each of these skill sets is valuable for effective security protection. But human nature compels us to focus on our specific strengths. We tend to spend time on what we understand or find interesting. This inclination produces a tendency to install security controls in our strongest disciplines, and then

Continued on page 13



INFORMATION SECURITY is caught between the high costs of effective security measures and the high costs of being hacked.

configure additional security controls that ultimately protect against the same or similar issues.

The result is that many of the unique defense-in-depth “layers” are often all within the same discipline. A UNIX administrator might configure a wide variety of security controls, but if the company doesn’t have a password policy, that system could still be easily compromised.

This reality is why it behooves security administrators to avoid operating in a limited set of security comfort zones. They need to work across all the relevant areas. Teams can do this better by dividing up responsibilities so that members can focus on what they do well, or what they like doing.

It Takes People to Run the Tools

WHILE COMPANIES CONTINUE TO INCREASE their capital budget for security tools, they’re not necessarily increasing the staffing budget to properly deploy and manage those products. The continually rising demand for security defensive strategies has led to a significant increase in the products available. Every year companies approve budgets for the next greatest security tool. This is often driven by audits or a network breach that occurred in the prior year. We have an extensive selection of security tools available to us, including one-time password generators, intrusion prevention systems, system hardening tools, and inline network antivirus appliances.

But security tools are almost always the exceptions to the “set it and forget it” approach that’s common in computing. It takes a surge of engineering resources to make the products effective by properly installing and customizing the solution to specific environments. Deploying new products also puts a drain on your ongoing support resources. If your environment is constantly changing, these tools must be constantly updated. On top of that, new vulnerabilities arise every week, signatures become available, and the tools must be made “aware” of new devices, such as Web servers, that are added to your network.

If inadequate time is put into a deployment or ongoing support resources are not available, many of these solutions functionally become shelfware. It’s common to look to technology to solve our problems. However, in both security and in network management, people, time, and system integration/tuning are also part of what’s required for effective solutions.

One question that I commonly ask engineers is “why” they last looked at their firewall, system, or IDS logs. The response that I almost always get is “Because I was troubleshooting a problem.” These tools were installed to monitor, protect, and notify us of security incidents, yet we’re not listening.

The Audit Trap

AUDITS ARE ARGUABLY the single best thing that has happened for information security awareness. Audits often lead to policy development, user education, better patching, and an overall improvement in an organization’s security posture. Unfortunately, audits are also training the industry to be “a better liar.”

Passing an audit is often critical to a company’s continued operation. This puts pressure on engineers to pass the audit at all costs, even if it means covering up known issues. When an audit is coming, the dust is blown off of a wide suite of tools, many of which were installed purely for audit compliance. Engineers claim that they’ve been reviewing and analyzing the data. The auditor is shown screens of alerts in the network operations center when the intrusion prevention system sees an issue (a screen that might only have been checked the morning before the auditor arrived to make sure it was functioning). Because the audit covers such a wide range of topics, the auditor is often not a qualified expert on each discipline being audited.

The bottom line is that information security is about discipline and process execution.

Audits also drive us to focus our energies inefficiently. One common audit requirement is that security personnel must initial each page of a log. When there are 100 or more pages to be initialed each and every morning, they are seldom being read in detail. Security event correlation and aggregation tools can be used to increase efficiencies, but they are still resource-intensive and may not meet audit requirements.

After the audit is complete and the results reviewed, one of two things most often occurs:

1. If there are no significant findings, the company goes back to business as usual, with a false sense of security that no serious issues still exist.
2. If there are problems, the results are waived in front of management, proclaiming that additional capital is needed to buy another tool to ensure that the company can pass the next audit.

Start Sleeping Better at Night

WE STILL HAVE A NUMBER of issues that need to be addressed even though there have been some improvements. The fundamentals of information security require knowing about risks and making informed decisions to react or accept the risks. Skills deficiencies, staffing shortages, and audit compliance pressures are preventing many of these issues from being appropriately evaluated and handled.

Often, consultants are engaged to help with the first phase of a new product deployment. Bringing in that initial surge of resources can help make certain that a product is deployed with the due diligence necessary to ensure proper integration. However, to be successful, effective handoff and training is required. And, most importantly, there needs to be sufficient full-time staff to manage and maintain these tools.

The bottom line is that information security is about discipline and process execution. It's a constantly changing and

improving process, caught between the high costs of effective security measures and the high costs of being hacked. We've made industry-wide progress and are beginning to reduce the risk and impact of attacks. But there is no one tool that will solve all security problems. We need to also focus on the "people and procedures" side of things.

The key to defense in depth is full cooperation and resolve from security, network, systems, and compliance professionals and their management. With a clear understanding of all the resources necessary—people as well as tools—a truly effective defense-in-depth security program can be implemented and maintained. ■

BILL YOUNG is a senior security consultant at Chesapeake NetCraftsmen, LLC (www.netcraftsmen.net), which delivers high availability solutions for network design, operating systems, applications, security, storage, and IP telephony.

Meeting US IPv6 Mandates

SIX STEPS TO ON-TIME, AFFORDABLE COMPLIANCE by tony hain

The US Office of Management and Budget (OMB) requires that government agencies enable their network infrastructures for IPv6 by June 2008. US federal agencies should not put off preparing to move to the next-generation routing protocol—but they don't need to panic, either.

The OMB mandate for IPv6 migration is unfunded, meaning that agencies will not get additional federal financial support for the transition, so it is prudent to incorporate IPv6 into your current network procurement, training, planning, and budgets to help meet the deployment deadline cost-effectively.

The 2008 conversion deadline applies only to network infrastructure equipment (backbone routers, switches, and hardware firewalls). This makes the required upgrade simpler and more affordable than if you had to IPv6-enable every component of your extended network.

Six basic steps will allow you to easily and cost-effectively meet the June 2008 deadline.

1 ASSESS YOUR IP INVENTORY. June 2006 was the OMB deadline for completing your inventory of existing IP-compliant devices and an analysis of the fiscal and operational impacts that the migration will have on your agency.

2 INCORPORATE IPV6 SUPPORT INTO PLANNED PRODUCT REPLACEMENTS. Folding IPv6 support into normal lifecycle product replacements that are already in your existing IT capital budget will help you avoid spending money explicitly to meet the IPv6 migration mandate.

3 EVALUATE YOUR EXISTING INFRASTRUCTURE HARDWARE FOR UPGRADEABILITY. If you haven't done so already, make it a priority to take inventory of network infrastructure hardware that is limited for use with IPv4.

The hardware most likely in question is routers at the very high and low ends. High-end routers, for example, often include acceleration hardware

that might be limited to IPv4 32-bit addresses. Less expensive routers might not have sufficient memory. In addition, firewall hardware and encryption accelerators are often IP version-specific, so check those.

4 USE TRANSITION TECHNOLOGIES WHERE IT MAKES ECONOMIC SENSE. Some of your backbone equipment that qualifies for the June 2008 deadline for IPv6 might already be scheduled for replacement soon thereafter—perhaps in late 2008 or early 2009. In cases where mandates allow discretion, consider temporarily using IPv6-in-IPv4 tunneling technology until the device's lifecycle has naturally ended.

Tunneling involves routing IPv6 packets over virtual paths in the backbone by encapsulating them in IPv4 network address headers. The IPv6 packets are delivered intact to end points, thus making the network appear as an IPv6 service. In this unfunded situation, the tunneling approach might minimize the need for short-term reprogramming while it maximizes your investment in existing backbone equipment.

5 ADD IPV6 TRAINING FUNDS INTO YOUR IT BUDGET AND PROCESSES NOW. Training is likely to represent a fairly high portion of your IPv6 migration costs, so it is advisable to integrate it into the IT training budget and process as soon as possible. Although the fundamentals are the same, it is wise to consider IPv6 as a completely different protocol than IPv4—one that will take your staff some time to learn. Among places to turn for IPv6 training are Cisco training partners, Native6 (www.native6.com), and Sunset Learning (www.sunsetlearning.com).

6 INCLUDE DETAILED IPV6 EXPECTATIONS ON ALL RFPs. Adding IPv6 as a criterion to all current and future IT requests for proposals (RFPs)—even beyond the core network—will help ensure that forthcoming OMB deadlines are met. This step should contain a detailed list of expected IPv6 features to help the industry avoid an impasse,

Further Reading

- IPv6 Requirements Memo to CIOs from Karen Evans, Administrator for Electronic Government and Information Technology, US OMB, August 2005
cisco.com/packet/183_4c1
- US OMB IPv6 Transition Strategy, June 2005
cisco.com/packet/183_4c2
- Cisco Systems Response to the US Department of Commerce IPv6 RFC
cisco.com/packet/183_4c3
- Cisco IPv6 Solutions
cisco.com/packet/183_4c4

whereby equipment vendors may not deliver the expected IPv6 feature set on certain devices due to lack of details on RFP checklists at the same time that agencies find themselves unable to deploy certain products because they do not deliver the appropriate IPv6 capabilities.

Although the IPv6 deployment process is fairly straightforward, it will take some time for staff training; equipment, operating system, and application updates; and procuring enhanced management tools. Integrating IPv6 procurement planning and training into your existing IT processes will help you meet your deadlines while avoiding unnecessary costs. ■

TONY HAIN is senior technical leader for IPv6 technologies in Cisco's Academic Research and Technology Initiatives group. He can be reached at ahain@cisco.com.

Gaining an Edge on the CCIE

NEW LAB SIMULATES RIGOROUS CCIE EXAM EXPERIENCE.



One of the most notoriously challenging certifications in the high-tech industry is Cisco's CCIE expert-level certification. An industry-leading certification for 13 years, the CCIE features an eight-hour exam where candidates are required to perform hands-on acts of networking wizardry that mimic real-life scenarios. But the reality is, very few CCIE candidates actually pass the exam the first time they take it.

"You can't simply cram for this exam over the weekend and hope to pass," says Kathe Saccenti, Cisco program manager for CCIE Assessor. "You have to be able to get in knee-deep to really prepare for something like this."

And now you can. The CCIE Assessor Lab, introduced last December, presents a series of technical scenarios and related questions as intense and realistic as the exam itself. In a four-hour session, students access live, remote Cisco equipment from their desktops and tackle problems designed to simulate the actual exam. Unlike the exam itself, the CCIE Assessor Lab provides detailed feedback and correc-



tions on each answer, and also suggests resources for areas that might require further study.

"In all respects, the experience is uncanny in its approximation to the real lab," says Systems Engineer Vernon Thaver, who took the CCIE Assessor Lab and subsequently passed his CCIE exam. "The information provided and the actual questions are very much what a CCIE candidate will experience at the various CCIE testing centers."

For more details about the CCIE Assessor Lab, visit cisco.com/web/learning/le3/ccie/preparation. **P**

POP QUIZ

Level: CCSP Security

- What design features enable a Cisco security appliance, such as the PIX Firewall, to outperform conventional application firewalls?
 - The Adaptive Security Algorithm
 - Super-packet filtering
 - Purpose-built, real-time operating environment
 - Hot standby proxy processing
 - Cut-through proxy support
- A Cisco security appliance can be configured to send syslog messages to all of the following except which one?
 - Console
 - Telnet session
 - Serial port
 - Syslog server
 - Answers a, b, c, and d are correct.
- Why is it difficult to penetrate a security appliance over UDP port 53?
 - The security appliance allows multiple outbound queries but randomizes the UDP sequence numbers.
 - The security appliance allows queries to go out to multiple DNS servers but drops all but the first response.
 - The security appliance allows responses only to outbound DNS queries.
 - All of the above
- Which command lets you create a network object group?
 - object-group network *group-id*
 - enable object-group network *group-id*
 - create network object-group
 - network object-group enable
- What is the size of the output for a MD5 hash?
 - There is no fixed size.
 - 256 bits
 - 255 bits
 - 128 bits
 - None of these answers are correct.

ANSWERS: SEE PAGE 82.

Source: CCSP SNPA Official Exam Certification Guide, 3rd Edition

readertips

THANK YOU FOR YOUR TIP

Each quarter we receive many more tips than we have space to include. 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.

SUBMIT A TIP

Help your fellow IT professionals by submitting your most ingenious technical tip to packet-editor@cisco.com. When submitting a tip, please tell us your name, company, city, and country. Tips may be edited for clarity and length.

Configuring a SOHO Router

If you are using a single IP address but you require more hosts to access the Internet, consider a small office/home office (SOHO) router (for example, a Cisco 800 Series) and use the following configuration:

```
!
ip dhcp pool Mypool
    network 10.10.10.0
    255.255.255.0
    default-router 10.10.10.1
!
interface Loopback0
ip address 10.80.80.1
```

Using Switch Clustering for Remote Configuration

When the console port of a switch is inaccessible and remote administration is not possible, it can be difficult to change an existing configuration. To reclaim the command-line interface (CLI) of the switch, use the switch clustering functionality available on most Cisco switches. A switch cluster consists of a command switch and up to 15 member switches. Use the Cisco Discovery Protocol to configure the command switch to discover the affected switch (candidate) and add it to the cluster. Then use `rcommand` to access the CLI of the new member switch as follows:

```
Command(config)# cluster enable CorpLAN
Command(config)# cluster discovery hop-count 5 (default 3)
Command(config)# exit
Command# show cluster candidates
MAC Address Name Device Type PortIf FEC Hops SN PortIf FEC
00d0.7961.c4c0 Affected WS-C2950-24 Fa0/5 1 0 Fa0/3
Command# conf t
Command(config)# cluster member mac-address 00d0.7961.c4c0
Command(config)# exit
Command# show cluster members
SN MAC Address Name PortIf FEC Hops SN PortIf FEC State
0 0002.4b29.2e00 Command 0 Up (Cmdr)
1 00d0.7961.c4c0 Affected Fa0/5 1 0 Fa0/3 Up
Command# rcommand 1
Affected>
You can add IP configurations and other appropriate settings (enable and vty passwords, HTTP server, SNMP, etc.) to the affected switch to enable the required remote management channels.
```

BENEDICT MUNYAO, *Netwise Associates Ltd, Nairobi, Kenya*

Editor's note: Because clustering will cause all the cluster members to have the same passwords, be sure you use good, strong passwords.

```

255.255.255.252
ip nat inside
ip tcp adjust-mss 1452
!
interface Ethernet0
ip address 10.10.10.1
255.255.255.0 secondary <-----IP
addresses assigned to clients
ip address 172.20.3.40
255.255.255.0 <-----the unique IP
address assigned
ip nat outside
ip policy route-map NAT
!
!
ip classless
ip route 0.0.0.0 0.0.0.0 172.20.3.1
!
ip nat inside source list NAT
interface Ethernet0 overload
ip access-list extended NAT
permit ip 10.10.10.0 0.0.0.255 any
!
access-list 101 permit ip
10.10.10.0 0.0.0.255 any
!
route-map NAT permit 10
match ip address 101
set ip next-hop 10.80.80.2
!
!
end

```

CARLO POGIORELLI, *Nuova Tin.it Srl,*
Rome, Italy

Troubleshooting

Debugging crypto on a Router

On a busy virtual private network (VPN) hub, much of the debug output is related to the normal management of other crypto sessions. It is not easy to follow a sequence of messages, because most output does not specify the crypto peer IP. In Cisco IOS Software Release 12.3(2)T, you can use the `debug crypto condition` command to filter debug output to a specific peer. You can use many criteria to

customize the output, including peer IP, SPI, connid, flowid, etc. This filter causes the router to show only the crypto messages for the peer you are troubleshooting. Syntax:

```

debug crypto condition [connid
integer engine-id integer] [flowid
integer engine-id integer] [fvrf
string] [ivrf string] [peer [group
string] [hostname string] [ipv4
ipaddress] [subnet subnet mask]
[username string]] [spi integer]
[reset]

```

Define a condition with a peer IP:

```

debug crypto condition peer ipv4
2.2.2.2

```

Identify which debug conditions are active:

```

sh crypto debug-condition

```

After you create the condition, start the debug:

```

debug crypto isakmp or debug crypto
ipsec, or debug crypto engine
When you are done, turn off the
debug condition:
no debug crypto condition peer ipv4
2.2.2.2

```

To stop the debug, you must turn off debug and remove the condition. If you do not, the condition will remain after you log off, which can make the next person's job very difficult. Always check for predefined conditions if you do not see expected output from a debug. For more information, refer to cisco.com/packet/183_4d1.

KEVIN MILLER, *Herman Miller Inc.,*
Zeeland, Michigan, USA

Resolving IP Telephony Voice-Mail Problems

Customers with IP telephony solutions (IP phones, Cisco CallManager, and Cisco Unity servers) sometimes experience intermittent problems when trying to access voice mail from outside the office. When they dial a number, there might be only one ring and then the call drops; at other times, there is dead

techtips

ASSIGN AN OUTSIDE INTERFACE IP ADDRESS.

This document provides a sample configuration for a Cisco PIX Firewall to dynamically obtain an IP address for the outside interface, using either Dynamic Host Control Protocol (DHCP) or Point-to-Point Protocol over Ethernet (PPPoE).
cisco.com/packet/183_4e1

TROUBLESHOOT AN SIP CALL BETWEEN TWO ENDPOINTS.

View a sample configuration of two fax machines that demonstrates how a Session Initiation Protocol (SIP) call takes place between two gateways. This document also describes the output of the debug ccsip messages command for troubleshooting SIP call failures.
cisco.com/packet/183_4e2

REPEAT A GREETING IN CISCO UNITY.

If an invalid key is pressed after the greeting plays in a Cisco Unity voice-mail system, the user is taken to a default greeting. This document describes how to configure the call handler to repeat the greeting each time the call is exposed to an error.
cisco.com/packet/183_4e3

CONFIGURE A WIRELESS LAN CONTROLLER AND LIGHTWEIGHT ACCESS POINT.

This document shows a basic configuration example of a lightweight access point that is connected to a Cisco Wireless LAN (WLAN) Controller through a Cisco Catalyst Switch.
cisco.com/packet/183_4e4

TROUBLESHOOT CISCO CATALYST 4500 SERIES SWITCHES.

Learn how to troubleshoot hardware problems and related issues that are common on Cisco Catalyst 4500/4000 series switches with Supervisor Engine II+, III, IV, and V modules.
cisco.com/packet/183_4e5

air. There can be many reasons for this problem: one is that after Cisco CallManager sends an ALERTING message to PSTN, it sends a NOTIFY message containing the display name of the ringing party, which in this case is voice mail. The problem is that Cisco Unity answers the call, then Cisco CallManager sends CONNECT to the PSTN and moves into the Connected State. The PSTN sends a STATUS message to Cisco CallManager, referencing the NOTIFY message it received. However, sending a status message to a Q.931 device that is in the connected state violates the protocol; therefore, Cisco CallManager disconnects the call. The problem is that Cisco Unity answers too quickly. For working calls, the NOTIFY/STATUS exchange takes place between ALERTING and CONNECT, and therefore does not violate the protocol. To solve this problem, uncheck the Display IE Delivery and/or

Redirecting IE Delivery Outbound boxes on the Gateway configuration page in Cisco CallManager. This requires you to reset the Media Gateway Control Protocol (MGCP) gateway on both Cisco CallManager and the router.

AHMED BAHER, *Equant,
Cairo, Egypt*

Editor's note: Status is a valid message in any call state. Q.931 TELECOMMUNICATION (05/98) "3.1.16 STATUS This message sent is by the user or the network in response to a STATUS ENQUIRY message or at any time during a call to report certain error conditions listed in 5.8. See Table 3-17."

The problem is a race condition, where Cisco CallManager sends a NOTIFY message and transitions state before the peer switch sends the STATUS message. The sta-

tus message from the peer switch includes "Call state," which indicates the call state when the peer switch received the message that generates the STATUS message. Cisco CallManager sends Notify, for example, in state IncomingCallProceeding, and then transitions to state Active. The peer switch sends STATUS message for the notify message indicating the call state when the offensive message was received was IncomingCallProceeding. Cisco CallManager sees a state mismatch and disconnects the call. The same problem can occur with the eServices/CRA/IPCC Express AutoAttendant script. Changing Display IE Delivery and/or Redirecting IE Delivery Outbound can affect other call flows. In the case of eServices/CRA/IPCC Express, add a 1-second delay before the accept step to work around the problem.

Continued on page 81

5 Ways to Ensure VoIP Reliability

ALL THE REDUNDANT HARDWARE IN THE WORLD WON'T COMPENSATE FOR POOR DESIGN. by ron trunk

IP telephony (IPT) systems and voice over IP (VoIP) have become commonplace applications in corporate IT departments. But along with the acceptance of IPT, come the demands that users place on the system, most of all that it be just as reliable as their old PBX system. The telecoms have done a great job of making telephones reliable—when you pick up the phone, you get a dial tone. Your users might tolerate occasional problems with their computers, but they'll have zero tolerance for problems with their phones. The reliability of IPT hardware and software has improved significantly, but one fundamental fact remains: IPT systems rely on your data infrastructure. The reliability of the data infrastructure is often the weakest link in the system. Unlike a traditional PBX, IPT systems have lots of “moving parts:” call servers, phones, gateways, routers, switches, TFTP servers, etc. A problem with any one of them can affect the availability of the overall system.

Reliability needs to be built in at the design stage. That means proper operational practices and procedures, not just lots of hardware. All the redundant hardware in the world will not compensate for poor design or poor practices. If you already have IPT on your network, there are some basic things you can do to improve its overall reliability. If you're still in the planning stages for IPT, these tips will help you start off on the right foot.

1 KEEP YOUR VOICE AND DATA SUBNETS SEPARATE. Using separate virtual LANs (VLANs) for voice and data—logically separating your voice and data traffic—is probably the most important thing you can do. By logically isolating your voice devices (phones, gateways) from your data devices (workstations, servers), you can decouple interactions between them. That can go a long way to improving your reliability—and security. Place your voice devices

and data devices on separate VLANs and allocate IP addresses for them out of separate (and summarizable) address blocks. Separate VLANs will also allow you to easily apply different quality of service (QoS) and security policies to voice than you do for data. There's no need for phones to talk to PCs or vice versa. By preventing traffic from flowing between your voice and data subnets, you can eliminate potential security vulnerabilities, misconfigurations, and operator errors. The one exception can be management workstations to administer your system. The same rule of thumb applies: place those workstations in a separate VLAN and only allow that VLAN to access voice subnets.



Implementing access control lists (ACLs) or other filter mechanisms to isolate voice and data traffic is much easier when you've allocated addresses for your voice subnets out of a separate (and summarizable) address block from your data subnets. With separate address spaces, the access list often can be simplified to a single line. If you haven't used

separate address spaces, you might seriously consider renumbering. If you use registered addresses for your data network, be sure your IPT system uses private (RFC 1918) addresses. There's no reason to have your phones use globally unique addresses.

And, of course, call servers should also be placed in their own separate VLAN, which allows you to filter traffic to and from the servers. Because the call servers are the heart of your IPT system, you need to protect them from unexpected events. Apply access lists to only allow the necessary traffic (typically call setup and management traffic) to reach the servers. Better still, if your budget allows, place a stateful firewall between the call servers and the rest of the network to prevent unexpected traffic.

2 APPLY QoS CONSISTENTLY. Don't count on high-bandwidth links to eliminate the need for QoS.

DESIGN FOR SUCCESS
Reliability needs to be built in at the design stage.

It's always important to develop a QoS policy and implement it consistently throughout your network—not just on a WAN link or two. Take a holistic, system-wide approach and apply QoS end to end. Tag voice traffic as soon as it enters the network and apply policies on every interface. Remember that networks frequently change, and a low-used link can suddenly become a highly-used one. The lack of QoS is often the cause of intermittent voice quality problems as network use varies during the course of a day. As you develop your QoS policy, consider other delay-sensitive applications such as videoconferencing or streaming media. Plan for future growth, and you won't be making major changes at the last minute.

3 DON'T FORGET THE BACKUP POWER. You no doubt have UPS systems to protect your data servers. But have you considered how a power failure would affect your IPT system? For safety reasons, if nothing else, you'll want the phones to work if the power goes out, typically for at least 30 minutes. Carefully consider all the devices that need to remain powered for the IPT system to remain functional: call servers, routers, switches (don't forget your wiring closets), power injectors, gateways, etc. This might mean a significant increase in the size and number of UPS systems. But it's an insurance policy you won't want to be without.

4 KEEP YOUR DIALING PLAN SIMPLE. Organizations often renumber their phones when they implement IPT. If you do, follow the same principle as with IP addresses: keep it simple. Allocate numbers in summarizable blocks. A simple dial plan has two benefits: it's easy to understand and easy to troubleshoot. Design your dial plan so that calls can always find a way out to the PSTN. In case of failure, your dial plan should route calls to any available gateway, even if that means incurring long-distance charges. It's better to pay a little more than not have your calls go through.

5 DOCUMENT BOTH YOUR NETWORK AND IPT SYSTEM. The simplest, yet most effective thing you can do to improve reliability and uptime is to have good documentation. Remember that reliability is not only a matter of having redundant components, but also being able to make repairs quickly when things break. With a well-documented system, you can much more easily diagnose problems when they occur. The faster you can repair things, the happier your users will be. ■

RON TRUNK, CCIE, CISSP, is a senior consultant at Chesapeake Net-Craftsmen, LLC (www.netcraftsmen.net), which delivers high availability solutions for network design, operating systems, applications, security, storage, and IP telephony.

Securing the Edge

SOLVING THE TRUST ISSUE WITH BGP by steven moore and russ white

S

election of the best routing protocol for an enterprise network typically driven by specific business requirements, including nimble response to change, quick convergence, open communication relationships (trust), and minimal configuration. However, connections outside the enterprise network have a completely different set of considerations. / When you connect to networks outside your administrative control, security and policy increase in importance, while convergence speed decreases. You will want to choose a different routing protocol—even a different type of routing protocol—to carry routing information. / Border Gateway Protocol (BGP) offers a wide array of tools for enterprise networks and is an ideal choice for an externally facing protocol at the network edge. Here's why.

Network Edge Types By far the most common type of external connection that comes to mind is the Internet—that “network of networks” which is almost magical in its power to reach around the world. But there are several other external connections to consider as well, for example, an extranet connection to a partnering company, whether a supplier, customer, or some other type of financial partner.

A less common type of external connection might be the connection of a business unit to a corporate backbone in a large, diversified company. The corporate backbone acts as a sort of service provider within the company, connecting units together, and to commonly shared services.

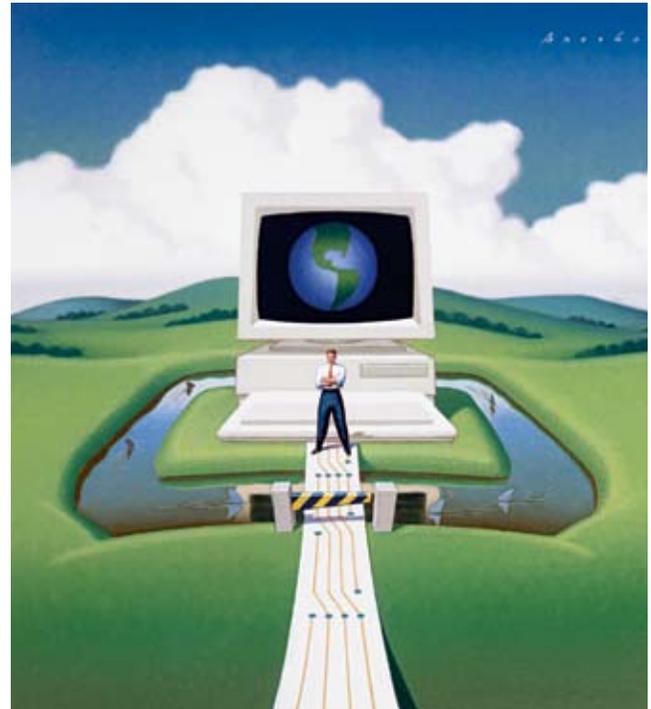
The common thread in all three of these cases? When connecting to networks outside your administrative control, you must solve the trust issue. Can you trust the routing information you are receiving across this connection? Do you need to consider policies? While you have probably put a lot of thought into protecting your data plane (your traffic), you might not have thought about protecting your control plane, or your routing system.

Routing Problems at the Edge The following situations involving incorrect routing information and flapping routing information can have a negative impact on your internal routing.

INCORRECT ROUTING INFORMATION

Consider two hypothetical companies: one called BigShoes, which uses the IP address 10.1.0.0/16, and one called Medium-Socks, which uses 10.2.0.0/16 (see the figure on page 25).

Continued on page 25



WHEN CONNECTING
to networks outside your administrative control, you must solve the trust issue.

The two companies have recently formed a partnership to sell shoes and socks at the same outlet stores. BigShoes also partners with other companies, such as SmallFeet, and all of these partners are connected using redistribution between their internal Interior Gateway Protocols (IGPs). Potential problems in this network include the following:

- SmallFeet injects 10.2.1.0/24 into the BigShoes network; the route leaks into the MediumSocks network, possibly causing the best route to some destinations to be through BigShoes, rather than to the local resource.

- BigShoes learns a route from MediumSocks, say 10.2.2.0/24, and advertises it into the Internet. The edge routers at MediumSocks learn this route from the Internet and believe the best path to 10.2.2.0/24 is through its ISP, disrupting its internal routing.

- BigShoes misconfigures its routers, injecting the entire Internet routing table into the MediumSocks IGP. This overwhelms the MediumSocks routers, causing a major outage in its network.

There is no easy way for MediumSocks to defend itself against these types of problems, whether they are malicious attacks or unintentional misconfigurations, using an IGP.

FLAPPING ROUTING INFORMATION

Suppose MediumSocks' network engineers have rolled out voice over IP (VoIP) throughout the network, and have tuned the network to provide the fastest possible convergence, including fast timers and exponential backoff.

If BigShoes injects routes into the MediumSocks network that change on a regular basis, what will the result be in the MediumSocks network? The IGP will interpret the constant changes as an indicator of network instability and will back off the fast convergence timers. The result is that constant changes in the BigShoes network have an impact on the convergence time, and thus the performance, of the MediumSocks network.

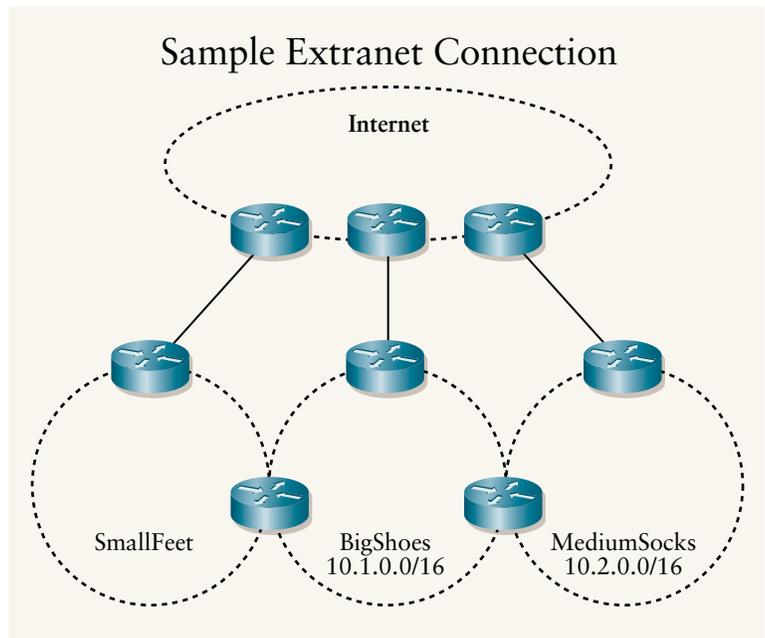
How can MediumSocks protect itself against this type of problem? By using BGP at the edge.

BGP Versus IGP

FROM ITS INCEPTION, BGP has been designed to address this specific routing problem: connecting networks together to create internetworks. While the IGP provides a toolset geared towards internal information exchange, BGP better addresses the unique issues that exist when crossing the boundaries of trusted internal relationships.

Routing Solutions to Secure the Edge

BGP HAS A PLETHORA OF knobs and tunable options, which allow you to build a secure, well-designed connection to



TYPICAL SET OF NETWORK EDGES, including an extranet and connections to the global Internet.

a network outside of your administrative control. The key functions that BGP offers include policy, protection, and peer-based management.

POLICY

Several BGP policies can prevent problems for MediumSocks. Let's return, for a moment, to the situation where 10.1.1.0/24 is being leaked through BigShoes from SmallFeet, causing misdirected traffic. While a simple prefix filter might work in this situation, BGP provides many other complementary tools.

- Use a prefix list to limit the injection of longer prefix (more specific) routing information into the MediumSocks network. The local policy across the partner peering sessions could only permit prefixes with a length of /17 or less, preventing routes more specific from being accepted at the network edge.

- Use an AS PATH filter list to prevent routes not originated within a peer's network from being injected into the MediumSocks network at the edge. In this case, MediumSocks could filter so only routes originating within BigShoes will be accepted at the edge.

You can attach specific communities to BGP routes, as well, indicating what you want the "other" network administrator to do with the route. For example, if MediumSocks does not want its networks to be visible to SmallFeet through BigShoes, it can set the NO_EXPORT community on the routes it advertises to BigShoes. Any route marked with NO_EXPORT should not be readvertised outside the routing domain.

Finally, always filter out bogon routes at the edge of your network. Bogon routes are known bad routes within the Internet. For example, filter out all private networks along

their Internet edge, although some private networks might be allowed, through prior arrangement, at private peering points. Address space reserved for research projects, or multicast use, and address space known not to be allocated to anyone, are also bogons, and generally should not pass through administrative domain boundaries.

The following sample configurations that provide suggestions on how MediumSocks can use BGP to protect its internal routing infrastructure in these areas, see cisco.com/packet/183_5a1. These examples assume that BigShoes is AS65000 and MediumSocks is AS65001.

BGP Router Peering with BigShoes

```
router bgp 65001
 neighbor <bigshoes> remote-as 65000
 neighbor <bigshoes> route-map filter-partner-in in
 /* inbound route filter, described below in the
 route-map
 /* filter-partner-in configuration
 neighbor <bigshoes> route-map filter-partner-out out
 /* outbound route filter, described below in the
 route-
 /* map filter-partner-out configuration
 ....
 route-map filter-partner-in permit 10
 match ip address prefix-list partner-routes-in
 /* any routes permitted by the prefix list partner-
 /* routes-in
 match as-path 1
 /* any routes permitted by the as path access-list 1
 will
 /* be accepted
 ....
 route-map filter-partner-out permit 10
 set community no-export
 /* prevents BigShoes from readvertising routes
 learned
 /* from MediumSocks, and from transiting traffic to
 /* MediumSocks
 ....
 ip prefix-list partner-routes-in seq 10 deny
 192.168.0.0/16 ge 15
 /* denies bogon routes in the range 192.168.0.0/16
 ip prefix-list partner-routes-in seq 20 deny x.x.x.x/xx
 /* deny other bogon routes here
 ip prefix-list partner-routes-in seq 10 permit
 0.0.0.0/0 le 18
 /* permit any routes with a prefix length less than
 /17
 /* prevents longer prefix routes from causing local
```

```
/* routing problems
!
ip as-path access-list 1 permit ^65000$
 /* denies any routes originated outside the peering
 AS
 /* including BigShoes' partners and routes BigShoes
 is
 /* learning from an ISP
```

BGP Router Peering with the Internet Service Provider

```
router bgp 65001
 neighbor <ISP> remote-as <ISP AS>
 neighbor <ISP> prefix-list isp-routes-in in
 neighbor <ISP> route-map filter-isp-out out
 ....
 ip prefix-list isp-routes-in seq 10 permit x.x.x.x/xx
 /* deny bogon routes here
 ....
 route-map filter-isp-out permit 10
 match as-path 2
 ....
 as-path access-list 2 permit ^$
 /* permits only routes originating within
 MediumSocks, so
 /* MediumSocks doesn't transit to BigShoes
```

PROTECTION

BGP also provides protection in the case where a flapping route in the BigShoes network impacts the convergence speed and stability of the MediumSocks network, and in the case of BigShoes injecting a full routing table along its peering edge.

The first protection is offered by a BGP feature called *Route Flap Dampening*. Route Flap Dampening works by applying a penalty to a route each time the route flaps, or changes. If the penalty applied to the route rises above a specified number, the route is ignored for a short time period. Route Flap Dampening is commonly used by Internet service providers to protect against constantly changing routes.

Route Flap Dampening parameters are considered aggressive if they dampen a route after a small number of changes in a long period of time, and not aggressive if they only dampen a route after many changes in a short time period.

JOIN THE DISCUSSION

Ask your peers and Cisco experts questions, or share your own knowledge about BGP and other routing protocols, at the Cisco Networking Professionals Connection "Network Infrastructure" forum: cisco.com/discuss/infrastructure

BGP is an ideal choice for an externally facing protocol at the network edge.

```
router bgp 65000
....
bgp dampening
bgp dampening 1000 2 2000 750 60
```

You can also dampen flaps on different prefixes at different rates by using a route-map. For example, if reachability to 10.1.1.0/24 is considered critical, and reachability to 10.1.2.0/24 is not, you can dampen 10.1.2.0/24 more aggressively. To determine what routes have been dampened, use the **show ip bgp dampened-paths** command.

Generally, you want Route Flap Dampening to be fairly aggressive in private peering relationships; there is little reason to accept a large number of route changes over short periods of time in private peering.

For MediumSocks to protect itself against BigShoes flooding its network with too many routes and possibly causing a routing failure, the MediumSocks network engineers can configure a route count limit, as follows:

```
router bgp 65000
neighbor <bigshoes>
neighbor <big shoes> maximum-
prefix 100 restart 30
```

This configuration causes BGP to start its session when more than 100 routes are received and allows the session to restart after 30 seconds of idle time.

Another BGP protection is to prevent a peering router from advertising routes that do not include the peering AS number. For example, if someone with access to the BigShoes network attempted to advertise a route, 10.1.1.0/24 with a spoofed originating AS, the MediumSocks BGP speaker could reject this advertisement, because it does not contain the BigShoes AS number. This feature, Enforce the First Autonomous System Path, is enabled in more recent versions of Cisco IOS Software.

This BGP feature allows you to control the number of routes you accept across a given peering session. BGP can react to additional routes either by providing a warning that the router is receiving too many routes, or it can actually close the BGP session down, providing absolute protection for the network.

Beyond these types of protections, BGP also provides protection against more direct attacks on your network. BGP was designed to operate in lower trust environments, where the

link to the outside network may not be secured, or able to be secured, in an easy way. For instance, BGP can easily hop over firewalls, because it uses a unicast TCP session to transfer routing data. This allows you to use a firewall to control the flow of data traffic through the network, and use BGP to control the exchange of routing information.

Routers that run BGP are protected by several BGP mechanisms, such as the Generic Time-to-Live Security Mechanism (GTSM), described in RFC 3682 at cisco.com/packet/183_5a2.

To configure this protection, use the following:

```
router bgp 65000
neighbor <bigshoes> incoming-++1
<minimum ++1 to accept>
```

In summary, never use an IGP to receive or transmit live routing data between two routing domains. Instead, use BGP and rely on the protections and trust level of a protocol designed to provide the types of protections you need. ■

STEVE MOORE, CCIE NO. 4927, is an engineer with the IP Routing Scalability, Performance and Integration Testing team, within a part of the Network Software and Systems Technology Group at Cisco. He has been with Cisco for 10 years, with expertise in routing protocols, WAN technologies, and optical networking. He can be reached at smoore@cisco.com.

RUSS WHITE, CCIE NO. 2635, is a frequent contributor to *Packet* and *IP Journal*, as well as a regular speaker at Cisco Networkers. The co-author of six books on routing protocols and routed network design, he is currently at work on a new Cisco Press book on Cisco Express Forwarding. He is a technical lead in the Routing Protocols Design and Architecture team within Network Software and Systems at Cisco, and can be reached at riw@cisco.com.

Further Reading

- Best Practices for Securing Routing Protocols
cisco.com/packet/183_5a3

Managing MPLS

AN OVERVIEW OF MPLS OAM TOOLS, TECHNIQUES,
AND STANDARDS by monique morrow and thomas nadeau

As carriers and service providers worldwide converge disparate networks and the services offered over those networks onto a common Multiprotocol Label Switching (MPLS)-based infrastructure, MPLS operations, administration, and maintenance (OAM) functionality is a critical infrastructure component for enabling this transition. In particular, it affords operators the insight into how their network is functioning (or not), allowing them to provide different service-level agreement (SLA) guarantees, service assurance, quality of service (QoS) assurance, predictable downtime management, and overall internetworking service management.

Providers can further use OAM functionality to offer premium managed services based on enhanced SLAs. To realize these benefits, network operators need the ability to reliably conduct SLA testing, detect MPLS control-plane and user-plane defects, and check MPLS forwarding path integrity in real time and in a scalable manner.

Cisco is demonstrating leadership in several MPLS OAM areas: Label Switched Path (LSP) ping and trace for Label Distribution Protocol (LDP) and traffic engineering (TE), IETF Virtual Circuit Connection Verification (VCCV), ITU-T Y.17fw, MPLS MIBs, and others. This article focuses on how these OAM mechanisms help operators manage and troubleshoot MPLS networks.

First, What Is MPLS OAM?

MPLS OAM TOOLS AND TECHNIQUES apply to all applications of MPLS (see Figure 1 at cisco.com/packet/183_5b1 for an overview of MPLS services and transport network management). In addition to maintaining core integrity, the primary objective of an MPLS OAM strategy is to reduce costs by minimizing service interruptions. Based on this, MPLS OAM addresses the following requirements:

- Determining consistency between MPLS control and data planes
- Detection, diagnosis, and localization of broken LSPs

- LSP trace capability
- Support for equal-cost multipath (ECMP) constructs
- Backward compatibility and support for existing infrastructure as well as new applications
- Support by the OAM mechanism for SLA measurement

While MPLS provides native resiliency facilities such as Interior Gateway Protocol (IGP) fast convergence, Fast Re-Route (FRR) fault recovery, and LDP graceful restart, these mechanisms cannot detect all faults, nor can they diagnose faults and their locations within the MPLS network. Furthermore, when the data plane is not in sync with the control plane, these mechanisms cannot recover. This is sometimes referred to as a data plane “black hole,” where traffic traversing an LSP continues until it is either mis-

MPLS OAM affords carriers and service providers the insight into how their network is functioning (or not).

directed or thrown away at the point of malfunction. Such mechanisms are required to operate seamlessly within a MPLS network.

Now let’s look at some of the OAM tools and techniques for use in MPLS networks. These mechanisms are available on Cisco routers.

LSP Ping/Trace

THE MPLS PING/TRACE TOOL is modeled after the IP ping and traceroute paradigm: ping (ICMP echo request) and trace (UDP packets with incremental time-to-live, or TTL, values) are used for connectivity verifications.

LSP ping and trace functionality diagnoses and

localizes LSP failures. Before this functionality was available, operators had to use tedious hop-by-hop show commands to isolate an LSP failure. With LSP ping/trace, operators can also glean more information about the nature of a failure condition (for example, MTU mismatch conditions often are reported as network degradation problems).

LSP ping has been standardized within the IETF MPLS Working Group as RFC 4379. LSP ping tests the connectivity integrity of an LSP by sending test messages known as echo requests that are encapsulated in precisely the same manner as the data traffic transmitted over the LSP under test. The only difference is that the packet payload contains special information allowing intervening Label Switch Routers (LSRs), or the ultimate LSR, to process the packet. These messages are correspondingly replied to using an MPLS echo reply message.

When the LSP ping echo function is invoked on an LSR, the originating LSR sends an MPLS echo request to the target LSR. When testing an LSP, the first task is to look up the Forwarding Equivalence Class (FEC) to label stack mapping for the LSP under test. This provides the information necessary to encapsulate the remainder of the MPLS echo request packet so that it is handled in a manner consistent with data sent on the LSP. See Figure 2 at cisco.com/packet/183_5b1 for an example of the LSP ping echo function.

Now, let's take a look at diagnosing problems in pseudowire tunnels, and the role of VCCV.

Virtual Circuit Connection Verification

PSEUDOWIRES CARRY EMULATED SERVICES such as Ethernet or Frame Relay over MPLS or IP networks. Cisco's original pseudowire feature is referred to as Any Transport over MPLS (AToM). However, pseudowires have since been expanded in Cisco routers to be capable of running over L2TP transports as well. VCCV establishes an in-band control channel between the pseudowire endpoints (or "Martini circuit" as it is sometimes called) and is used to convey, among other things, OAM information between the endpoint provider edge (PE) routers. In a nutshell, VCCV is a tool that allows operators to perform a connectivity verification operation on a pseudowire (see Figure 3 at cisco.com/packet/183_5b1). Unlike MPLS LSP ping alone, VCCV provides the capability of checking one specific pseudowire.

VCCV is being standardized

within the IETF Pseudowire Emulation Edge-to-Edge (PWE3) Working Group, and is designed to run over any supported transport technology, although at present, MPLS is currently only supported.

When VCCV is first run, it transmits a "capability advertisement" to the remote PE router via an extended interface parameter TLV contained in the LDP setup message. When signaling capabilities, an LSR indicates which of the various connectivity check types it will support being sent from the peer. The router should indicate all of the methods it supports to promote the highest possibility of the peer supporting one of the modes.

VCCV can support multiple types of payloads and/or operations, but the latest draft of the specification indicates that only one method can be used after it is successfully transmitted. For example, if the LSP ping mode is chosen and that type of payload is transmitted and replied to, the sender must not send any other types until the pseudowire is re-signaled. This is done to simplify the state machine used to process the requests by requiring only a single mode of operation once started.

The following enumeration lists each of the connectivity check (CC) types: MPLS LSP ping and Bidirectional Forwarding Detection (BFD):

```
0x00 None.
0x01 ICMP Ping
0x02 LSP Ping
0x04 BFD for PW Fault Detection only
0x08 BFD for PW Fault Detection and AC/PW Fault
      Status Signaling
```

Other Proposed MPLS OAM Mechanisms: Y.1711 and Y.17fw

THE ITU-T Y.1711 RECOMMENDATION is based on connectivity verification packet flows, which are inserted in the

Figures List

The following figures referenced in this article are available at Packet Online, cisco.com/packet/183_5b1:

- Figure 1 "MPLS Services and Transport Network Management"
- Figure 2 "LSP Ping Example for MPLS OAM Troubleshooting"
- Figure 3 "Virtual Circuit Connection Verification"
- Figure 4 "The Cisco Auto IP SLA Feature"
- Figure 5 "Cisco MPLS Diagnostics Expert Tool"

network at the LSP headend. These packets are checked at the tail end. If a faulty condition is detected at the LSP, notifications are sent back to the headend. Each LSP requires a state machine at its terminating LSRs (both headend and tail-end LSR), which keep track of the default condition status. OAM packets are uniquely identified via a special MPLS label (14).

Y.1711 has limitations. The OAM label can break the commonly used ECMP algorithm, resulting in false positives and limited coverage of the ECMP tree. Y.1711 does not apply to networks using the very common penultimate hop popping (PHP) feature of MPLS. These two conditions constrain Y.1711 to point-to-point constructs such as TE tunnels. In addition, Y.1711 implementations will require hardware changes to achieve full levels of functionality due to heavy processing needed for sequence numbers and packet volume. And a probe LSP every 1 second (another Y.1711 requirement) will place a significant load on the network.

To date, there is no industry traction for Y.1711.

On the other hand, the ITU-T Y.17fw recommendation provides a framework for MPLS network administration and maintenance and aligns with the work done in the IETF, such as LSP ping, LSP trace, BFD, and VCCV. Cisco is co-editor of the Y.17fw recommendation.

Cisco Auto IP SLA

CISCO HAS ENHANCED the MPLS OAM standards-based tools by wrapping them with automation of redundant tasks, such as the trigger of periodic connectivity tests as well as triggering actions based on detected failures. In particular, the Cisco Auto IP SLA feature integrates the power of IP SLA probe scheduling and optimization with the LSP ping/trace functions and a PE next-hop discovery function to automatically verify all equal-cost paths between two or more PE routers supporting VRFs within the same VPN. See Figure 4 at cisco.com/packet/183_5b1 for an example of how the Auto IP SLA feature works.

Cisco MPLS Diagnostics Expert

TO THIS POINT, we have discussed MPLS OAM mechanisms on Cisco routers. But there are additional functions that are useful for operators that exist “outside of the box.” The Cisco MPLS Diagnostics Expert (MDE) tool integrates the MPLS OAM mechanisms for Layer 3 VPNs that exist on Cisco devices with intelligent post-failure detection troubleshooting algorithms.

Before MDE, operators had to either manually take over from the automated embedded tools or trigger scripts based on

Standards Status

Cisco is involved in several OAM standards areas: LSP ping and trace for LDP and TE, VCCV, and Y.17fw, among others. Following is the status of some of these efforts:

MPLS Ping/Trace	RFC 4379
VCCV	Last call
Y.17fw	Pending consent 2006 (New Rec y.1714 as of July 2006)

the probe failure notification. MDE was designed to be triggered based on the probe failure indication from a router, and then engage troubleshooting algorithms used in dozens of well-known troubleshooting scenarios. This data was gleaned from the Cisco Technical Assistance Center (TAC), as well as from operational staff at major service providers deploying MPLS. The algorithms are periodically enhanced and updated to improve MDE’s troubleshooting techniques. Figure 5 at cisco.com/packet/183_5b1 depicts the steps required to troubleshoot failed LSPs with and without the use of the Cisco MDE tool. **p**

The authors would like to thank Laure Andrieux, Stephen Speirs, and Hari Rakotoranto for their contributions to this article.

MONIQUE MORROW is a Distinguished Consulting Engineer at Cisco with more than 20 years experience in IP internetworking that includes design and implementation of complex customer projects, and service development for service providers. She can be reached at mmorrow@cisco.com.

THOMAS NADEAU is a principal engineer at Cisco responsible for operations and management architecture and network management of MPLS-related components. He can be reached at tnadeau@cisco.com.

Further Reading

- MPLS OAM Tools for Troubleshooting MPLS Networks
cisco.com/packet/183_5b2

I WANT MY IPTV!

IN FACT, I WANT
EVERYTHING,
ON ANYTHING.



ADOLESCENTS AND YOUNG ADULTS—"Generation Y," as the US calls this current throng—typically do their homework with MP3 player headphones in their ears. They game over the Internet, sometimes linking two game boxes and TV sets together, and talk to friends and opponents across town or on other continents. They instant message over PCs and cell phones and download music over the Internet. They capture films from their TV sets and burn them onto DVDs.

by *janet kreiling*

DWIGHT ESCHLIMAN



These Generation Y offspring are having a very different experience with communications and entertainment—and different expectations—from their parents, or even from Generation X, about ten years older. Very soon, about the time they begin earning salaries, Generation Y members are going to want their IPTV, or TV over Internet Protocol, along with IP everything else. A good many early adopters from Generation X already do.

“We’ve already seen a transition from passive watching of programs according to network schedules to people picking when and what they want to watch with video on demand [VOD],” says Pankaj Gupta, senior manager for service provider marketing at Cisco. “The next transition, which is already occurring, is to interactive services—the end user socializes with others through the TV or Internet—gaming with others online and participating in community networks such as MySpace.com and Wikipedia, for example. The final step is consumer empowerment, when individuals create and remix content.”

The Connected Home and Glass to Glass

Just in time, the technology is ready to give Generation Y and other early adopters what they want. IPTV might well be the first IP service to make a big splash. Some providers are already offering it over fiber to the home (FTTH) and even high-definition TV (HDTV) over IP (see related article, page 63). But once high-bandwidth IP gets into the home, it’s likely to penetrate everywhere, tying all of the home’s entertainment and communications activities together onto one network.

Consumers will have a *Connected Home* with an all-IP network throughout that allows them to get any information, from any device, in any room (see adjacent sidebar). And service providers will have an end-to-end IP network from head-end/central office to the home.

According to market research firm In-Stat, the networked entertainment market had reached US\$3.9 billion in revenue worldwide at the end of 2004 and is expected to grow to \$16.1 billion by 2009. In-Stat also estimates that networked entertainment devices

will be used in more than 38 percent of home networks within the same period.

Because everything runs over the same infrastructure, across country as well as in the home, all the communications streams—voice, video, and data, even wireless—are coordinated with each other and follow the same priorities. Video gets the right quality of service (QoS), as does voice. So do file transfers, music or video, whatever the end user orders. This type of network is increasingly being referred to as *glass to glass*: from the glass of the video camera to the glass of the TV or PC screen.

From glass to glass, Cisco equipment is everywhere: in the provider’s headend and in the home itself through Cisco’s recent acquisition of Scientific Atlanta and the 2003 acquisition of Linksys.

First, IPTV

As far as wanting your IPTV, what’s possible now? Most of it, starting with the network infrastructure to get multi-megabytes of bandwidth to the home. HDTV, combined with voice and data over one line into the home, requires some 20 Mbit/s; FTTH isn’t a prerequisite, though. Wideband for DOCSIS and

What’s Possible in the Connected Home?

In the Connected Home, end users can get any information, on any device, in any room. Or even on the outdoor patio for that matter. Among the things they could do:

- Order a VOD on the family’s main TV set and finish watching it on a set in the kitchen or in the bedroom
- Sort through and display a batch of photos stored on the PC in glorious color on the large TV screen in the living room
- Stream music downloaded onto a PC to the home stereo surround-sound system
- Check on the baby sleeping upstairs, or even on a second home hundreds of miles away, from any Web browser
- Get a caller ID announcement of a call from a daughter away at college, pause the show in progress, whether it’s a movie or scheduled program, take the call on the TV—and if she’s using a video phone, hold a video-conference with everybody in the family. And then go back and pick up the TV program where it was paused.

For a visual representation of the Connected Home, see page 43.

DSL2 or VDSL can deliver up to 50 Mbit/s, and service providers on several continents are already using these technologies.

What about other networked entertainment in the home? Just about any household with a router linking a couple of computers and peripherals is already bringing an Ethernet signal over IP into the residence. Building up the home network requires, probably, a heftier router along with adapters to enable phones, game boxes, and other devices to handle IP. Scientific Atlanta and Linksys have the requisite devices now.

The home network also requires a means of distributing signals throughout the house, either wireless or wired. Wireless home networks are already popular, and, while some people run Category 5 twisted pair cables over the floors between rooms (one hopes temporarily), more appropriate wiring networks can be built using coax and twisted pair installed in the walls and, most recently, the near-ubiquitous electrical wiring. The latter is still in the early-adopter stage and can be noisy, but the technology is improving.

In the near future, Scientific Atlanta's set-top boxes could cooperate with Linksys routers to distribute video content throughout the home, not just to the TV sets they top but to any device with a display that can take an IP signal—what Jim Strothmann, director of product strategy and management, North American cable video products, at Scientific Atlanta, calls “many services to many screens.”

As Strothmann explains, “Caller ID from the phone can show up on the TV set. So can digital photos from a PC. Or they can be sent to a portable video player somewhere else in the home.” He sees the set-top box playing a role in social networking, too: “Most TV programs have Websites. You could click on an icon on the TV set and be linked to the Website, perhaps with a discussion group.”

SCIENTIFIC ATLANTA, which already sells a device with an integrated DVD burner, announced in April 2006 that it is pointing its boxes toward the home IP network and a broader user experience. New set-top boxes will have the processing power and connectivity to share content within the home, including with devices purchased from consumer electronics stores, such as portable media players and PCs. In addition, its new Explorer 940 Compact Digital Only Interactive Set-Top Box can receive introductory digital video service and support pay-per-view and VOD services.

Thinking Outside the [Set-Top] Box

Linksys offers a wide range of adapters to attach popular home devices to the broadband service. Its media adapters connect a TV and stereo to the home network, moving video, music, and

photos throughout the home, according to Chris Dobrec, director of business development at Linksys. Its analog telephony adapters give regular telephones an IP connection so they can make use of voice-over-IP (VoIP) services. A Wi-Fi phone extends VoIP both in the home and at public hotspots. A game adapter links popular game consoles to the Internet wirelessly—saving that Cat 5 cable over the floor. A wireless print server lets everyone in the family print from any PC in the home.

Two additional Linksys products especially exemplify the Connected Home: the NSLU2 Network Storage Link and the WVC54GC Wireless-G Compact Internet Video Camera. The camera can be placed anywhere for surveillance: in the baby's room, by the front door or another entry point, or in a second home, small business, or even the backyard. The



Network Storage Link lets users attach an inexpensive USB hard drive to the home network. Everyone on the network can share movies, pictures, music, or other digital content. The hard drive is also useful for backing up data from PCs in the home.

At the center of this Connected Home is the Linksys router. Already an IP-ready device, the router will gain in capacity and ports, says Dobrec. “We see four categories of applications running off the router—personal computing, communications, entertainment, and home controls such as remote control of heating or lighting products.” Even now, he adds, devices in all four categories are being introduced that communicate via Ethernet or IP rather than proprietary systems.

Linksys recently introduced a line of routers and products based on the draft 802.11n standard. This new generation of wireless-N devices boasts improved range and dramatically faster speeds, enabling video, voice, data, and other content to move wirelessly—all at the same time.

The User Experience Starts with the TV

Of all the devices deployed around the home, the TV is the one with the most impact and the one to which consumers will most often look for the Connected Home's advantages, simply because it has a big screen that can be shared with others.

As Gupta points out, “IP video is key. VOD, digital video recording, video streaming, gaming, videoconferencing and videophones, buying movies online and burning them to DVDs all rely on video.”

Consumers want video of all types to be reliable and simple to use, with quality equal to what they've been getting. This means a reliable network infrastructure. Studies show that consumers will tolerate no more than one screen artifact per two-hour movie. That one artifact can be caused by the loss of just one packet. They want channel changes to take place

Further Reading

- Introduction to the Connected Home
cisco.com/packet/183_6a1
- “The Connected Life: Enabling the Transition from Service Provider to Experience Provider”
cisco.com/packet/183_6a2

instantaneously, and they want to be sure that if they have paid for a program, QoS will be maintained throughout its length. They want variety, in both local and national programming.

THE SERVICE PROVIDER'S network must have capabilities such as instantaneous channel changing, easy insertion of local content and ads, interactivity, video admission control, and security for the network and for content. Preferably, it operates at Layer 3 rather than Layer 2, Gupta says. “Layer 3 network intelligence has advantages that improve security and increase capacity and resilience. For example, a packet need travel only from point A to point B, rather than all the way around a ring, saving considerably on bandwidth.”

The Experience Provider

While the Connected Home will enable consumers to use their networks to do many things, whether they do or not comes down to whether they have a good experience. To this end, service providers, who have traditionally been known as cable companies or telcos, now need to think of themselves, as *experience providers* (see adjacent sidebar).

The service provider experience matters, too, of course. IPTV and all the other features of the Connected Home must be delivered reliably and cost effectively. Service providers want a proven, scalable, end-to-end infrastructure, timely deployment, and experienced partners that can imagine and deliver

the next moves in creating the winning experience for end users. In addition to Scientific Atlanta and Linksys, in September 2005 Cisco also acquired KiSS Technology A/S in Hørsholm, Denmark, a leading maker of home networked entertainment devices that include DVD players and recorders. Additionally, Cisco along with Intel Corporation has invested in content provider MovieBeam, Inc. (see page 4 for more on MovieBeam).

“Companies such as Google and Yahoo are already fronting challenges to the traditional model of a service provider,” says Wayne Cullen, senior manager in the Service Provider Routing and Switching Group at Cisco, “by providing content in addition to their Internet services. Traditional communications providers will need to put together entirely new models of what business they’re in and how they do it. The Connected Home, starting with IPTV, is a great place to begin.” ■

Differentiating Your IPTV Service

Three primary factors will rule the consumer's choice of service provider for IPTV and the Connected Home: the degree of empowerment, quality of service, and content. Content is a new arena for traditional telcos, one they'll need to master, but cable companies will need to innovate as well.

When comparing content, consumers want two things: more and exclusive. Because content can be stored on the network, the communications provider is in an ideal position to deliver a very wide variety. Consumers respond: As Cisco's Pankaj Gupta points out, “One-fifth of Netflix rentals are titles other than its 3,000 most popular ones, and Rhapsody streams more songs outside its top 10,000 than within that group. Many providers are finding that niche content is helping drive their growth.” So, many people will choose a provider that offers all of Alec Guinness's 1940s comedies and other movies they don't find at even the big-chain video stores.

Service providers should also make it easy for consumers to enjoy IPTV and its possibilities. The provider can set up the in-home IP networks that empower customers and supply the router gateways and the set-top boxes with DVR and DVD-burning capabilities. Home networking enables triple or quadruple play over one line coming into the home and one core infrastructure, so providers can offer truly bundled services cost efficiently—bundles that create sticky retention. The opportunity is here. Most major markets in the US now have from 14 to 17 HDTV channels, and about 40 percent of US homes are expected to have at least one HDTV set by 2007. According to industry research, of consumers with DVRs, more than 80 percent report using the device to watch a recorded program at least several times a week, and more than 70 percent report watching one program while recording another. And as consumers use more services that they control, they report notably greater satisfaction with their provider.

a view from the net

VIDEO
SURVEILLANCE
JOINS THE
IP NETWORK.

VIDEO SURVEILLANCE delivered by analog closed-circuit television (CCTV) technology has long been an integral component of an organization's physical security strategy. Surveillance applications, both old and new, promise to become still more valuable as video recordings move onto the corporate IP network.

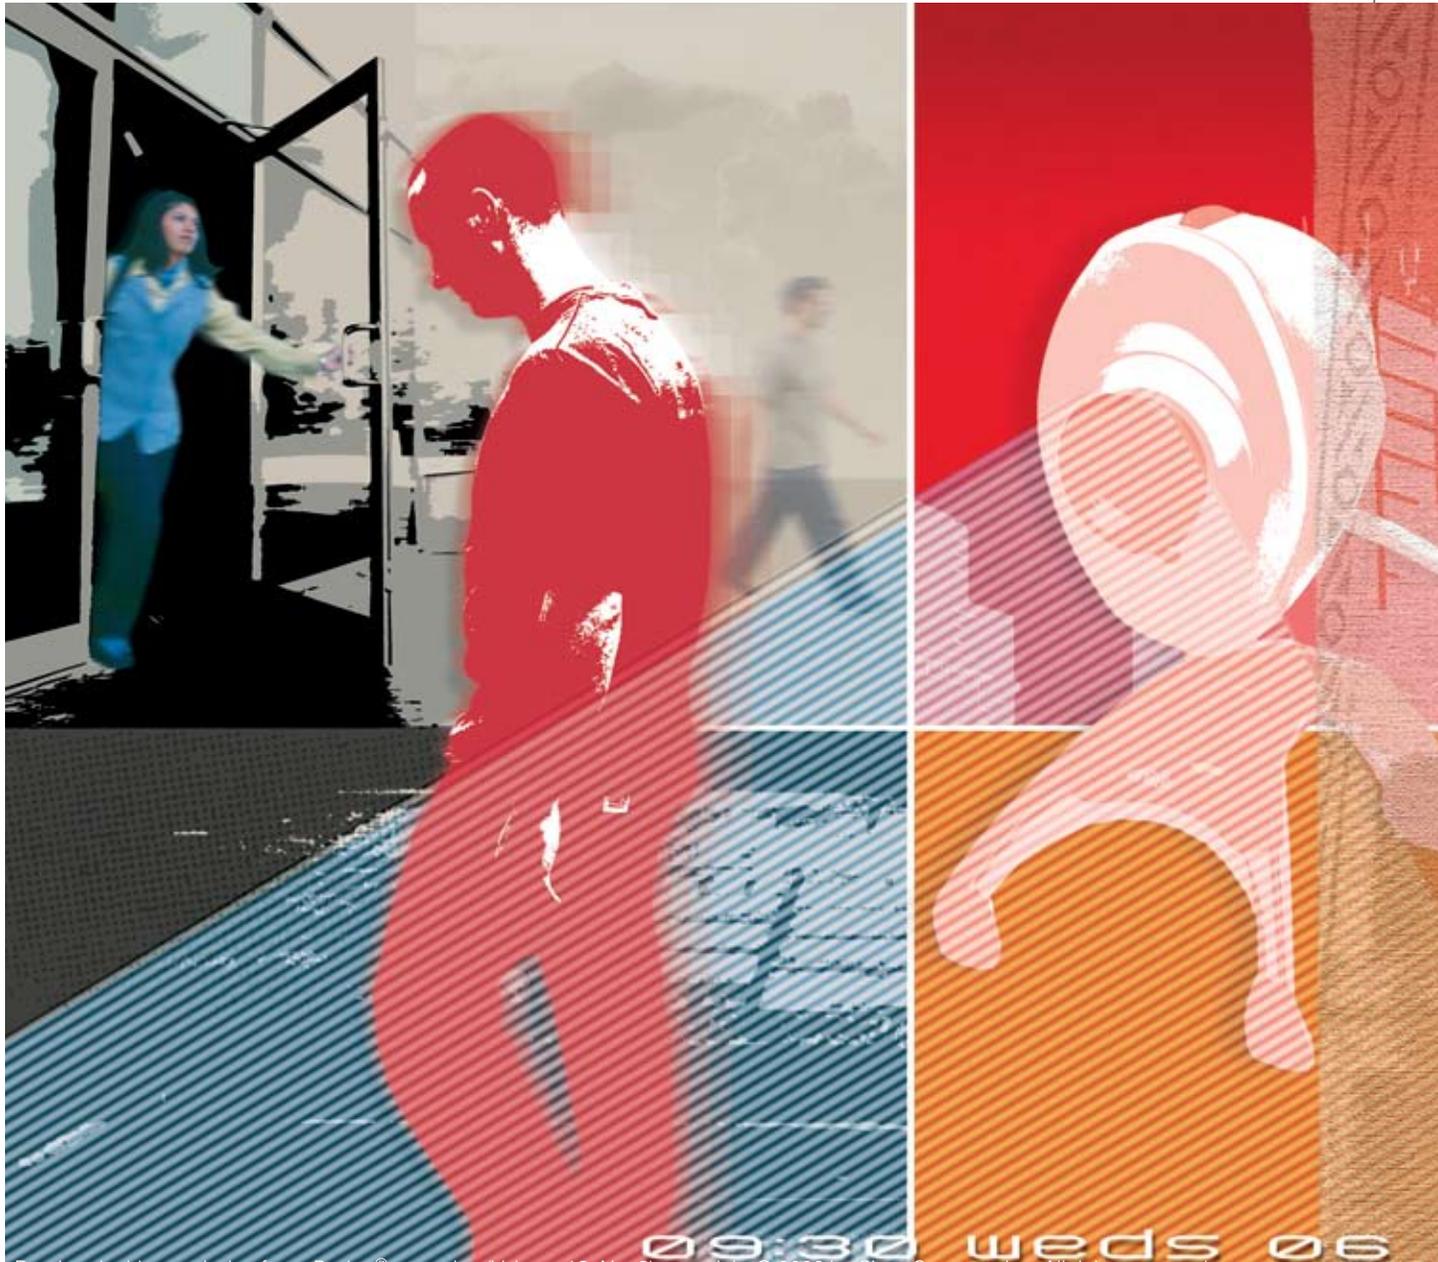
New uses for video surveillance beyond security are emerging. For example, local retail store managers use CCTV systems to identify the need to open or close checkout lines based on the length of customer queues. They can also determine whether a merchandising display is successful by observing real-time customer behavior and taking action accordingly. Transportation companies, for their part, use recorded video to help track and validate the movement of cargo.

Being able to conduct these functions remotely across an IP network—rather than having to be located in a special control room—means that personnel in distributed offices are empowered to make surveillance-centric retailing, merchandising, and tracking decisions, too. Security personnel gain the ability for real-time response, investigation, and resolution.

Why Deploy IP Surveillance?

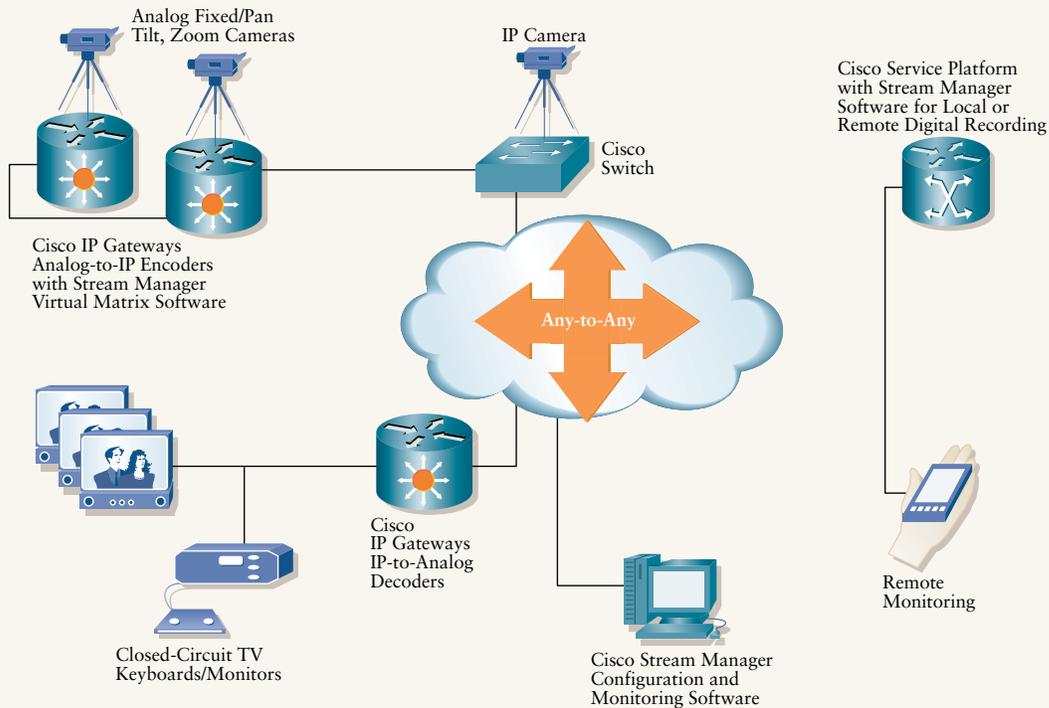
Analog CCTV surveillance systems have traditionally been standalone systems falling under the purview of the company facilities manager. However, organizations that have appointed chief security officers (CSOs) are beginning to blend the facilities and IT departments in efforts to move surveillance onto the corporate network. The goal is to deliver surveillance content as another IP network service—one that enables remote real-time viewing from any network-connected location and no longer requires users to hunt through banks of tape cassettes to find specific recorded material. In the meantime, making the surveillance capabilities IP-centric enables different vendors' once-proprietary equipment to work with one another.

New IP-based surveillance systems work something like TiVo. . . . Events being recorded can be viewed from any network-connected location.



MATT HERRING

Intelligent Converged Environment For Video Surveillance



IP NETWORK-CENTRIC VIDEO SURVEILLANCE Cisco IP video surveillance enables multivendor device interoperability.

With surveillance video accessible over the corporate IP network, network-connected personnel, such as security guards with mobile devices walking the grounds, can view video as it is recorded in real time and respond to incidents more rapidly. Video recordings can also be accessed later from hard-disk storage.

“This approach consumes less real estate, and finding a specific video segment is faster than manually storing hundreds or thousands of tape recordings and having to wind through them,” says Steve Collen, marketing director in the Cisco Converged Secure Infrastructure Business Unit.

New IP-based surveillance systems work something like TiVo—the popular brand of digital video recorder (DVR) that has revolutionized how many people watch television programs. Events being recorded can be viewed from any network-connected location. Content watched “live” can be paused or rewound to repeat a sequence.

Because there are efficiency, safety, system security, and other enterprise benefits to IP-enabling and network-connecting

surveillance systems, IT and network managers who work in companies that are early technology adopters are gaining at least partial responsibility for the surveillance discipline. Collen estimates that 5 to 20 percent of large organizations have teamed the IT and facilities departments on physical security. In its 2006 report, “World IP Surveillance Markets,” international researcher Frost & Sullivan determined that the need for extending remote accessibility of real-time content and the inability of existing proprietary, self-contained surveillance systems to deliver it will be the two strongest reasons for deploying IP-based surveillance systems for the next two years.

Let’s look at how to migrate to a network-based surveillance system in the context of these drivers and benefits.

From Analog to IP

Historically, analog CCTV surveillance systems have entailed setting up a fixed control center for viewing recordings. Individuals not physically in that control center can’t view what is

being recorded. Typically, dedicated communications links exist between a fixed camera(s) and the control room. And generally, CCTV systems are proprietary, so encoders/decoders, cameras, keyboards, and monitors from different equipment makers have not been interoperable.

This situation has limited surveillance to those viewers physically present in the control room, and has required multiple dedicated links for viewing in more than one control room. It has confined solutions to single-vendor components and has made it cumbersome and time-consuming to find specific segments of recorded video.

Using IP as the standard communications protocol, however, different vendors' video surveillance equipment can interoperate, giving enterprises a broader choice of component suppliers (see Figure, page 40). By putting an Ethernet interface and an IP protocol stack onto the surveillance equipment, the surveillance function feeds into the corporate network for remote viewing in real time and for remote access to recorded content.

It's not necessary to eliminate existing analog equipment to merge surveillance onto their networks. Analog-to-digital convergence equipment from Cisco, in the form of a video

surveillance gateway, allows companies with traditional CCTV systems to convert the analog video format to a digital format and IP-enable it for transmission across the corporate IT network. The Cisco product portfolio includes analog-to-digital video encoders/decoders, analog video transmission equipment, video surveillance recording servers, and video surveillance management software.

These components, gained in Cisco's April 2006 acquisition of SyPixx Networks, work with the advanced features and functions of Cisco IP network switches and routers. For example, companies can integrate video surveillance with access control and intrusion prevention system (IPS) capabilities. Integration with network IPS capabilities managed by the IT staff protects the surveillance system from viruses and other malware alongside other network server components, rather than leaving the surveillance system individually exposed.

Integration of the surveillance system with alarm systems helps alert personnel to problem situations and then provides critical video details. This can help cut the expense of responding to false alarms, which can exceed 90 percent in many organizations, according to Collen. For example, with video surveillance, a security officer could determine that a "door-forced" alarm was triggered by a gust of wind rather than an intruder.

The Cisco encoders/decoders run Cisco Video Surveillance Stream Manager Gateway software, enabling every Cisco Video Surveillance IP Gateway to become part of a highly available, distributed "virtual matrix switch" formed by Cisco Catalyst Ethernet switches. Standalone analog systems required a separate and costly analog video matrix switch. Other equipment, such as fiber-optic distribution amplifiers and multiplexers, are also eliminated.

Starting at Home

Cisco itself began its own deployment of a network-centric IP video surveillance system in 2003. Cisco now uses the IP network for video transport and has replaced analog video cassette recorders (VCRs) with network-centric digital video recorders (NVRs), which can be viewed and controlled anywhere on the network.

"With our IP-based system, we can create command center operations in

Lessons Learned: Cisco Experience

Like any technology migration, migrating video surveillance from a stand-alone analog system to an IP network-based service carries its fair share of bumps and bruises. Bill Jacobs, senior manager of Risk Technologies at Cisco, has overseen this evolution at Cisco.

IT departments that are unfamiliar with video surveillance may be initially cautious about adding video to the network, says Jacobs. But they needn't be. First, video surveillance doesn't have to be streamed in real time, and it doesn't require 30 fps quality—lessening its impact on the network. Also, once a detailed network design is done, IT will be able to visualize the video flows and see that not all video will be transmitted across the WAN. Most will stay on individual LAN segments.



Jacobs urges that those embracing IP video surveillance deploy a system that is scalable for growth and is based on open systems for interoperability. "Be sure the video surveillance you deploy can take advantage of [ITU-T] H.264, the digital video codec standard for very high data compression, and 802.1af authentication [an extension of the 802.1X framework], now under development."

Further Reading

- Cisco Video Surveillance
cisco.com/go/videosurveillance
- ITU-T H.264 Home Page
www.itu.int/rec/T-REC-H.264/en
- IEEE 802.1af Draft Standard
www.ieee802.org/1/pages/802.1af.html
- Cisco IT at Work Case Study
cisco.com/packet/183_6b1
- Cisco IT at Work Video
cisco.com/packet/183_6b2

real time, which we never could do before,” says Bill Jacobs, Senior Manager of Risk Technologies at Cisco. “We can respond to incidents much faster. The [return on investment] is tremendous, because we can centrally manage, predict, and even respond to service calls from locations as far away as

Gateway encoders/decoders, however, use high-performance digital signal processors (DSPs) and application-specific integrated circuits (ASICs) to ensure low-latency, broadcast-quality video. The system delivers 30 fps in National Television System(s) Committee (NTSC) format and 25 fps for phase-alternating line (PAL), Common Intermediate Format (CIF), 2CIF, 4CIF, and D1 formats.

In addition, the MPEG-4 compression format cuts the 4- to 5-Mbit/s streams generated by older MPEG-2-based equipment to about 3.5 Mbit/s, says Bob Beliles, manager of product marketing in the Cisco Converged Secure Infrastructure Business Unit. He also notes that Cisco IP video surveillance products and software take advantage of IP Multicast technology to save additional bandwidth. Rather than transmitting a unicast stream to each recipient wishing to view content, multicast technology sends a single multicast stream that can be offered through subscription to as many users as desired. The network infrastructure, preferably at the point closest to each subscriber, handles the replication of the video to multiple devices, minimizing bandwidth consumption over shared links.

With our IP-based system, we can create command center operations in real time, which we never could do before. The return on investment is tremendous, because we can centrally manage, predict, and even respond to service calls from locations as far away as Beijing from a central location.” —Bill Jacobs, Senior Manager of Risk Technologies, Cisco

Beijing from a central location, such as Cisco’s San Jose, California, headquarters.” (See sidebar, “Lessons Learned: Cisco Experience,” on page 41.)

The equipment can generate multiple streams at different frame-per-second (fps) rates, which Deon Chatterton, a program manager in the Risk Technologies team at Cisco, sees as valuable. “We’ll be able to view one frame rate and resolution with live video, but record at a slower frame rate for storage. The new equipment will allow us to maintain a good compression and storage rate but still have excellent live video quality.”

Technologies and Formats

Significant technology advances of the last few years have helped accelerate the movement of video surveillance to IP networks. For example, many traditional video surveillance products rely on software-based compression implementations that can produce poor video quality. The Cisco Video Surveillance IP

The Future: Video Analytics

As common formats for video and control signals mature, third-party vendors have begun writing new applications that allow greater intelligence to be applied to the video. One is video analytics.

For example, in a high-value area such as a specialized warehouse or other secure buildings, rules could be written into the video software so that if movement is detected a few hundred feet from the building, the camera would zoom in. Then, using more specialized intelligence and pattern matching, it could further determine if a human or animal were causing the movement. If human, before setting off an alarm—a standard response—the application could issue a pre-alarm that would generate a voice-over-IP (VoIP) message across the network before the intruder reaches the building that says “you’re on private property.” ■

THE CONNECTED

HOME

**Networked
Entertainment
That Puts You
in Control**



IN THE CONNECTED HOME, YOUR network is about so much more than simply sharing the Internet and connecting computers together. It's about getting and sharing multimedia and being able to enjoy it from any room in your house. It's about self-management of audio and video content connected to your home entertainment center. It's about remote monitoring of things and places in your home that are important to you. On a basic level, it's about choice and mobility—getting the information you want, when you want it, on any device (see Cover Story, page 32).

In the Connected Home, entertainment and communications activities are married onto one network, preferably a wireless one for the greatest mobility. On the next few pages are some of the activities possible in a Connected Home today—from sharing digital content, to online gaming, home surveillance, and voice over IP—all untethered. The home network in the not-too-distant future will undoubtedly bring us much more.

ARNE HURTY

by joanna holmes

time

JONATHAN BARKAT



to migrate?

MAKING WAY FOR THE NEXT GENERATION OF ROUTERS

ACCORDING TO MOORE'S LAW, rapid advancements in processing power typically render PCs and servers obsolete every 12 to 18 months. A similar dynamic applies to networks, pushing older switches and routers towards obsolescence after a number of years. True to this natural progression, Cisco recently announced the pending retirement of the Cisco 1700, 2600, and 3700 series multiservice routing platforms.

But with extensive support and useful programs from Cisco combined with a highly evolved, versatile new generation of routers, users can expect an easy migration to their next router platforms. In March, Cisco announced to customers the end-of-sale (EOS) and end-of-life (EOL) schedules for Cisco 1700, 2600, and 3700 series multiservice routing platforms, with the first of these milestones taking

place in March of 2007. These platforms began shipping in the late 1990s, and over the past 18 months many Cisco users have already begun replacing them with the newer, more services-ready Integrated Services Router (ISR) product portfolio (see the sidebar "End of Sales Details" on page 48 for specific information on the EOS and EOL dates).

The announcement reflects a need for customers to continually evolve their networks to meet new business requirements, according to Michael Shorts, a marketing manager in Cisco's Solutions Marketing Team. "The branch or access platform you invested in a few years ago was best-in-class at the time, but today it can't keep up with the new demands being placed on networks," Shorts explains.

Through several programs and tools, Cisco is easing the migration path for customers that are ready to move to an ISR platform. "With more than one million Integrated Services Routers shipped since their launch in September 2004, the ISR has transformed the access router into a services platform," says Dave Frampton, vice president of marketing in the Access Technology Group at Cisco. "By deploying Integrated Services Routers, our customers can converge their infrastructure and enable new applications and services at branch and remote sites."

Networks in Days of Yore

TO FULLY APPRECIATE CISCO'S end-of-sales announcement, Brian Ryder, a product line manager in the Access Technology Group at Cisco, refers to the networking world's status quo when these products were being developed. "This really is a story of network evolution," Ryder says.

Take, for example, the Cisco 1600 and 3600 router platforms, designed a decade ago. "In 1996, state-of-the-art branch office networking comprised 56K leased lines, Frame Relay, or maybe 128-Kbit/s ISDN BRI. That was your primary access for your branch office—and that was high speed," says Shorts. As for security requirements, "Well, there was no security, because you weren't connected to a shared public network of any kind."

Network services were a different matter, too. In 1996, the industry focus was on reliable connectivity for IBM terminals. File transfers tended to be the most bandwidth-intensive applications at the branch. Integrating services into routers and switches was an idea whose time had yet to come. When the design team for the Cisco 1600 integrated a 56K CSU/DSU into the product, says Ryder, "It was pretty radical."

VoIP Evolution

VOICE OVER IP (VOIP) became a hot topic in the late 1990s, and Cisco introduced VoIP support in 1998 on its new Cisco 2600 platform. "We supported voice on the first shipments of these products, but it was an add-on to the product design—it wasn't integrated," explains Mark Monday, vice president for voice

marketing in the Cisco Access Technology Group.

In Cisco branch office products, VoIP as an integrated service debuted with the 1999 introduction of the Cisco 1750 router. "This was the first product based on Cisco IOS Software that was designed from the ground up for VoIP," says Monday.

The next stop along the evolutionary path was the Cisco 3700 Series Router, whose development began around 2000. "The Cisco 3700 introduced new voice features, such as the time-division multiplexing (TDM) bus, which let you send voice streams throughout the system—and these features were very important for a voice-enabled router," recalls Monday. However, this platform fell short of accommodating digital signal processors (DSPs) for voice directly on the motherboard. "The Cisco 3700 lacks the level of voice integration we have today," he notes.

Security Needs

BY THE LATE 1990s, security needs in the branch had gained momentum. In 2000 Cisco released its first hardware-based encryption card for the Cisco 2600, which helped secure the new surge of virtual private networks (VPNs). But with limited expandability, processing power, and memory, Cisco 2600 users couldn't run voice and security services concurrently. (A "midlife refresh" of this product in 2002 produced the Cisco 2600XM router models, which offered extra memory and expanded the platform's ability to handle new services.)

When Cisco began development on the ISR platform around 2002, the state of networking had changed drastically since the days of designing the Cisco 1600 and 2600. Two major changes were the pervasiveness of the Internet and, accordingly, low-cost Internet connectivity. "You could get cheap Internet access at Point A and Point B, and then create a VPN tunnel and do encryption across it," says Jennifer Lin, director of marketing in the Access Technologies Group at Cisco. "But when you did that, suddenly you were connected to the public Internet, and

End-of-Sales Details for the Cisco 1700, 2600, and 3700 Series Routers

Cisco released information regarding the end of sale of all Cisco 1700 Series fixed and modular routers, all Cisco 2600XM Series, and the Cisco 2691 (note the exclusion of the 2621XM-DC service provider platform), and all models of the Cisco 3700 Series router. The 36-port EtherSwitch modules and first-generation T1/E1/J1 digital voice network modules (NM-HDV), J1 voice interface cards, and associated DSM (PVD-12) are all affected. Spares will be sold for one year past chassis end-of-sales. The EOS announcement was made on March 27, 2006, and these platforms will reach end of sale status on March 27, 2007. Software maintenance will end no sooner than March 27, 2010, and the products reach final end of supported life on March 27, 2012.

the security concerns multiplied a hundred fold.”

Whereas in the days of the Cisco 1600 and 2600 it was enough to simply provide connectivity, product design requirements were now calling for encryption, a firewall, intrusion detection and prevention systems that could run effectively alongside other network services, and applications. “With the arrival of things like denial-of-service (DoS) attacks, the list of required security features just exploded,” says Lin—and these features are very processor- and memory-intensive.

“Products that were developed back in the 1990s, when security wasn’t a big deal, don’t have the processors or memory resources to keep up with the threats that are on modern networks,” Lin observes.

When the Cisco 3700 routers were launched in early 2002, they presented a good set of security capabilities, but these capabilities weren’t integrated. It was nevertheless a significant product introduction. With its high performance and flexible modularity, the Cisco 3700 was the first single-unit platform that successfully brought IP telephony and security together in the branch—and it did so with agreeable performance. “That was important, because it enabled new applications,” says Lin.

Enter Integrated Services

CISCO’S DEVELOPMENT OF the Integrated Services Router portfolio and the subsequent phasing out of the Cisco 1700, 2600, and 3700 series are a logical progression in this continuum. “We’ve arrived here through a long evolution of these end-of-sales products,” says Shorts. “The Integrated Services Routers continue their legacy, picking up where the other products leave off.”

New benchmarks for branch routers were defined when Cisco launched its Integrated Services Router portfolio in late 2004. Foremost among these trends is the integration of services into the router (thus doing away with the complexity of extraneous devices). ISRs can replace functionality that was previously provided by other external devices, which offers an array of benefits to IT staff. “We provide, for instance, Cisco CallManager Express on the ISR products to take over all capabilities of a PBX or voice switch,” says Robert Checketts, a marketing manager for enterprise routing and switching at Cisco. “Call forwarding, picking up and sending calls, placing calls on hold, voice mail, autoattendant—all these features run right on the router itself.” This integration of services extends to security and voice capabilities and goes a long way toward reducing network complexity and costs of ownership.

Another factor in the ISR’s design was its ability to support multiple concurrent services at wire speed. “You can turn on all these services—security, IP telephony, wireless capabilities, and applications networking services—and still keep your WAN pipe completely filled,” says Checketts.

Some elements of the earlier router platforms will ease cus-

tomers’ migration to the Integrated Services Routers.

For instance, the Cisco 1600 and 3600 introduced the form factors of the WAN interface card and the network module, and those same interfaces are in use today on the Integrated Services Routers. “That provides a lot of investment protection,” comments Checketts,

“because IT staff can continue to use those same familiar interfaces, and they can share interfaces between product lines.”

Programs in Place

“WE PROVIDED A ONE-YEAR EOS NOTICE because we want Cisco users to become comfortable with the new product and to finish their final rollouts of projects they already have in place,” says Shorts. A general recommendation when preparing to purchase new equipment: For projects that call for EOS products, use the coming months to complete deployments with your existing platform. For any new projects, plan to move to a product platform with greater flexibility for new services.

Several Cisco resources to assist users with their platform migration strategies:

- The Cisco Technology Migration Plan provides customers with a trade-in credit toward the purchase of any new Cisco product. The program underscores Cisco’s commitment to provide effective migration options in the face of continuously changing network requirements.

- The Cisco Discovery Tool is a free network-profiling tool that can assess all components in a Cisco network and quickly analyze and identify the location of Cisco equipment. It creates a detailed report of all connected Cisco devices, including what versions of Cisco IOS Software are in use, and what products are EOL or EOS.

“Cisco is responsibly retiring aging products from the market with a tried and true mechanism that gives customers six years to develop and implement a technology migration plan,” writes Joel Conover, principal analyst for enterprise infrastructure at research firm Current Analysis. Conover describes Cisco’s product retirement mechanism as “one of the most open and customer friendly in the industry.”

In the coming years, networks will be tasked with a raft of new demands to evolve toward service-oriented architectures. “If you’re designing a new network today,” says Checketts, “you’ll want a foundation that fully supports the features and services your business will need to turn on in the next few years.” ■

Further Reading

- Cisco Technology Migration Plan
cisco.com/packet/183_6d1
- Cisco Discovery Program
cisco.com/packet/183_6d2
- Cisco Routers
cisco.com/packet_183_6d3

3 Steps to Network Virtualization

TAKE ALL THREE WITH THE CATALYST 6500 SERIES SWITCH by lori gadzala

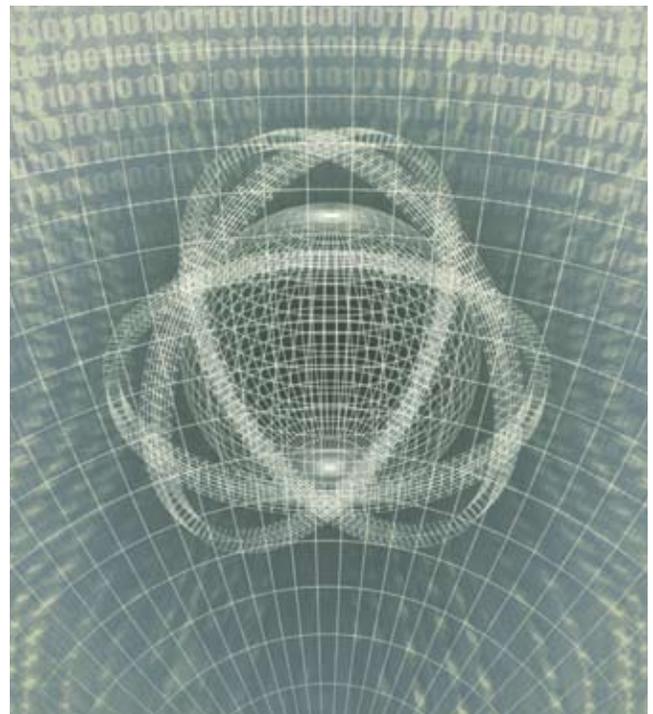
A

s companies expand their networks and require faster application deployment and more storage, the number of devices in the data center noticeably spikes. Yet many of these devices remain underused, siphoning both time and money from the IT budget. So what can your company do to allocate these network resources more efficiently? Transform physical network devices into virtual resource pools. / *Virtualization*—the logical segmentation of a single physical network—began with segregation of enterprise LAN traffic into virtual LANs (VLANs). Rapid growth in storage demands and capacity resulted in the development of storage switches, such as the Cisco MDS 9000 Series Multilayer Switch, to aggregate multiple disk drives into a single storage-area network (SAN) with partitioning by virtual SANs (VSANs). / Today, network virtualization is even more intelligent. As companies expand globally with larger and more complex data centers, deploy dozens of applications, partner with others, and comply with regulatory rules, their networking needs become more complex. Virtualization features in the Cisco Catalyst 6500 Series Switch address the challenges associated with deploying application services and security policies in a scalable, distributed environment.

New functionality in the Catalyst 6500 Series provides large-scale virtualization of application delivery and security services. Released in April 2006, the Cisco Application Control Engine (ACE) consolidates the functions of multiple network devices and allows logical partitioning of the ACE physical resources into virtual contexts (see figure, page 52). Network virtualization can be achieved in three steps: *access control*, *path isolation*, and *policy enforcement*.

First Step, Access Control Increasing collaboration with customers and partners requires multiple levels of access privileges to a range of applications. Visitors to corporate offices often expect to have “guest” wired or wireless access to the Internet. Corporate suppliers, such as contract manufacturers, might also work for a competitor, extending traffic segmentation outside of the enterprise. This complex web of access requirements is often addressed with multiple physical networks, creating significant management complexity and duplication of physical devices and services.

IEEE 802.1X port authentication standards in the Catalyst 6500 Series extend access control to the media



SHARE NETWORK
resources with
secure separation
between applica-
tions, groups, or
individuals.

layer. The identity of incoming users or machines can be used to permit or deny access and apply traffic policies. Users are tightly associated with their VPN or network partition and thus are confined to permitted areas.

Rice University in Houston, Texas, uses this functionality to separate student Internet access from internal department traffic and inter-university research requirements. “As our faculty research and student expectations have grown, so has our need to deploy safe, reliable Internet access,” says William Deigaard, director of networking telecommunications and data center operations at Rice University. “We are employing the Catalyst 6500 to partition our network into multiple networks and apply unique policies to each. We use the network virtualization capabilities to manage and protect the campus network, including differentiating who people are and supporting visitors in the friendly, collegial fashion.”

Second Step, Path Isolation

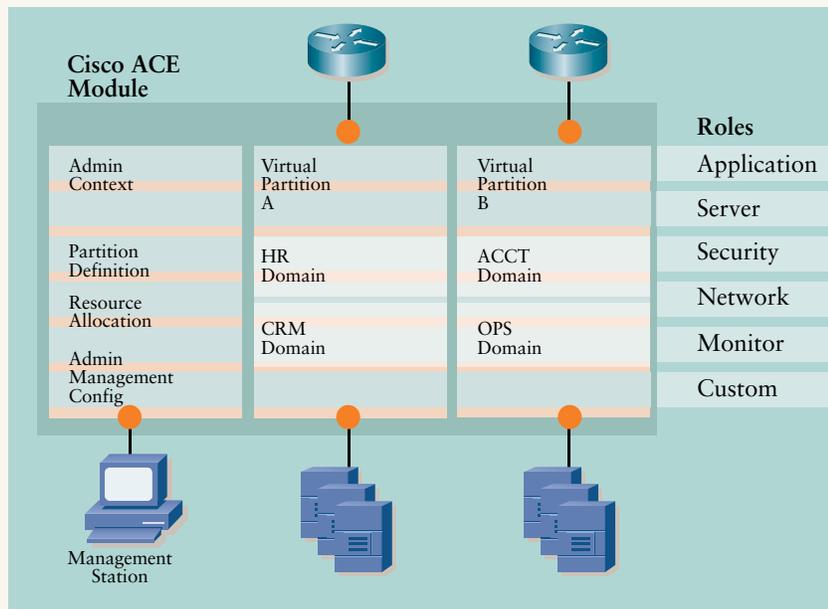
THE SECOND STEP IN NETWORK VIRTUALIZATION is to isolate various traffic flows. Regulations and privacy concerns require that departmental applications, such as finance and human resources, be segregated. The Catalyst 6500 Series handles this segregation through VLANs and Layer 2 or Layer 3 switching, and also enables multiple types of path isolation for closed user groups. These include Generic

Routing Encapsulation (GRE) tunnels for creating a small number of closed user groups on the campus network (e.g., guest VLAN access); virtual routing and forwarding (VRF)-lite for campus segmentation where IP addresses can be overlapped among the VPNs (each group can independently use private IP addressing); and Multiprotocol Label Switching (MPLS) for establishing closed user groups through VPNs transported independently over the network core (any VPN can be configured to connect users and resources at any location in the network).

Unique Zurich Airport, operator of the Zurich Airport in Switzerland, needed a network to support a wide range of applications, including public Wi-Fi access, airline operations, and tightly secured air traffic control, for 180 different companies. Typical enterprise approaches suffered from scalability or troubleshooting issues. Service providers have used MPLS VPNs to provide this type of service for many years, but most enterprise switches lacked this functionality. The virtualization features available in the Catalyst 6500 Series enabled Zurich Airport to use MPLS VPNs to support a wide range of network connectivity and performance requirements across multiple virtual partitions, all on a single physical infrastructure.

“The Cisco Catalyst 6500-based network at Zurich Airport allows us to offer ‘carrier-grade’ network services to our Zurich Airport customers including airlines, airport operations, and

Cisco ACE Application Infrastructure Control Features



ACE IN THE HOLE Cisco ACE consolidates the functions of multiple network devices.

Modular Growth with Cisco ACE

Cisco ACE for the Catalyst 6500 can slice resources into 250 virtual partitions. Each partition can be defined by customer, business organization, or application, and resource allocations such as bandwidth or number of connections can be defined for each partition.

Role-Based Access Control (RBAC) in Cisco ACE allows each virtual partition to be managed by the appropriate business or IT team. In addition, within each virtual partition, up to ten management domains can be created, providing further granularity for controlling resources within that virtual partition. The network group can configure Layer 3 variables. Security can be applied

from centralized policy or specified for individual partitions based on application, business organization, or user. Application and server departments can monitor and manage their virtual servers without risk to other IT configurations. Servers can be taken in or out of rotations for maintenance by the application or server owners based on group ownership.

This RBAC flexibility enables faster service deployment, simplifies workflow within IT, and reduces configuration errors. Centralized security enables consistent enforcement throughout the organization, and reduces the complexity and operational expenses of maintaining multiple policies.

Furthermore, Cisco Application Networking Manager (ANM) software (a server-based management package) simplifies management of Cisco ACE virtual partitions.

Cisco ANM provides a single interface for configuration, maintenance, operations, and performance monitoring of virtual partitions within and across Cisco ACE modules. Template-based configurations enable organizations to rapidly partition applications. Multiple concurrent administrators can be active across partitions and modules. Partitioning of functions simplifies and shrinks the configurations and reduces the probability of errors.

additional services—a typical service provider technology at the price point of an enterprise network,” says Peter Zopfi, head of communications engineering at Unique.

More Applications, Greater Availability

A SIGNIFICANT CHALLENGE AFFECTING application performance is the increasing number of services being performed by servers. Activities such as Secure Sockets Layer (SSL) encryption/decryption, TCP optimization, multiple levels of security checks, and rich-media processing are siphoning server capacity.

The new ACE service module in the Cisco Catalyst 6500 can intelligently load balance application traffic to server farms with market-leading throughput, connection setup rates, and performance scalability via software licenses rather than truck rolls.

With the introduction of virtual partitions, up to 250 per module, the Cisco ACE allows exceptional control of the application delivery infrastructure. For each virtual partition, administrators can tune the processing resources—bandwidth, connection setup rate, SSL transaction rate, syslog rate,

etc.—as well as many memory resources—number of concurrent connections and access control lists (ACLs), etc. Thus, business organizations, customers/subscribers, and applications can all share a physical ACE module with complete isolation among them.

Most importantly, virtual partitions empower operators with the ability to turn on a new application or service with a few clicks rather than going through the tedious, time-consuming process of selecting, qualifying, deploying, and troubleshooting a new device.

The Cisco ACE employs a variety of optimization techniques. TCP connections can be pooled to individual servers so that new client TCP connection requests pose no additional server overhead. SSL sessions can be processed directly on Cisco ACE, significantly increasing SSL scalability and decreasing server load.

Application or module redundancy can be configured within a single chassis, across chassis within a data center, or between data centers. ACE offers the unique ability to protect at an application-by-application level across a pair of ACE modules using virtual partitions.

Third Step, Policy Enforcement

INTEGRATED SERVICES MODULES in the Catalyst 6500 Series help enable centralized policy enforcement. For example, security can be virtualized. ACE supports hardware-accelerated inspection and fixup of popular data center

Further Reading

- Cisco Network Virtualization Solutions
cisco.com/packet/183_7a1

SONA: Adapting Form to Function

Virtualization is a key component of the Cisco Service-Oriented Network Architecture (SONA). SONA provides a framework for helping enterprises make their networks more intelligent, less complex, and more efficient—allowing capital and operating budgets to shift away from infrastructure and toward applications and services that enhance productivity and competitiveness.

Designed for modularity, enterprises can migrate to SONA with incremental investments, preserving existing network designs. To learn more, visit cisco.com/go/sona.

protocols. It can scale up to 1 million Network Address Translation (NAT) entries and up to 256,000 access control elements. These can be divided up across many virtual partitions. For a broader array of protocol support for Internet/Intranet firewalling, the Cisco Catalyst 6500 Firewall Services Module (FWSM) also delivers multiple logical firewalls on one physical hardware platform using virtualization techniques. The ACE module works with the FWSM to load balance firewall traffic or to seek the FWSM's help in safeguarding against protocols not supported natively by ACE.

Overall, the virtualization features in the Cisco Catalyst 6500 Series Switch allow your company to share network resources while maintaining secure separation between applications, organizations, groups, or individuals. Network traffic for different departments, customers, and suppliers can be logically separated without having to build overlay networks or deploy separate devices. Critical applications can be deployed with fewer resources but improved deployment times. And the application infrastructure can be readily managed according to assigned roles in the IT department. ■

7200 Gets a Boost

PERFORMANCE, CAPACITY GAINS MEET RISING WAN/MAN SERVICES AGGREGATION NEEDS.

Pushing intelligent network services from consolidated data centers out to many distributed locations is on the rise because it gives organizations geographic hiring flexibility and real-estate cost advantages. As a result, the cumulative processing power and capacity required in the router at the corresponding aggregation site, or *head end*, across the WAN or MAN has also begun to rise. To handle the additional processing burden, IT departments running the premier Cisco enterprise aggregation workhorse—the Cisco 7200VXR Series Router—can simply upgrade their systems for greater performance with multiple services enabled rather than having to qualify and install a whole new platform.

Network managers can plug newly available hardware components into the system's chassis to aggregate greater volumes of traffic and integrate more services from more locations. The new components for the Cisco 7200VXR router chassis collectively double overall router performance and triple the IPsec VPN processing capability. They also add 50 percent more bandwidth and slot capacity and deliver speeds from OC-3 (155 Mbit/s) to sub-rate Gigabit Ethernet (up to 1,000Mbit/s) with multiple IP services enabled, such as security and voice, says Afaq Khan, Cisco 7200 Series technical marketing engineer.

The modular enhancements enable the administrators of the hundreds of thousands of Cisco 7200VXR routers currently installed worldwide to create a more powerful aggregation device out of their existing platform. Using this approach, existing Cisco 7200VXR router customers avoid making complete equipment upgrades, which add capital costs, cause network downtime, and require months of testing and staff retraining.

Saving "Millions"

EXTENDING ITS ORIGINAL ROUTER investments using the new modules "has saved us millions of dollars," says John Burns, vice president of network services at Wachovia Corporation, a US-based, US\$507 billion diversified bank holding company that offers various

brokerage, banking, and other financial services domestically and internationally. The company plans to deploy several of the recently released Network Processing Engine-Generation 2 (NPE-G2) processors for its Cisco 7200VXR Series routers to increase Compressed Real-Time Transport Protocol (cRTP) throughput. CRTP, an Internet Engineering Task Force (IETF) standard, decreases the size of IP, UDP, and RTP headers to accelerate latency-sensitive voice delivery.

A typical Wachovia branch design supports T1 access with two 384-kbit/s permanent virtual circuits (PVCs), and Wachovia aims to carry eight to 12 calls on-net at each site. The company has installed about 150 Cisco 7206VXR head-end routers, which can each support 120 remote sites, for a potential total of 960 to 1,440 concurrent cRTP flows. The NPE-G2 will help the company support this volume of calls.

"The more we can take a platform and scale it without having to retrain staff, the more efficient that is for us."

JASON SMITH, NETWORK MANAGER,
WACHOVIA CORPORATION

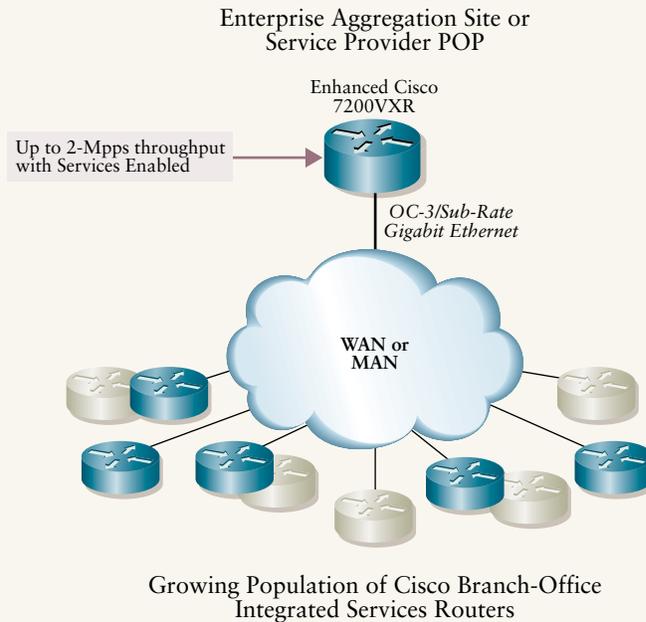
"Our older [7206VXR] platforms were just hitting their peak," explains Jason Smith, Wachovia's manager of network testing and certification. "The more we can take a platform and scale it without having to retrain staff, the more efficient that is for us."

Modular Improvements

SPECIFICALLY, THE CISCO 7200VXR ROUTER has gained the following optional plug-in modules that boost the horsepower of the device:

- Cisco 7200 NPE-G2—The new processor supports throughput of up to 2 million packets per second (Mpps).

Services Aggregation Model



THE ENHANCED CISCO 7200VXR ROUTER packs the power to process routing, voice, encryption, MPLS, and secure IP Multicast traffic generated by a growing number of distributed sites with upgrade modules, rather than a product replacement.

- Cisco 7200 VPN Services Adapter (VSA)—It scales encryption performance to 500Mbit/s with 300-byte packets and can be used to map IPsec sessions to MPLS for extending the MPLS network securely off-net. Like its predecessor, the SA-VAM2+ card, the VSA supports all key sizes of both Triple Data Encryption Standard (3DES) and Advanced Encryption Standard (AES) encryption—the same encryption supported in Cisco branch-office Integrated Services Routers (ISRs). The VSA installs directly into the router chassis' I/O controller slot so as not to interfere with any existing adapters. The VSA triples IPsec performance, compared to the SA-VAM2+, with either

3DES or AES encryption in use, even when voice sessions are running concurrently at the various locations.

- Cisco 7200 Port Adapter Jacket Card—It increases slot density and overall available bandwidth within the system chassis. The card holds a single port or service adapter, which installs into the router chassis' I/O controller slot. By taking advantage of the dedicated PCI bus that connects the I/O controller slot to the Cisco 7200VXR Series NPEs, the Cisco 7204VXR and 7206VXR chassis can increase their slot capacity by one while increasing the overall PCI bus capacity by 50 percent.

The jacket card supports the following port adapters:

- Cisco VPN Acceleration Module 2/2 + for use with NPE-G1
- Cisco VPN Acceleration Module 2+ for use with NPE-G2
- 2-Port Packet/SONET OC3c/STM-1
- 2-Port Channelized T3 Serial Port Adapter Enhanced
- 1-port multi-channel STM-1 multi- and single-mode port adapter

Network administrators do not have to reconfigure any existing interfaces after installing the jacket card, says Stefan Dyckerhoff, director in Cisco's Midrange Business Unit. "To migrate to the upgraded platform, you take out the old engine and put in the new one. You leave the existing cabling in place, which reduces network-change errors."



Increased Aggregation Requirements

OVER THE PAST 18 MONTHS, organizations have been ramping up their efforts to distribute intelligent network services across various locations, such as enterprise branch offices. The reason? Nearly 90 percent of employees work in branch locations, away from the headquarters facility, according to Nemertes Research, a firm specializing in quantifying the business impact of technology.

Putting distributed workers on a par with centralized employees from a network services perspective has been accomplished, in large part, with intelligent, multiservice branch-office routers such as the Cisco ISR family. The devices combine routing, data security, voice processing, wireless, and other capabilities in a single platform. Approximately 1 million ISRs have been installed and about 330,000 Cisco 7200VXR routers are in use at enterprise (and service provider) aggregation sites. A large percentage of the Cisco 7200VXR installed base aggregates ISR

Further Reading

- Cisco 7200 Series Routers
cisco.com/packet/183_7b1
- Cisco 7200 Series Routers data sheets
cisco.com/packet/183_7b2
- Miercom MPLS Diagnostics Expert benchmark report
cisco.com/packet/183_7b3 (click on report, "Package Speeds MPLS Diagnostics," November 2005)
- Services Aggregation in the WAN and MAN
cisco.com/go/servicesaggregation

traffic and services, so the modular hardware components help increase response times for the distributed, branch-office workforce (see Figure on page 56). Services integration in the Cisco

7200 Series Routers enables network operators to reduce operating costs and simplify installation, maintenance, and network management while protecting existing aggregation router investments as traffic volume grows. The platform remains the appropriate Cisco product to select for aggregating traffic up to OC-3 speeds with multiple services enabled. For aggregating multi-speed traffic from distributed sites at speeds above OC-3, the Cisco 7304 router, the 7600 Series Router, or Cisco Catalyst 6500 switch are higher-end options.

"The 7200 will remain part of the core network for years," says Tom Nallen, a manager in Cisco's enterprise routing and switching group. "We're continuing to invest in it and allowing customers to make modular upgrades to protect their investments." ■

A Word About the Software

The new Cisco 7200VXR Series hardware runs Cisco IOS Software Special Release 12.4(4)XD, which is based on 12.4(4)T and will later merge into the main Cisco IOS Software 12.4T train.

Meanwhile, enhancements have been made to Cisco IOS Software Release 12.4(6)T, which runs on the earlier versions of the Cisco 7200VXR router hardware as well as the Cisco branch-office ISRs, allowing end-to-end services delivery and aggregation.

For example, the Cisco 7200VXRs and ISRs running Cisco IOS Software Release 12.4(6)T both support Secure Multicast, an industry first. This Cisco IOS Software feature enables a router to apply IPsec encryption to IP Multicast traffic without having to configure overlay tunnels. Secure Multicast is of particular interest to organizations and applications supporting real-time, broadcast communications such as stock trading and video conferencing, says Stefan Dyckerhoff, a director in Cisco's Midrange Business Unit.

Cisco IOS Software running on the 7200, ISRs, and other platforms includes support for MPLS VPNs, as well as optional MPLS Diagnostic Expert software, available separately. The diagnostics tool speeds troubleshooting of MPLS networks for service providers or large enterprises self-managing their MPLS networks.

The tool's performance increase has been benchmarked by Miercom, a network consultancy based in Cranbury, New Jersey. The company's report states that diagnosing MPLS problems using the tool is 10 times faster than manual troubleshooting.

Can We Talk?

EFFECTIVE CALL ADMISSION CONTROL FOR COMPLEX NETWORKS by karl kocar

To successfully deploy IP-based telephony and video solutions in the enterprise, the network must provide appropriate quality of service (QoS) guarantees. While packet-queuing technologies vary, generally you must place delay- and jitter-sensitive voice and video within the highest priority queue (PQ). Police the PQ to ensure that both traffic types do not exceed configured bandwidth allocations. You generally mark data as “best effort” or place it in one of the lower priority queues. Adopting such a Differentiated Services (DiffServ) QoS policy, which guarantees bandwidth for different traffic classes, is an effective way to protect IP communications traffic from conflicting data traffic.

However, in the traditional circuit-switched world, the number of calls that can be supported between two endpoints is gated by their physical interfaces. For example, an E1 trunk between two private branch exchanges (PBXs) never carries more than 30 simultaneous calls. Thus it is comparatively easy to provision the correct amount of bandwidth required to transport all of these calls across the WAN. By contrast, in the IP world, no similar physical limitation exists on the maximum number of calls that can be attempted between two devices across a WAN circuit, and this has the potential to severely disrupt the entire service. For example, if the PQs on a pair of opposing WAN routers are configured to support a maximum of ten simultaneous voice calls, an eleventh call causes indiscriminate packet drops that negatively affect all conversations.

Call admission control (CAC) solutions enable you to protect voice against voice, and video against video in an IP WAN environment.

CAC Approaches

ONE APPROACH TO PROVIDING CAC is to query the network quality in real time using ping-type probes

before and during a call. If certain delay and jitter criteria have been exceeded, the call can be rerouted over a PSTN trunk. However, this approach has potential challenges. To be truly effective, every system endpoint must support the rerouting mechanism, and unless this type of support is standardized, an organization will be forced to use voice and video endpoints from a single supplier. Also, because this approach is reactive, if the network becomes over-subscribed every user might experience a period of poor call quality until remedial action is taken.

Another CAC solution uses the concept of a virtual trunk configuration on the IP PBX. The administrator manually configures the maximum number of calls that can be simultaneously transported between any two locations. These might be entered as bandwidth statements for the voice and video services. On the WAN routers, the relevant queues must be allocated the equivalent number of resources to match the IP PBX configuration. This manual approach to CAC works very well in static hub-and-spoke topologies and has been proven in many very large real world deployments.

Dynamic Call Admission

HOWEVER, IN CERTAIN SCENARIOS, a static CAC overlay might not be optimal. These scenarios include complex multi-hop topologies and topologies with redundant paths between sites. These topologies commonly occur in real life networks and require a dynamic, topology-aware call admission mechanism. If multiple circuits exist between sites, IP communications traffic must be able to take advantage of all the available bandwidth during normal operation. If one or more links become unavailable, the total number of available call admissions must be “throttled back” accordingly. During a circuit failure between two adjacent locations, the call admission solution must also support call routing through alternate paths in



**CALL
ADMISSION
CONTROL**
is essential
to guarantee good
voice quality in
a multisite deployment
involving
an IP WAN.

the network, but must still ensure that the number of calls established does not result in packet drops caused by queue overflows in the intermediate WAN routers.

The IntServ Approach

AN INTEGRATED SERVICES (INTSERV) approach can solve the call admission problems described above. The IntServ model, in existence since the mid-1990s, uses Resource Reservation Protocol (RSVP) to communicate an application's QoS needs and make call admission decisions along a routed path through the network. RSVP works with any existing routing protocol. Path and reservation messages are sent to set up simplex session flows. A path message, which requests a bandwidth reservation for a Router Transfer Protocol (RTP) stream, is initiated by the sender and follows the route to the destination provided by the routing protocol. At each hop along the way the routers store the "path state" for the flow. The reservation message, which is used to confirm the bandwidth request, is sent by the receiver back to the sender via the reverse route that the original path message took. At each intermediate router a call admission decision is made depending on the amount of bandwidth currently available for voice or video traffic. Two separate simplex sessions are required for a bidirectional communications path. If the reservation is accepted for both receive and transmit directions, the network has guaranteed that sufficient bandwidth will be available for the duration of the call. One caveat is that unless all applications on the network are RSVP enabled, only RTP streams under the control of the IP PBX should have access to the WAN router's voice and video queues.

You can use an IntServ model for call admission with an existing DiffServ QoS model. In this deployment type, the DiffServ implementation still provides the Differentiated Service Code Point (DSCP)-based queuing mechanism at each hop of the call path through the network and RSVP only for policing call admissions. Scalability is the main benefit of this approach.

An historical challenge with using RSVP for IP PBX CAC has been to coordinate the network resource reservations with call control signaling; in other words, to ensure that the path reservation between two IP endpoints has succeeded before the call is alerted. However, the majority of endpoints currently have no native support for RSVP. The logical solution to these problems is to embed RSVP agents into proxy devices such as routers and add the intelligence that allows them to interact with the call control platform to provide the signaling synchronization mentioned earlier. Thus, for any call across the WAN a minimum of two RSVP agents is involved in the reservation process. Adopting a proxy solution allows any existing voice or video IP device to use the dynamic CAC service. Currently, the interaction between the IP PBX and RSVP agent is vendor specific, but at a minimum it should have the following attributes:

Further Reading

- Cisco Unified Communications Solution Reference Network Design Guide based on Cisco Unified CallManager 5.0
cisco.com/packet/183_7c1
- Cisco Unified Communications
cisco.com/packet/183_7c2

- The RSVP agent should be able to simultaneously support both voice and video clients and provide differentiated handling of each. If using a DiffServ queuing model, the voice and video traffic should be marked differently. Ideally, this should be done by the endpoints under the control of the IP PBX. However, an RSVP agent should be able to police and remark DSCP for misbehaving applications.

- The dynamic CAC solution should be protocol independent and work with standards-based signaling such as Session Initiation Protocol (SIP), H.323, and Media Gateway Control Protocol (MGCP).

- The call-control signaling to the RSVP agent should cater for supplementary service support, such as diversions, transfers, and conferencing. Any necessary mid-call reservations should be signaled so that new call legs can be added and removed dynamically.

- RSVP reservation failure should be handled correctly by the IP PBX. The system administrator should be able to dictate whether the call will be rerouted across an available PSTN trunk or reclassified into the best-effort network queues.

- The RSVP agent architecture should be able to support multiple application types. In the future it may be desirable to treat contact center traffic differently from back-office voice applications or provide priority to the company's executive videoconferencing service over desktop video.

In summary, to provide effective call admission control for complex network topologies, standards-based solutions are needed that allow an IP PBX to dynamically react to changes in network topologies. These standards-based solutions must also ensure that any new call admission solution is compatible with existing endpoints regardless of the protocol used to signal calls. Intelligent communication between the IP infrastructure and the IP PBX call control function is the key enabler for the successful evolution to a ubiquitous call admission capability, and further increases the importance of the network in ensuring a successful migration from time-division multiplexing (TDM) to IP-based communications. ■

KARL KOCAR is a Cisco consulting systems engineer based in the UK. He can be reached at kkocar@cisco.com.

First HDTV over IP in the US

WITH IP OVER FIBER EXTENDED TO THE HOME, SUREWEST COMMUNICATIONS DELIVERS HDTV AND OTHER DESIRED SERVICES NOW.

S

ureWest Communications, based in Roseville, California, is aggressive about serving customers, from marketing what it calls guerrilla style—going into neighborhoods and knocking on doors—to being the first in the US to offer high-definition television (HDTV) over an IP network. / Its business model is based on symmetric Ethernet bandwidth of 100 Mbit/s over an IP network that goes all the way to the home over a single fiber. The company began to offer IPTV services in 2004. The offering included multiple streams of standard-definition TV (SDTV), voice, and very high-speed symmetrical Internet access. Early in 2006, it added MPEG-2 HDTV at 19.4 Mbit/s. With 100 Mbit/s to the home, SureWest is in position to add new services seamlessly. / Although the company has a 90-plus-year history, SureWest has the energy of a youngster. In 2002, it bought a local cable TV provider to add high-growth video services to its data and voice business. Not content to base its growth on the hybrid-fiber-coax (HFC) network it acquired, it began to plan for more services and the bandwidth and flexibility they would require. The first goal was triple play—voice, video, and data over one network and one connection to the home or business. And not just any video, but HDTV.

The company currently offers more than 275 standard video and audio channels, 17 HD channels, and video on demand (VOD) totaling about 900 hours at any point in time. Going forward, the all-IP network simplifies the provisioning of data, voice, and video services, and also streamlines the processes required for service delivery, service monitoring, and reacting to network performance. A Cisco foundation has also enabled enhanced data services, with the standard offering delivering 10-Mbits symmetrical data rates and optional 20-Mbit/s rates. The high bandwidth and low latency of the Cisco switches provide a level of responsiveness that appeals to gamers and other customers.

“Channel-change time is a good indicator of the overall network responsiveness,” says Scott Barber, vice president of network operations at SureWest. “Some carriers struggle with this. Our network is very responsive—more so than satellite and comparable to cable. We can take advantage of multicast for bandwidth efficiency, without compromising performance or degrading the user experience.”



FOR YOUR VIEWING PLEASURE
Delivering the experience that customers want is number 1 on SureWest's priority list.

PHILIPPE GELOT/GETTY IMAGES

A Network with Staying Power

THE IP NEXT-GENERATION NETWORK (IP NGN) SureWest chose employs Cisco Catalyst 6500 and 4500 Series switches for core and distribution tasks (see figure), and it includes pluggable optics for accommodating changes without requiring the expense of additional fiber. In the video headend/central office, two redundant Catalyst 6509 switches are fed video signals from Scientific Atlanta digital and analog satellite receivers and the VOD server, along with IP voice and data signals from other devices.

Each Catalyst 4510 Switch (providing 1550/1310 bidirectional single-fiber connectivity) delivers 100 Mbit/s bidirectional Ethernet to 384 residential customers. Located in cabinets in residential neighborhoods, the Catalyst 4510 remote terminals connect to a primary hub housing Catalyst 6509 switches for 40,000 homes. The Gigabit Ethernet uplink connection between the Catalyst 4510 and Catalyst 6509 switches can be upgraded to 10 Gigabit Ethernet as deployment densities increase, and SureWest also plans to upgrade the core switches to 10 Gigabit Ethernet. The primary hubs are connected upstream to the primary core locations, which also house Catalyst 6509 switches.

SureWest takes advantage of integrated quality of service (QoS) and security features in the Cisco gear, and the Cisco IOS multicast capabilities help to deliver IPTV and HDTV over IP efficiently. While traditional cable infrastructures deliver all channels to all houses, the Cisco Catalyst switches forward a single copy of the channels down to the remote Catalyst 4510 switches. At the most remote switch, each channel is multicast to only the homes that are watching it. At any point in the network, only one copy of each channel is being forwarded (see sidebar, “Multicasting with a Lot Less Bandwidth”).

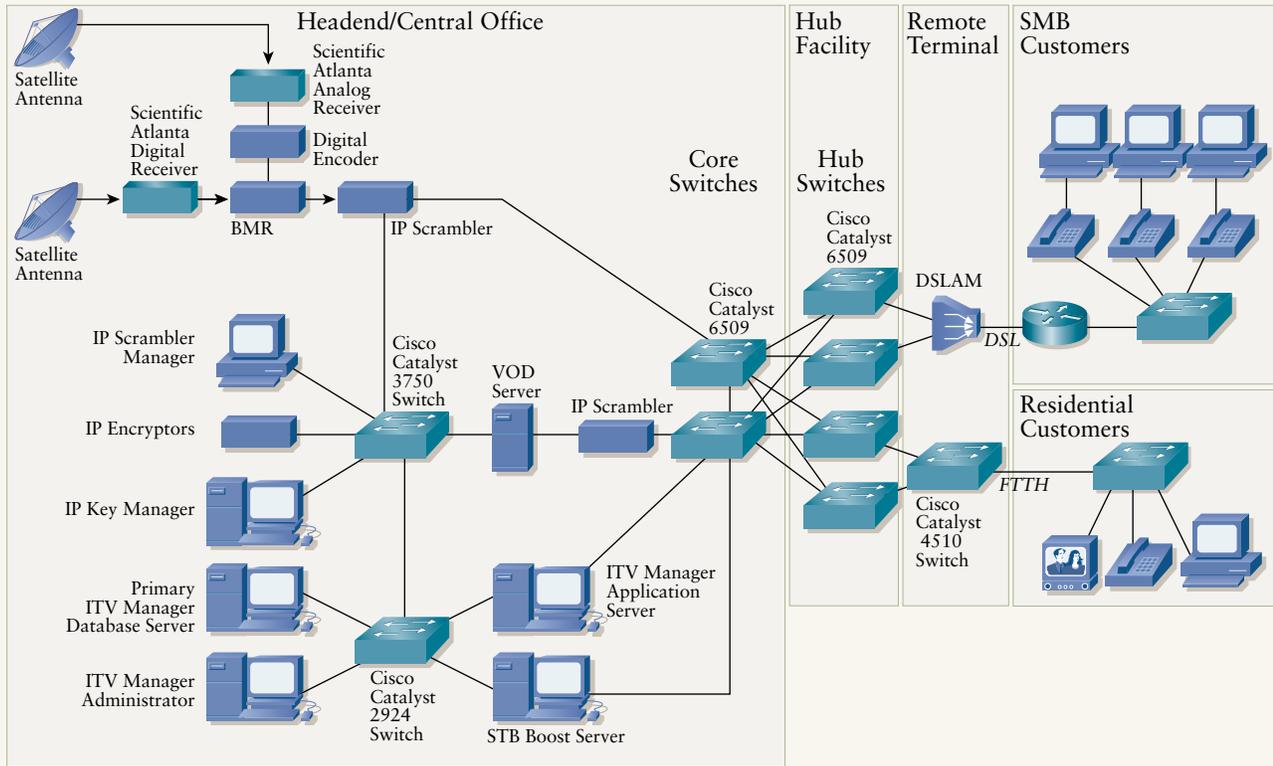
Further Reading

- Video/IPTV Solutions for Wireline Carriers
cisco.com/packet/183_8a1
- IP Multicast
cisco.com/packet/183_8a2

The Implementation

SUREWEST'S NEW CORE and distribution upgrade was deployed in 2003. Alpha trials were carried out within a controlled subset of the network, and a beta test began at the end

SureWest FTTH Distribution Network



VIDEO READY Among the advantages of SureWest's FTTH distribution network is its density—the ability to serve a large number of customers from small chassis footprints that deliver plenty of functionality.

of the year, providing IP video services to employees and some customers. By the beginning of 2004, the new Cisco IP NGN and FTTH architecture allowed SureWest to roll out IPTV services. The initial IP video service was provided to a new market in Sacramento, with the traditional cable infrastructure still supporting the installed base.

To date, SureWest has worked with leading IPTV companies to strengthen the network and expand services. “Developing our IP video business encompassed lots of integration,” says Barber. “The network, services, middleware, set-top boxes, and other components affect the overall viewing experience. Most of the work was focused on the video components—the Cisco core and distribution solutions worked from day one. Our new foundation and integration work has paid off, and our IPTV services have been well received by customers.”

As mentioned, in January 2006 SureWest became the first provider in the US to introduce HDTV over IP. From its inception, the Cisco switching architecture was designed to meet the requirements for HDTV, including:

- **Bandwidth**—The Cisco switches allow SureWest to provide a bidirectional 100-Mbit/s connection to FTTH customers. This supports up to two streams of MPEG-2 HDTV at 19.4 Mbit/s each, up to six MPEG-2 SDTV streams at 3.5 Mbit/s each, and up to 20 Mbit/s for data services.
- **Small footprint**—The Cisco Catalyst 4510R chassis, with 100BASE-BX-D 48-port line cards, provides a high-density (384 ports) solution for minimizing the amount of fiber required to pass all the homes.
- **Single fiber, dual direction**—Fiber costs are further reduced. For example, a new line card gave SureWest the single-fiber capability that it required and lowered fiber costs compared to dual-fiber alternatives.

The Fruit of SureWest's Efforts

SINCE SUREWEST TRANSFERRED TRAFFIC to the new Cisco network, its service quality and customer metrics have been heading upward:

- **Disconnects**—The number of customer disconnects has been cut in half compared to equivalent traffic volume on the previous network.
- **Average revenue per user (ARPU)**—With the introduction of video services, SureWest's ARPU rose to nearly US\$100, and that doesn't include revenue from HDTV.
- **Revenue-generating units (RGUs)**—SureWest tracks the number of services purchased by customers as RGUs. With the

Multicasting with a Lot Less Bandwidth

SureWest transmits just a single copy of each channel being viewed along a given network path, saving considerably on bandwidth over providers who must carry multiple copies, even when no one in a neighborhood is watching. The key is its use of Layer 3 Internet Group Management Protocol (IGMP) for video distribution. Managing video at Layer 3 enables smarter distribution. For example, SureWest can employ network access control to ensure that customers already viewing a program do not have their experience diminished by more subscribers requesting the program than the network bandwidth can support.

In an IP NGN design such as SureWest's, the Cisco Catalyst 6509 switches in the hub and Catalyst 4510 switches in the remote terminal have the intelligence to replicate a program as needed and route it to the other remote terminals and individual homes depending solely on customer requests for the program. Because the program is distributed only when and where someone wants to see it, providers can save up to 50 percent of the bandwidth needed for video.

new network, RGUs have increased from 2.30 to 2.37 per customer, meaning that customers are signing up for more services on the network.

- **Market penetration**—Penetration rates are increasing overall as a result of SureWest's expanded service portfolio. Though still in its infancy, HDTV service sales have been strong. Of the homes purchasing this service, early reports show that there is an average of 1.6 high-definition set-top boxes per home. SureWest's network provides plenty of bandwidth for the provider to take advantage of this growing base of HDTV viewers.

By creatively bundling telephone, high-speed Internet, and digital TV services, SureWest has increased “stickiness” (customer loyalty) and attracted new customers. Subscribers are given the convenience of a single bill and a single point of contact for all services. As a result of the expanded bundles, SureWest reports that turnover rates have dropped to 1.2 percent.

The converged core and distribution network also lowers operating expenses (OpEx) for SureWest, providing an integrated foundation that streamlines provisioning and support functions. Traditional neighborhoods are being upgraded to the new IP architecture to further reduce OpEx over time, and increase the performance and responsiveness compared to the acquired RF infrastructure.

“Our top-priority goal was to improve our ability to deliver services,” says Bill DeMuth, senior vice president and chief technology officer at SureWest. “A stable, reliable core network was the essential step toward that goal. With the Cisco IP NGN, we gained a solid foundation and were able to get to market first with HDTV over IP.” ■

High Availability for MPLS

INCREASING SERVICE AVAILABILITY THROUGH FAST RECOVERY FROM NETWORK DISRUPTIONS by santiago alvarez

Converged networks allow you to run multiple services over a single network. But that also means a single network disruption has the potential of impacting multiple services, and an increasing number of customers are paying for those services. Fortunately, if you built your converged network with Multiprotocol Label Switching (MPLS), there are now some important tools to help you ensure network availability.

MPLS High Availability (MPLS HA) from Cisco rapidly restores network disruptions by reducing single points of failure in both hardware and software. It complements other network-level mechanisms such as Fast Re-Route (FRR), Interior Gateway Protocol (IGP) fast convergence, Border Gateway Protocol (BGP) enhancements, and Bidirectional Forwarding Detection (BFD).

This combination of system- and network-level resiliency mechanisms, along with proper network design and operational procedures, can help you achieve the highest levels of service availability.

SSO, NSF, NSR, and ISSU

STATEFUL SWITCHOVER (SSO) is one of the main features that enable MPLS HA. It preserves state information associated with control traffic, across an active and a standby route processor on the same system, so a switchover between route processors does not require re-initialization of the control plane state. In addition to SSO, Non-Stop Forwarding (NSF) and Non-Stop Routing (NSR) ensure the operation of the forwarding plane during a route processor switchover. These features also enable in-service software upgrades (ISSU) without interrupting traffic forwarding.

SSO, NSF, and NSR features have been available for IP protocols in Cisco IOS Software and are being extended in Release 12.2S to support additional protocols and services in MPLS networks. The

extensions include support for Label Distribution Protocol (LDP), MPLS VPN, and Any Transport over MPLS (AToM). In many cases, MPLS networks benefit already from the resiliency features available for IP protocols.

NSF/SSO: MPLS LDP, Graceful Restart

NSF AND SSO WORK IN CONJUNCTION to minimize network downtime caused by a disruption in a device main route processor. To benefit from this functionality, a device requires redundant route processors, support for the continuous synchronization of state information about these processors, and protocol extensions to maintain proper traffic forwarding during a switchover.

In normal operation, one of the processors runs in active state while the other one remains in standby state. HA-aware protocols are being constantly synchronized between route processors.

When an event in the device gives control to the standby processor, the control plane protocols perform a graceful restart while the device temporarily continues to forward traffic using the stale state information.

Figure 1 (page 68) depicts a device with dual route processors and NSF/SSO support.

As mentioned, NSF/SSO support for MPLS in Cisco IOS Software includes LDP, MPLS VPN, and AToM. LDP can recover from a protocol or session disruption without losing label bindings and while maintaining packet forwarding. LDP sessions might have been established to a directly or non-directly connected (targeted) neighbor.

MPLS VPN can retain VPN labels and continue traffic forwarding during a processor switchover, including inter-autonomous system (Inter-AS) and Carrier Supporting Carrier (CSC) configurations. This functionality requires NSF/SSO support for the routing protocol running between the MPLS and



ENHANCING AVAILABILITY
Cisco MPLS High Availability reduces single points of failure in both hardware and software.

ISSU enables full-version software upgrades while minimizing the impact on packet forwarding.

the customer network. It also requires NSF/SSO support for the IGP and label distribution mechanism in the MPLS network.

NSF/SSO for AToM maintains attachment circuit and pseudowire information across route processors to preserve packet forwarding during an LDP graceful restart.

BGP Non-Stop Routing

BGP NSF REQUIRES that peers provide assistance during the graceful restart of the protocol (NSF-aware peers). This requirement can limit your deployment, particularly for a MPLS VPN provider edge (PE) where a customer edge (CE) device might be unmanaged or not all peering devices are NSF-aware.

BGP NSR synchronizes state information across route processors and maintains BGP sessions during a switchover without any special protocol requirements on the BGP peer. It provides the benefits of NSF without imposing any special requirements on other peers. However, BGP NSR does not preclude the graceful restart procedures with those NSF-aware peers. A device using BGP NSR automatically detects NSF-aware peers and performs a graceful restart with those peers during a processor switchover.

In-Service Software Upgrades

ISSU ENABLES FULL-VERSION SOFTWARE upgrades while minimizing the impact on packet forwarding. It reduces the downtime associated with planned outages required to introduce software fixes or new features.

ISSU relies on NSF/SSO functionality and configuration. This implies that the device requires redundant route processors and the versions of software must support NSF/SSO and ISSU in particular. Current support for MPLS ISSU in Cisco IOS Software includes MPLS VPN and LDP.

Cisco IOS ISSU introduces a four-step procedure to perform the software modification. The procedure is implemented using exec commands, and no ISSU-specific configuration commands are required.

Figure 2 depicts the IOS ISSU steps:

1. The *load* step verifies the proper configuration of SSO and the existence of the new software version in the file system of both the active and the backup route proces-

sors. If these conditions are met, the standby processor is booted with the new version.

2. The *run* step forces a switchover. The standby processor running the new version becomes the active processor. NSF procedures have been active during the switchover to provide uninterrupted traffic forwarding.
3. The *accept* step confirms acceptance of the new configuration, stopping a rollback timer. This timer defaults to 45 minutes, and if it expires, the ISSU process reverts to the old software version.
4. The *commit* step completes the procedure by loading the new version of software on the now standby route processor.

ISSU provides continued packet forwarding for protocols and features that are ISSU-capable. Therefore, it is important to verify that the old and new software versions support ISSU for the protocols and features of interest. The Cisco Feature Navigator tool (cisco.com/go/fn) supplies this type of compatibility information. It classifies a pair of images as compatible, base-level compatible, or incompatible according to their support for the required high availability functionality.

Cisco Hardware Platforms

CISCO IOS SOFTWARE SUPPORTS many high availability features in a range of hardware platforms, from the Cisco 1700 Series to the 12000 Series routers. While MPLS networks benefit already from resiliency features in IP protocols,

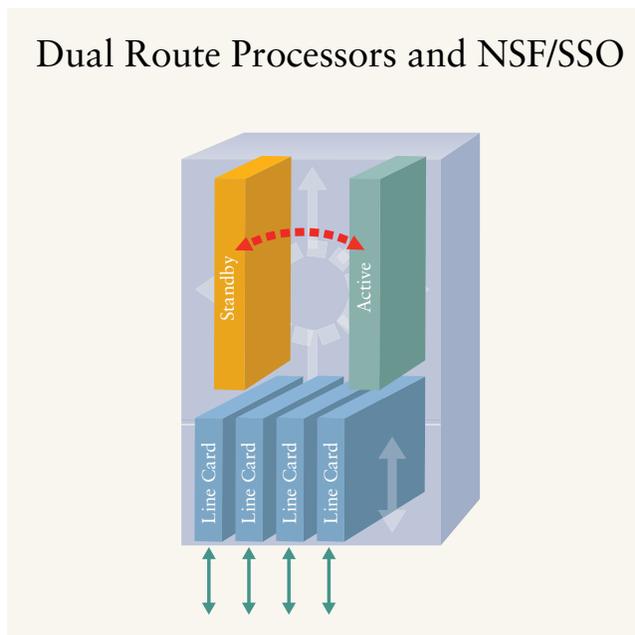


FIGURE 1 NSF and SSO work together to minimize network downtime.

Cisco IOS ISSU Procedure

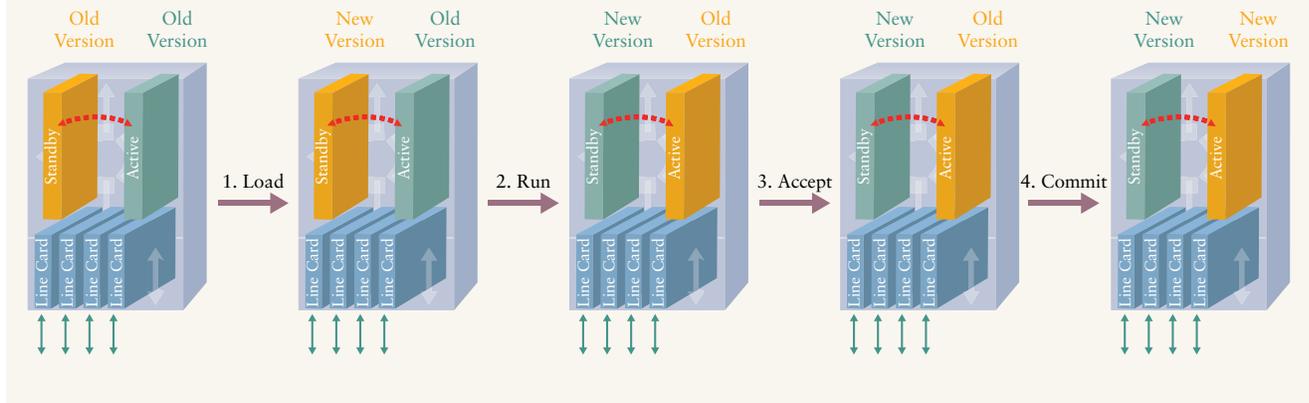


FIGURE 2 Using exec commands, no ISSU-specific configuration commands are required for this procedure.

Avoiding single points of failure is increasingly important as you combine more services and customers on a single MPLS network infrastructure.

LDP and BGP, enables a route processor to recover from service disruption without losing its LDP bindings, MPLS forwarding state, or VPN prefix information. In addition, ISSU enables network operators to perform a full-version software upgrade across redundant processors while protecting network traffic. ■

SANTIAGO ALVAREZ, CCIE No. 3621, is a technical marketing engineer in Cisco's Internet Technology Division and focuses on MPLS and QoS technologies. He can be reached at saalvare@cisco.com.

the support for MPLS protocols and features introduced in IOS Release 12.2S increase availability even further.

Specifically, MPLS LDP and VPN and AToM NSF/SSO are supported on the Cisco 7500 Series; MPLS LDP and VPN NSF/SSO are supported on the Cisco 7304 and Cisco 10000 and 7600 series; and MPLS LDP and VPN ISSU are supported on the Cisco 10000 Series. In addition, Cisco 12000 Series routers and the Cisco CRS-1 Carrier Routing System can provide some MPLS HA functionality using Cisco IOS XR Software.

Keeping Network Disruptions at Bay

MPLS NETWORKS DEMAND HIGH availability. Avoiding single points of failure is increasingly important as these networks combine more and more services and customers onto a single infrastructure.

Cisco IOS Release 12.2S supports MPLS protocols and features including NSF/SSO for LDP, MPLS VPN, and AToM. Integrating SSO and graceful restart with key protocols, such as

Further Reading

- Cisco MPLS High Availability
cisco.com/packet/183_8b1
- Cisco IOS High Availability
cisco.com/packet/183_8b2
- Cisco IOS MPLS
cisco.com/packet/183_8b3
- Cisco Feature Navigator
cisco.com/go/fn

beyond speeds+feeds

Unified WLAN in Access Layer

NEW CATALYST WLAN CONTROLLER MAKES SCALING AND MANAGING WIRELESS NETWORKS AS EASY AS WIRED NETWORKS. by gene knauer

S

AY GOODBYE TO THE COMPLEXITY of managing multiple wireless access points in different locations and integrating LAN and wireless LAN (WLAN) features. The new Cisco Catalyst 3750G Integrated WLAN Controller is the first WLAN controller that unifies wired and wireless functions in the access layer. It is a standalone, plug-and-play extension of the Cisco Catalyst 3750G Series Switch. "Many organizations want to put intelligent features at the edge of their networks, leaving the core and distribution layers to handle the movement of packets," says Chris Kozup, manager of wireless mobility marketing at Cisco. "The Catalyst 3750G Integrated WLAN Controller is a great solution for medium-sized companies and enterprise branch offices. They get switching ports and WLAN functionality and don't have to retrain personnel on a separate WLAN controller."

Continued on page 76



INTELLIGENT AND RELIABLE

The 3750G WLAN Controller works with Cisco lightweight access points, the Wireless Control System, and the Wireless Location Appliance to support mission-critical wireless data, voice, and video applications.

Catalyst 3750G, Continued from page 71

“Before, you had to manage each individual wireless access point. Now, each controller can aggregate up to 200 Cisco Aironet lightweight access points across a single campus or in multiple locations,” explains Matt Glenn, product manager for the Catalyst 3750G WLAN Controller. “From one central location, you can scale and manage a wireless network with up to 3,600 access points if you use Wireless Services Modules [WiSM], or 3,350 access points if you mix and match WiSMs and the 3750G.”

Supporting Cisco Unified Wireless Network Software

The Cisco Catalyst 3750G Integrated WLAN Controller makes full use of the new Cisco Unified Wireless Network Software Version 4.0—which brings together enterprise-class LAN and WLAN security, deployment, management, and control features for the entire Cisco product line of WLAN controllers. These include the Cisco 2000 and 4400 Series WLAN controllers, the WLAN Controller Module for Cisco Integrated Services Routers, and the Wireless Services Module for Cisco Catalyst 6500 Series Switches.

“Cisco’s strategy is to unify wired and wireless networks to allow our customers to more easily and affordably take advantage of intelligent

features behind a lot of helpful applications,” says Kozup. He points to an April 2006 study conducted by Forrester Consulting and commissioned by Cisco. The study uncovers major return-on-investment scenarios in four WLAN application categories: advanced security, guest access management, location-based services, and voice.

According to Kozup, “We’ve always promoted the increased productivity from wireless—its ability to keep users and guests connected longer—but now we can also show an array of applications that cut costs and provide additional value.”

Security Enhancements

Cisco Self-Defending Network features for intrusion detection and intrusion prevention are enabled with the new release of the Cisco Unified Wireless Network software.

“Users can come to a campus or branch office, log on and be authenticated, but if they try to access the Oracle financials database, a shun request will automatically quarantine their device,” says Glenn.

Other advanced security features include the detection and mitigation of rogue access points.

Guest Access Enhancements

Features for guest access enhancements make it easier

and less expensive than traditional adds, moves, and changes, to set up network access policies for guests onsite. Location services allow administrators in a WLAN equipped with the Cisco 2700 Series Wireless Location Appliance to track any mobile device within the WLAN, from wireless laptops to devices equipped with IEEE 802.11 radio frequency identification (RFID) tags such as hospital equipment or inventory on retail shelves. And the availability of voice-over-WLAN services can cut cellular phone costs on dual-mode Wi-Fi and cellular phones.

WLAN Grows in Popularity

The WLAN has become a mainstream feature at more than 60 percent of enterprises in North America and Europe, according to an independent survey conducted by Forrester in May 2005, “Network and Telecommunications Benchmark North America and Europe.” This figure is expected to increase to 75 percent by the end of 2006 and is also growing among smaller businesses and municipalities.

“Cisco is designing products and solutions for this culture of mobility,” says Glenn. “The convergence of wired and wireless networks will blur past distinctions, and consumers will expect wireless anywhere, anytime.” ■

Further Reading

- Cisco Catalyst 3750G Integrated WLAN Controller Data Sheet cisco.com/packet/183_9a1
- Cisco Catalyst 3750G Integrated WLAN Controller Q&A cisco.com/packet/183_9a2
- Cisco Unified Wireless Network Overview cisco.com/packet/183_9a3
- White Paper: “The Benefits of Centralization in Wireless LANs” cisco.com/packet/183_9a4
- White Paper: “Five Steps to Securing Your Wireless LAN and Preventing Wireless Threats” cisco.com/packet/183_9a5

newproduct dispatches

Core Routing

Cisco Routers and Catalyst 6500 Series Switches: New Shared Port Adapters

Twelve new shared port adapters (SPAs) offer enhanced feature support for selected Cisco core and edge routers and Cisco Catalyst 6500 Series switches. Among the new cards, the Ethernet Version 2 SPAs support new Metro Ethernet features such as QinQ termination, Layer 2 access control lists (ACLs), and bridge protocol data unit (BPDU) filtering. The OC-3, OC-12, and OC-48 packet over SONET (POS) SPAs provide modular POS interfaces. The OC-48 and OC-192 SPAs are multiprotocol-capable, with built-in support for POS Dynamic Packet Transport/Spatial Reuse Protocol (DPT/SRP) and IEEE 802.17 Resilient Packet Ring (RPR). Additional new SPAs support different port densities, link types, and network types, including Gigabit Ethernet and Fast Ethernet.

cisco.com/go/spa

Edge Routing, Access, and Aggregation

Cisco 7200 Series Routers: New Engine and Adapters

Three new products for the Cisco 7200 Series Router enhance routing capabilities

at enterprise headquarters. The Cisco 7200 NPE-G2 routing engine for the Cisco 7200VXR chassis offers performance of up to 2 million packets per second for services aggregation across a WAN or LAN at OC-3 speeds with Gigabit Ethernet support. Additional features offer threat defense, highly secure VPN connectivity, Network Admission Control (NAC), and voice/IP-to-IP gateway functions. The Cisco 7200 VPN Services Adapter provides encryption performance at up to 500 Mbit/s, supports up to 5,000 simultaneous IP Security (IPsec) tunnels, and provides other VPN security capabilities. The Cisco 7200 Port Adapter Jacket Card enables the router's I/O slot to hold a single port adapter for greater density.

cisco.com/go/servicesaggregation

Cisco 3800 Series, 2800 Series, and 1841 Model Integrated Services Routers: New VPN Advanced Integration Modules

The Cisco VPN Advanced Integration Modules (AIM) accelerate VPN performance for Cisco Integrated Services Routers and optimize IPsec and Secure Sockets Layer (SSL) Web/VPN deployments on a single platform. The VPN AIM modules offer improved

SPOTLIGHT ON

Cisco Application Control Engine for the Catalyst 6500 Series Switch

The Cisco Application Control Engine (ACE) provides large-scale virtualization of application delivery and security services. Cisco ACE consolidates the functions of multiple network devices and allows logical partitioning of the physical resources of ACE into virtual contexts. And this new module can intelligently load balance application traffic to server farms with exceptional throughput, connection setup rates, and performance scalability via software licenses rather than truck rolls.

Cisco ACE reduces the time and resources needed to deploy and manage the network application infrastructure. With the Cisco ACE virtual partitioning capability, IT administrators can tune both processing and memory resources for each virtual partition, up to 250 per module. Administrators can also guarantee resource levels and apply functions to each virtual partition. These capabilities allow administrators to quickly add or change applications, simplify system and network topologies, consolidate resources, and respond rapidly to business demand.

In addition, new application security software for the Cisco AVS 3100 Series Application Velocity System works seamlessly with Cisco ACE to add bidirectional application inspection and protection.

Role-Based Access Control (RBAC) in Cisco ACE allows each virtual partition to be managed by the appropriate business or IT team. RBAC flexibility enables faster service deployment, simplifies IT workflow, and reduces configuration errors. Additionally, networks with multiple Cisco ACE modules can be centrally managed and monitored using the new Cisco Application Networking Manager (ANM) application. Cisco ACE for the Catalyst 6500 Series is covered in greater detail on page 51.

Cisco ACE: cisco.com/go/ace

Cisco AVS software: cisco.com/go/avs

performance over the built-in IPsec encryption and software-only performance of SSL Web VPN connections. Cisco Integrated Services Routers with AIM-VPN/SSL bring the flexibility of both IPsec and SSL VPNs to small and mid-sized businesses and enterprise branch offices. Service providers can also use this combination to offer managed security services.

cisco.com/packet/183_npd3

Switching

Cisco Catalyst 2960G-48TC Switch

The fixed-configuration Cisco Catalyst 2960G-48TC Switch accelerates deployment of Gigabit to the Desktop (GTTD) by providing 48 ports of Gigabit Ethernet in a single rack unit. The switch supports integrated security features such as Network Admission Control (NAC) and sophisticated access control lists (ACLs), as well as advanced quality of service (QoS) and resiliency features. Designed for networks serving midsized businesses and branch offices, the Cisco Catalyst 2960G-48TC Switch provides four dual-purpose ports and comes with a limited lifetime warranty.

cisco.com/go/catalyst2960

Cisco Catalyst Blade Switch 3020 for HP

The Cisco Catalyst Blade Switch 3020 for HP is an integrated switch for the HP c-Class BladeSystem that dramatically reduces cable complexity. This switch also

offers a complete set of intelligent services that support security, QoS management, and availability in a server farm access environment. The Cisco Catalyst Blade Switch 3020 for HP provides 16 internal 1000BASE ports that connect to servers through the c-Class BladeSystem backplane; up to 8 external Gigabit Ethernet uplink ports; and 4 external dual-media Ethernet interfaces. Interfaces can be either 1000BASE-SX SFP or 10/100/1000BASE-T ports. The switch also provides four external 10/100/1000BASE-T ports; two of these ports can connect an additional switch.

cisco.com/go/cbs3020

Security and VPNs

Cisco NAC Appliance Version 4.0

The Cisco NAC Appliance provides Network Admission Control (NAC) capabilities that authenticate, authorize, evaluate, and remediate users and their devices before allowing network access in a standalone product. Among the new features in Version 4.0, single sign-on allows the Cisco NAC Appliance to automatically authenticate users who are already authenticated to a Windows domain, which augments existing single sign-on capabilities for VPN and wireless users. Layer 3 support for out-of-band deployments reduces the number of Cisco NAC Appliance Servers required when serving multiple locations, and a new



CISCO CATALYST BLADE SWITCH 3020 FOR HP

“Super Manager” can manage deployments of up to 60,000 online and concurrent users. Corporate asset authentication features enforce policies on devices such as printers and guest kiosks that are not associated with a single user.

cisco.com/go/cca

Cisco Security Agent Version 5.0

Through integration with Cisco NAC with Trusted QoS, Cisco Security Agent Software Version 5.0 enhances endpoint/network collaboration to increase the functionality of Cisco network and security devices and improve the delivery of mission-critical traffic when the network is under heavy load or attack. Integration with Intel Active Management Technology provides the ability to track which media was used to boot an endpoint (disk, USB, etc.) and report suspicious activity. These advances help further integrate endpoints into a Self-Defending Network.

cisco.com/go/csa

Applications Networking

Cisco Wide Area

Application Engine: New Models and Wide Area Application Services Software

Two new models for the Cisco Wide Area Application Engine (WAE) are now available. The Cisco WAE-512 serves small to midsized branch offices with a 3.0-GHz Pentium 4 processor and configurable storage up to 500 GB on two internal disks. The Cisco WAE-612 serves regional offices or larger branch offices with a 3.0-GHz Pentium D (dual-core) processor and dual-disk options that provide storage capacity up to 600 GB. Both models can run Cisco Application and Content Networking System (ACNS), Cisco Wide Area File Services (WAFS), or Cisco Wide Area Application Services (WAAS) software. The new version 4.0 of Cisco WAAS software improves the performance of any TCP-based application operating in a WAN environment. With

newproduct dispatches

Cisco WAAS, enterprises can consolidate costly branch-office servers and storage into centrally managed data centers, while still offering LAN-like service levels for remote users.

Cisco WAE:

cisco.com/packet/183_npd10

Cisco WAAS:

cisco.com/go/waas

Wireless

Linksys Gigabit ExpressCard Adapter

The Linksys Gigabit ExpressCard Adapter (EC1000) provides new notebook and PC users with ExpressCard slots for higher I/O performance, hot-swappable functionality, and a simple, reliable way to connect to a Gigabit network. The EC1000 fits in either an ExpressCard/34 or ExpressCard/54 slot, and the RJ-45 connector is integrated into the card. The EC1000 automatically negotiates the best 10 Mbit/s, 100 Mbit/s, or Gigabit network speed. The card supports IEEE 802.1p traffic prioritization and has automatic MDI/MDI-X crossover detection at all speeds. It also reduces power consumption by drawing power directly from the ExpressCard slot, which eliminates the need for an external power supply and minimizes drain on the notebook battery.

cisco.com/packet/183_npd2

Network Management

Cisco Network Assistant Version 4.0

The PC-based Cisco Network Assistant application offers centralized management and configuration capabilities for networks with up to 250 users. Cisco Network Assistant Version 4.0 supports new features for color-coding virtual LAN (VLAN) devices on the topology view, single-click activation of a Telnet session on a device, and options for printing the content of a display window. Application users can also add names or descriptive text under individual device symbols that appear in the network topology view.

cisco.com/go/cna

CiscoWorks Network Compliance Manager

The new CiscoWorks Network Compliance Manager application tracks and controls configuration and software changes throughout a multivendor network infrastructure. The application provides visibility into network change that allows IT staff to easily identify and correct trends that could impact service interruption or network stability. The CiscoWorks Network Compliance Manager also helps enterprises track and enforce

network compliance against regulatory mandates, corporate IT policies, and technology best practices.

cisco.com/go/cwnncm

Voice and Video

Cisco Unified CRM Connector

Cisco Unified CRM Connector is a customer relationship management (CRM) application that is tightly integrated with Microsoft Dynamics CRM 3.0 to support call handling by small and mid-size businesses. When a call is received, the application automatically links to the Microsoft Dynamics CRM system and provides an onscreen window of customer information on the agent's PC. New customer data or phone call information can be saved in the application to enhance future interactions. This application is available for Cisco Unified CallManager Express, Cisco Unified CallManager, and Cisco Unified Contact Center Express; it also supports customer information displays on certain Cisco Unified IP Phones.

cisco.com/packet/183_npd4

Cisco Unified IP Phone Power Injector

The Cisco Unified IP Phone Power Injector, deployed between an Ethernet switch port and a Cisco Unified IP Phone, is a single port mid-span injector with integrated power supply. It has been specifically designed and tested to support all Cisco Unified IP Phones. The power injector can support a maximum distance of 100m between a Cisco Unified IP Phone and an unpowered Ethernet switch port.

cisco.com/packet/183_npd5

Scientific Atlanta OCAP Platform

The OpenCable Applications (OCAP) Platform allows cable operators to deploy interactive cable applications across their network regardless of the set-top device, TV hardware, or system software—eliminating the need to deploy different applications for multiple device types. Based on Java technology, the Scientific Atlanta OCAP solution includes the following components, which can be bought separately or as an end-to-end integrated package: OCAP



CISCO UNIFIED IP
PHONE POWER INJECTOR

Digital Network Control System Release 4.0, OCAP middleware and operating system, and OCAP-capable set-tops.

cisco.com/packet/183_npd12

Scientific Atlanta Explorer 940 Compact Digital Only Interactive Set-Top

The Scientific Atlanta Explorer 940 Set-Top is a compact, cost-efficient, digital-only solution to help cable operators migrate toward digital simulcasting and digital broadcast networks. This product enables operators to create an “enhanced basic” service or support an introductory digital video service tier to drive conversion of basic or expanded basic subscribers. The Explorer 940 Set-Top can also help operators expand the number of additional cable outlets in a subscriber’s home, as well as increase pay-per-view and on-demand transactions.

cisco.com/packet/183_npd13

Networked Home

Linksys Wireless-N Broadband Router and Notebook Adapter

The Linksys Wireless-N Broadband Router (WRT300N) and Wireless-N Notebook Adapter (WPC300N) support the IEEE 802.11n draft specification. Linksys Wireless-N products use multiple radios to simultaneously transmit two streams of data over multiple channels, which maximizes network performance. In addition, using

multiple spatial streams allows each 20-MHz channel to contain multiple data streams for greater capacity. Linksys Wireless-N products can transmit over two available channels simultaneously, effectively creating a 40-MHz channel for applications such as high-definition video, audio streaming, online gaming, and voice over IP (VoIP).

The router and adapter also provide mixed-mode operation and backward compatibility with 802.11g and 802.11b technologies.

Router:

cisco.com/packet/183_npd6

Adapter:

cisco.com/packet/183_npd7

Linksys Network Optimizer for Gaming and VoIP

The Linksys Network Optimizer for Gaming and VoIP (OGV200) helps to eliminate network lag in data streams that are sensitive to delays. The network optimizer is installed between a home router and a broadband cable or DSL modem to monitor data traffic on the network connection. Built-in quality-of-service (QoS) techniques enable the OGV200 to automatically distinguish between data that is time sensitive and data that can be given a lower priority. As a result, online games and video have a smoother appearance, and VoIP calls have clearer voice quality.

cisco.com/packet/183_npd8

Cisco IOS Software

Cisco Intelligent Services Gateway

The new Cisco Intelligent Services Gateway (ISG) software is a policy and subscriber management solution that can significantly accelerate new-service delivery while protecting a provider’s investment in its existing broadband infrastructure. An integral, modular component of Cisco IOS Software, Cisco ISG can act as a network-based, self-contained policy management and enforcement system or interoperate with external service control systems using an array of open protocols. Cisco ISG supports IP, Ethernet, ATM, Multiprotocol Label Switching (MPLS), and VPN architectures. Other software features include automated service provisioning and the per-flow granularity and dynamic control required for voice, data, and video services. With a feature for RADIUS Change of Authorization (RFC 3576), subscriber profiles can be changed dynamically by users through a Web portal or BSS process. Cisco ISG is available for Cisco 10000 Series, Cisco 7200 Series, and Cisco 7301 routers.

cisco.com/packet/183_npd9

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 May and July 2006. For real-time announcements of the most recently released products, see “News Archive, News Releases by Date” at newsroom.cisco.com/dlls/.

ABOUT SOFTWARE:

For the latest updates, versions, and releases of all Cisco software products—from IOS to management to wireless—registered Cisco.com users can visit the Software Center at cisco.com/kobayashi/sw-center/.

productreview

Cisco ASA 5500 Series Adaptive Security Appliance

THE FOLLOWING product review is excerpted from the Networking Professionals Connection Website and was submitted by Jack Ko, security and network consultant, Trilogy Computer Systems Pty Ltd., Australia. For the full review, visit cisco.com/packet/183_9c1.

Why did you choose the Cisco ASA 5500 Series?

As a managed service provider, instant technical support and real-time monitoring are top priorities, and this is achieved in the ASA 5500 Series by deploying multiple dedicated WAN links connecting to clients. It is recommended that each client WAN link connect to a single interface on a Cisco PIX 500 Series Security Appliance. Undoubtedly, the more interfaces involved, the more complicated to manage. In fact, a maximum number of interfaces is supported per chassis. The ASA 5500 Series tackles this complicated scenario by running in multiple context mode. In this mode, the appliance is virtually partitioned into multiple security contexts for individual use. Each security context would then be managed as a separate

appliance for each client WAN connection. This is an enormous improvement in ease of management. For instance, modification of one security context will not have any impact on the other security contexts. Imagine an appliance not running in multiple security context mode; a modification for a particular client may involve modifying a few lines out of a 100-line multiclient-purpose access control list (ACL), and there is simply no guarantee that the modification has no impact on the rest of the ACL either by accident or misinterpretation.

Any specific caveats you'd like to share about the Cisco ASA 5500 Series?

There are several limitations with this new product, for example, PPPoE, PPTP, and L2TP over IPsec are not supported. Further, the ASA 5500 Series running in multiple context mode is not able to terminate any VPN connection. Also, because there is only one available slot per chassis, it is not feasible to deploy both the Advanced Inspection and Prevention (AIP) Module (responsible for IPS) and the Content Security and Control (CSC) Security



THE MOST APPROPRIATE BALANCE OF TECHNOLOGIES, PERFORMANCE, AND COST AMONG CISCO SECURITY PRODUCTS.

Services Module (responsible for Anti-X) simultaneously.

What level of experience is needed to install this product?

For individuals who have experience with Cisco PIX 500 Series Security Appliances, installation should not be too complicated, because the codes are very much alike except the interface and VPN configuration. Alternatively, Cisco introduced the Adaptive Security Device Manager (ASDM) for those individuals who have minimal or no experience with Cisco PIX 500 Series Security Appliances. The release of ASDM is another successful factor of the ASA 5500 Series. ASDM is a mature GUI device manager that offers a platform to configure, manage, and monitor the appliance.

What types of networks would benefit most from the Cisco ASA 5500 Series?

Small to medium- and enterprise-sized networks would

benefit from this product due to the flexibility of the tailored marketing packages. For instance, a small or medium-sized business can deploy the business edition of the ASA 5500 Series with the CSC module. This would provide a security gateway for ordinary business activities, as well as securing the private resources. A global enterprise can deploy the enterprise/VPN edition, which offers thousands of VPN connections for LAN-LAN and remote access. Regardless of the size of the business, the return on investment for the ASA 5500 Series Adaptive Security Appliance is extremely high. ■

WOULD YOU LIKE TO SUBMIT A PRODUCT REVIEW?
Visit cisco.com/go/product_review for details.

asktheexpert

GOT A QUESTION?

Expert Darren Douglas will answer your questions about wireless security in a live discussion forum August 28 through September 8, 2006. Join your networking peers!
cisco.com/go/askexpert/packet

Wireless Security

The Networking Professionals Connection is an online community for Cisco experts and networking colleagues. Following are excerpts from a recent Ask the Expert forum, “Wireless Security,” moderated by Cisco’s Darren Douglas. To view the full discussion, visit cisco.com/packet/183_10a1. To join other live online discussions, visit cisco.com/discuss/networking.

[Cisco Clean Access is now only supported in-band. When will it be supported out-of-band for wireless networks?](#)

Cisco Clean Access is one of the potential Network Admission Control (NAC) solutions that can be used with Cisco wireless LAN (WLAN). It complements Cisco’s NAC Framework and is useful for clients that cannot support an IEEE 802.1X supplicant or Cisco Trust Agent. Currently, out-of-band NAC is unsuitable for shared access environments. There are not suitable per-user access controls in Cisco WLAN equipment other than 802.1X and Extensible Authentication Protocol (EAP). There are no specific plans to support Cisco Clean Access out-of-band deployment with WLAN.

[We are testing a wireless controller with access points \(APs\), model 1200 converted to Lightweight Access Point Protocol \(LWAPP\). We want to use the Web authentication feature without creating the local user database on the controller. We prefer to have the controller authenticate against our RADIUS server and existing database \(LDAP\). Is this possible?](#)

Yes, it’s possible to use an external RADIUS server for Web authentication. With the WLAN Controller Version 3.2 software on the Cisco wireless controllers, it is possible to use either Password Authentication Protocol (PAP) or Challenge Handshake Authentication Protocol (CHAP), which should be compatible with LDAP. The controller checks both the RADIUS server and internal database for authentication when a RADIUS server is configured.

[Must the Cisco WLAN Controller communicate with lightweight APs only?](#)

Yes, Cisco AP1000 Series or Cisco IOS APs that have been loaded (or ordered) with the Lightweight AP software are capable of communicating with the controller. With the Wireless LAN Controller Version 4.0 software, AP 1100, 1200, 1130, 1240, and 1300 all support lightweight AP operation. In standard mode, lightweight APs direct all traffic through the controller. For bridging of traffic local to the AP, a function known as *Remote Edge AP* (REAP on the 1030 or HREAP on the 1130 and 1240) permits local bridging of traffic.

[In a multiple VLAN SSID deployment using ACS, do you recommend assigning a VLAN with IETF option 81 or Aironet VSA for SSID?](#)

The use of RADIUS IETF attribute 64, 65, and 81 is probably more flexible than SSID assignment. Technically, when using Aironet VSA for SSID, the SSID is not reassigned, but rather is restricted. For example, it is not possible to move a client from SSID A to SSID B; it is only possible to restrict a client using SSID B. However, if the AP is not connected via 802.1Q, or a simple restriction versus VLAN assignment is required, the Aironet VSA can be employed.

[Does the Catalyst 6500 Wireless Services Module \(WiSM\) support LWAPPs in both Layer 2 and Layer 3?](#)

Some documentation says it will support both, and some says it will support only Layer 3. It also states that the WiSM supports all that the Cisco 4400 WLAN Controller does minus VPN.

The WiSM supports LWAPP APs in Layer 3 mode only. The WiSM does not have an ESM module available as the 4400 does, but it does have the capability of being used with the Catalyst IPsec VPN Service Module. **■**

DARREN DOUGLAS is a technical marketing engineer in Cisco’s Wireless Networking Business Unit.

Recovering IOS on a Cisco 2600 Series Router

We recently needed to copy a Cisco IOS Software image to a router in ROM monitor mode using Trivial File Transfer Protocol (TFTP). I copied a new IOS image onto the router and it turned out to be corrupted. The old IOS image had already been erased, so the router had no IOS image available for its use. The router returned a rommon > prompt.

By using the following commands, I was able to download a good IOS image from a TFTP server and we were back up and running in about 15 minutes. This procedure uses the first LAN port (Ethernet) and can only be used for downloading a file. It cannot be used to upload a file from the router. After you finish configuring the IP address and various related parameters, use the `sync` command to copy the settings to NVRAM. That way, if you have to repeat the procedure, you do not have to reenter all the settings. You can see what settings are already in place by using the `set` command in ROM monitor mode. Below are guidelines:

```
rommon 10 > IP_ADDRESS=192.168.0.1
rommon 11 > IP_SUBNET_
MASK=255.255.255.0 rommon 12 >
DEFAULT_GATEWAY=192.168.0.2 rommon
13 > TFTP_SERVER=192.168.0.18 rom-
mon 14 > TFTP_FILE=c2600-c-mz.123-
3h.bin rommon 15 > tftpdnld
IP_ADDRESS: 192.168.0.1
IP_SUBNET_MASK: 255.255.255.0
DEFAULT_GATEWAY: 192.168.0.2
TFTP_SERVER: 192.168.0.18
TFTP_FILE: c2600-c-mz.123-3h.bin
Invoke this command for disaster
recovery only.
WARNING: all existing data in all
partitions on flash will be lost!
Do you wish to continue? y/n: [n]: y
Receiving c2600-c-mz.123-3h.bin
from 192.168.0.18 ..!!!!!!!!!!!!!!
!!!!!!!!!!!!!!... File reception com-
pleted.
```

```
Copying file c2600-c-mz.123-3h.bin
to flash.
Erasing flash at 0x607c0000
program flash location 0x60440000
rommon 16 >reset[enter]
```

M. SALAHUDDIN JAWAD, *Document World
Pakistan (PVT) Ltd, Pakistan*

PACKET ADVERTISER INDEX

ADVERTISER	URL	PAGE
AdTran	www.adtran.com/info/wanemulation	2
Aladdin Knowledge Systems	www.Aladdin.com/Cisco	IFC
Boson Software	www.boson.com/p16	A
CIPTUG	www.ciptug.org	22
Cisco Press	www.ciscopress.com	B/61
Cisco Marketplace	www.cisco.com/go/marketplace/packet0806	F
Cisco Systems Networkers	www.cisco.com/go/nw06	54
Cisco Systems	www.cisco.com/securenetworks	12
Cisco Systems	www.cisco.com/poweredy	37
Citrix	www.citrix.com/cisco	28
Colt	www.colt.net	80
eIQnetworks	www.eiqnetworks.com/cisco	14
Energis	www.energis.com	50
Extraxi Ltd.	www.extraxi.com/packet	20
GL Communications	www.gl.com	10
Global Knowledge	www.globalknowledge.com/deliver	66
Hong Kong Broadband Network	www.hkbn.net	70
NetQoS	www.netqos.com	OBC
OPNET Technologies	www.opnet.com	24
Panduit	www.panduit.com/dc05	IBC
Solsoft	www.solsoft.com/packet2	16
Spanlink Communications	www.spanlink.com	6
Sprint	www.sprint.com/business	58
Statseeker	www.statseeker.com	62
Trendium	www.trendium.com	D
Websense	www.websense.com/security	78

cache**file**

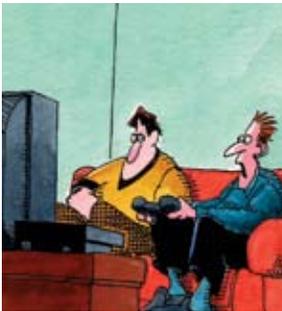
CYBER QUOTE “Computers are useless. They can only give you answers.” • Pablo Picasso, painter

Hidden Files in Computer Images

The emergence of software that enables hiding digital code in photographic images has given criminals a new way to disguise their activities, but researchers in the US are working to give law enforcement advanced tools of their own to sniff out such code. The Midwest Forensics Resource Center at the US Department of Energy's Ames Laboratory and Iowa State University are working on tools for what is called steganalysis. Steganography refers to the concealing of files in other files, such as JPEG images (the colors of a picture might be slightly changed to hide data). [Networkworld.com]

100 Million Mobile Voice-over-IP Users by 2011

While voice over IP (VoIP) is already spreading rapidly in homes and enterprises, there also will be 100 million users of mobile VoIP by 2011, according to a study by ON World. The study predicts that 36 percent of the devices that mobile users have in 2011 for accessing VoIP service will be entertainment devices such as



the 5th wave

“Wait a minute . . . this is a movie, not a game?! I thought I was the one making Keanu Reeves jump kick in slow motion.”

©The 5th Wave, www.the5thwave.com

Wi-Fi-enabled iPods. Skype will be a big winner in this market and, by 2011, will have 25 percent of mobile VoIP users throughout the world, according to the study. [TechWeb.com]

Blueprint for Invisibility Cloak

Three physicists at Duke University's Pratt School of Engineering and Imperial College London say they have developed the blueprint for an “invisibility cloak,” or fabric, to make physical objects appear as though they have disappeared when they are covered. The key to making the cloak work comes from an exotic artificial composite called “metamaterials,” which could have numerous uses, from defense applications to wireless communications. The cloak works on the principal of light as an electromagnetic wave, with a longer wavelength than X-rays and ultraviolet and shorter wavelength than infrared, microwaves, and radio waves. [TechWeb.com]

POP QUIZ

Level: CCSP
Security

Answers

QUESTIONS ON PAGE 17.

1. a, c, e
2. e
3. b
4. a
5. d

Source: CCSP SNPA Official Exam Certification Guide, 3rd Edition

WEBSITE AUTHENTICATION A “HOME-GROWN” ENTERPRISE

Website authentication takes backstage to e-mail authentication due to new initiatives using SIDF and other formats. A study conducted by Evans Data Corporation identified Web services security as a “home-grown” enterprise, stating that up to 23 percent of developers build their own authentication systems. The survey reports 22 percent use SSL, and only 9 percent use SOAP headers, both industry-standard protocols. One-quarter of the respondents cite authentication as the largest problem in Web services security. Developing authentication for an enterprise site creates resource demands, and finding IT professionals versed in Web services development is an issue for 19 percent of respondents. [Clickz.com]

Net Lingo

Cylences—Long gaps in phone conversation that occur while one person is reading e-mail or cybershopping simultaneously. [whatis.com]