CISCO SYSTEMS NETWORKING PROFESSIONALS MAGAZINE SECOND QUARTER 2006

PACKET

Connecting People Not Devices

Cisco Unified Communications Changes the Way We Communicate

Is Your Wiring Closet Ready for Real Time?

The ABCs of VPNs

4 Data Center Design Options

Lifecycle Services: Your Roadmap to Successful Deployments

CISCO SYSTEMS

contents

COVER STORY

FEATURES

KEVIN TWOMEY

departments

TECH TIPS + TRAINING

CHALK TALK

BEST PRACTICES

TROUBLESHOOTING

DESIGN STRATEGY

INFRASTRUCTURE

SWITCHING

OPTICAL

SERVICE PROVIDERS

BEYOND SPEEDS + FEEDS

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

MISHA GRAVENOR/GETTY IMAGES
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@cambeywest.com.

Welcome to the New *Packet* Magazine

IF YOU'RE A REGULAR READER OF PACKET, YOU PROBABLY ALREADY KNOW THAT WE'VE CHANGED more than the look of the magazine. We have rethought and restructured *Packet* from the ground up—improving on the best and getting rid of the rest—to bring you what we think is the industry's premier magazine for Cisco networking professionals. Whether you're a newbie or an expert, a network administrator or IT director, you're likely to find something of interest in the pages of *Packet*.

The cover story and feature articles (see the first page of the table of contents) highlight the latest technologies and industry trends of interest to the general reader, regardless of industry or technical level. Feature articles highlight the business benefits of particular networking technologies, answering the "what" and "why" but not necessarily the "how" of networking. For that, you'll want to turn to the "Departments" page of the table of contents.

Tech Tips+Training—Targeted to the newbie or the hands-on technical user who is new to a particular technology, this department offers entry points into various networking technologies, primers, tech tips, and even a quiz (see page 16). This section is also home to our highly popular "Reader Tips" column (see page 14).

Chalk Talk—For the more experienced networking professional, Chalk Talk articles are often authored by experts in the field and include in-depth discussions on routing and switching protocols, new networking standards, troubleshooting, deployment strategies, and best practices

Infrastructure—This section targets both the hands-on networking professional as well as technical managers, and focuses on systems and solutions as opposed to individual products.

Service Providers—This department highlights networking technologies, solutions, and services most relevant to telcos, ISPs, MSPs, and mobile operators.

Beyond Speeds + Feeds—This new department is all about products. It's home to our ever-popular New Product Dispatches column, as well as longer articles that focus on new or existing products from Cisco and its partners. A new Product Review rounds out this department.

There are other changes, but we'll let you discover them for yourself. We think you're going to appreciate the new *Packet*. Whether you do or don't, or fall somewhere in between, we want to hear from you. Please tell us what you think. Send an e-mail to packet-editor@cisco.com. **P**

David A. Ball

DAVID BALL
Editor in Chief
daball@cisco.com

mail

Long-Lasting Routers

I have seen numerous letters about uptime on routers in your past few issues. I couldn't help but notice that one of those readers was running Cisco IOS Software Release 12.1(2). Apparently, these folks are not keeping their routers up to date in response to all the vulnerability notices. If it were not for the vulnerability upgrades, 90 percent of my routers would be up for 3 to 4 years or more, too. It seems that a vulnerability is detected once or twice a year, making it necessary for me to upgrade more than 350 devices.

JAY E. DONOUGH
Verizon Harrisburg
Pennsylvania, USA

EDITOR'S NOTE We have received letters from other readers with valid concerns about network security. As reader Matt Carter pointed out, "In the UNIX world, uptime only reflects stability to a critical point, beyond which it just advertises to the world that you are running insecure systems."

Info on WAN Switches?

I have been a regular reader of *Packet* for many years, and would like to request that you include information and tech tips on Cisco WAN switches like the IGX and MGX. Many of my colleagues need more hands-on tips on Cisco WAN switches.

SANGAMESHWARA P.C.
Infosys Technologies Ltd.
Bangalore, India

EDITOR'S NOTE For switching support information, visit **cisco.com/packet/182_2a1**. The WAN switches area of this web page includes links to information that you may find helpful.

Littlest Packet Reader

I thought you might want to see that your readers are getting younger and younger. My daughter, Esther, at 9 months just can't put down my *Packet* magazine. She assures me that it is more than the brightly colored advertisements that she is interested in.

SCOTT DENHOLM
New Plymouth, New Zealand

Mixed Reviews for Digital

I have some concerns about your new digital version. I am a longtime print subscriber, and have enjoyed the PDF version of your magazine. I regularly download the PDF and use Adobe Acrobat extensively to make bookmarks, notes, highlights, etc. Then I send the PDF out to my network engineer peers so that they can add notes after trying out an implementation of something that they saw in *Packet*. This facilitates information transfer, and it is very easy to copy and paste information such as commands and scripts directly out of the PDF.

After reading the online digital version and downloading it to my PC, I found it impossible to use the same workflow. The format has these shortcomings:

1. Lack of tools in Macromedia (now Adobe) Flash
2. Size of file, which is 4 to 5 times larger than a PDF
3. Inability to download .EXE files from the Internet and execute them (because of corporate security policies)

I hope that Cisco continues to offer the PDF version of its fine magazine.

LUIZ DE PAULA JR.
TGS Management Corporation
Irvine, California, USA

EDITOR'S NOTE Thank you for your concerns. A PDF version is now available for back issues of *Packet* from 2004 to the present at **cisco.com/packet/182_2a2**.

I am a new subscriber to your digital edition, and I want to thank Cisco for providing information about new products and innovations. Having access to the digital version saves me from the stress that physical mailing causes in this part of the world. Cisco has once again proven to be responsive to enthusiasts like me.

KAYODE AFOLABI
A2International
Ilorin, Nigeria

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.

datagrams

Networking the Ocean

In the same way that Internet access is an integral part of nearly every research lab and office on land, extending that access to laboratories installed at sea is revolutionizing the way marine science will be conducted in the coming decades.

Internet-connected ocean observatories are already a reality at the Woods Hole Oceanographic Institution (WHOI), where oceanographers can sit in their labs ashore and communicate with instruments in the water, analyzing events such as hurricanes and earthquakes.

For its cyberinfrastructure, WHOI, a private, independent marine research, engineering and higher education organization in Falmouth, Massachusetts, deploys a Cisco-enabled LAN that includes a redundant core network of Cisco Catalyst switches and Cisco routers. A Cisco wireless network also overlays most of the Woods Hole campus. The network provides connections to both the Internet and Internet 2.

WHOI operates an ocean observatory that includes a shore station, underwater node, and offshore monitoring tower—all linked and conveyed ashore by an undersea fiber-optic cable. The institution's three ocean-going vessels also have LANs with Internet access.

WHOI Senior Scientist Alan Chave is principal scientist for the Laboratory for the Ocean Observatory Knowledge Integration Grid (LOOKING), a collaborative project of several scientific institutions for experimental wireless, optical networks, and grid technology. Through the project, communities of oceanographers will be linked via high-speed wireless and optical networks to observatories off the coasts of the US and Canada.

With LOOKING, Chave says, "Scientists worldwide will be able to access data on this growing global network, making much more information available to scientists, teachers, students, and the public."

For more on the LOOKING project and other WHOI research and activities, visit **www.whoi.edu.** P

AIR-SEA INTERACTION TOWER, Martha's Vineyard Coastal Observatory, Woods Hole Oceanographic Institution.

JAYNE DOUCETTE, WOODS HOLE OCEANOGRAPHIC INSTITUTION

How Does Cisco Do IT?

See how the latest IP networking technologies have been implemented at Cisco, by watching a new Cisco on Cisco Technology Seminar on your PC. Part of the Cisco IT@Work program, these one-hour video seminars offer insight into Cisco IT's own deployment of Cisco IP networking technologies—with leading Cisco experts relating how Cisco designs, deploys, and manages its own products and solutions. Gain insight from lessons learned at Cisco and find best practices for a range of current technologies, including application-oriented networking (AON), IP telephony, wireless LANs, data center design, security, storage-area networking, and more. Downloadable presentations on these topics are also available.

cisco.com/packet/182_3d1. P

Convenient New Linksys Cordless Internet Phone with Skype

Skype Internet calling is becoming as convenient as regular phone service with the recent introduction of the Linksys Cordless Internet Telephony Kit (CIT200)—a Skype phone from Linksys.

The sleek Linksys CIT200 uses a wireless DECT base station connected to the PC via a USB port to untether the handset from the computer. The Linksys CIT200 is a pleasure to hold, has a vibrant color screen, and offers excellent audio quality, even when using the built-in speaker phone. The menus are easy to navigate, allowing you to manage Skype remotely from the handset, which supports a variety of features for calling and messaging. Find more information at **cisco.com/packet/182_3c1.** P

SLEEK AND PRACTICAL
The Linksys CIT200 includes a cordless handset, charger, and base station that connects to a USB port on a PC.

FREE SMART BUSINESS GUIDE FOR SMALL AND MIDSIZED BUSINESSES

LEARN HOW THE networks from Cisco can deliver new ways of doing business for your small or midsized business. Download your Smart Business Guide, an in-depth look at how Cisco Smart Business Communications can help you take advantage of the network to strengthen your business in new ways. To learn more, visit **cisco.com/youinc.** P

Recently Announced Cisco Acquisitions

Acquired		Employees	Location
SyPixx Networks, Inc.	Network-centric video surveillance software and hardware that enable existing analog video surveillance systems to operate as part of an open IP network. This acquisition will enable Cisco to deliver video surveillance as part of an intelligent converged environment. As a result of this acquisition, physical security will become a new emerging technology area for Cisco.	26	Waterbury, Connecticut, USA

The ABCs of VPNs

A COMPARISON OF VIRTUAL PRIVATE NETWORKS by mark lewis

Virtual private networks have long allowed the provisioning of private network services across a shared public infrastructure such as the Internet or other WAN backbone. Over the years, however, the proliferation of VPN protocols and technologies, has made it challenging to differentiate between various VPN types and understand how they compare. • A wealth of technologies enable both site-to-site VPNs and remote access VPNs. Site-to-site VPNs allow connectivity between fixed, geographically dispersed sites (such as a head office and branch offices). Remote access VPNs allow mobile or home-based users to access an organization's data or other resources. • In both types of VPN, tunnels are created between locations by encapsulating users' traffic within other packets. For this to happen, the tunneled (encapsulated) traffic gains additional header(s), tags, or labels that correspond to the tunneling protocol. Through the encapsulation with an additional packet header, tags, or labels, a VPN gateway, customer edge (CE) device, or provider edge (PE) device can distinguish among customers or users. Therefore, tunneling keeps each organization's or user's traffic separate, and thus private, from other traffic flowing on a shared network.

Data encryption can be added to the mix to scramble data in transit for extra security. This happens frequently when the WAN used for VPN transport is the public Internet, which does not fall under the operational control of a single service provider and is thus considered an "untrusted" network. Many organizations that use network services offered across a carrier's backbone, such as Frame Relay, ATM, and Multiprotocol Label Switching (MPLS)-based services, opt not to use encryption, because carriers use Layer 2 virtual circuits, labels, or similar technologies to separate customer traffic. These VPNs are called *trusted VPNs*.

In a *secure VPN*, by contrast, customer data traffic is authenticated and encrypted. Examples of secure VPNs are IP Security (IPsec) VPNs, Secure Sockets Layer (SSL) VPNs, and Layer 2 Tunneling Protocol (L2TP) VPNs secured using IPsec.

In site-to-site VPNs, data traffic is either tunneled between CE routers or between the public network service operator's PE routers. The difference is that in a CE-to-CE configuration, the VPN tunnels and associated security extend across the WAN to the customer premises. In a PE-to-PE configuration,

GETTY IMAGES

TUNNELING ALLOWS organizations to keep data ccnfidential by erecting a virtual private network.

the tunnels are confined to the interior of the shared service provider network.

Service provider–provisioned site-to-site VPNs can be used to tunnel either Layer 2 or Layer 3 protocols. Technologies such as L2TP version 3 and Any Transport over MPLS (AToM) can be used to tunnel a variety of protocols such as Point-to-Point Protocol (PPP), Frame Relay, ATM, and Ethernet, while IEEE 802.1Q tunneling (Q-in-Q) can be used to tunnel Ethernet only. IPsec, Generic Routing Encapsulation (GRE), and MPLS are commonly used to tunnel IP or other Layer 3 protocols.

• Layer 2 site-to-site VPNs—These allow data-link-layer connectivity between separate sites and can be provisioned between switches, hosts, and routers. Communication is based on Layer 2 addressing, and PE devices forward customer network traffic based on incoming link and Layer 2 header information, such as a MAC address or a Frame Relay Data Link Connection Identifier (DLCI).

AToM (Draft Martini) and L2TPv3 pseudowires (emulated circuits) can provide the Layer 2 protocol point-to-point transport necessary to enable Virtual Private Wire Service (VPWS) Layer 2 site-to-site VPNs. AToM transports Layer 2 frames across an MPLS network, with the Label Distribution Protocol (LDP) signaling pseudowire capabilities and attributes, including a virtual circuit (VC) label that is used to distinguish Layer 2 traffic associated with different customer last-mile attachment circuits. L2TPv3 provides similar signaling capabilities, and uses "Session IDs" and cookies to associate customer Layer 2 traffic with appropriate attachment circuits. While VPWS Layer 2 VPNs provide point-to-point connectivity, Virtual Private LAN Service (VPLS) and IP-Only LAN Service (IPLS) Layer 2 VPNs are required for multipoint (any-to-any) connectivity. VPLS and IPLS take advantage of technologies such as MPLS and L2TPv3 pseudowires, as well as IEEE 802.1Q tunneling to allow multipoint Ethernet and IP-only connectivity respectively.

• Layer 3 site-to-site VPNs—These interconnect hosts and routers at separate customer sites. Customer hosts and routers communicate based on network-layer addressing, and PE devices forward customer traffic based on the incoming link and on the addresses in the IP header.

PROTOCOL BUFFET Many types of VPNs have evolved to serve different applications and to enable newer IP and MPLS networks to transport legacy user traffic.

Common VPN Alternatives

VPN Type	Application	Attributes	Who Provisions
BGP/MPLS IP VPN (RFC 4364 / RFC 2547bis)	Site-to-site; multipoint	Layer 3; typically enables full-mesh connectivity (hub-and-spoke, partial-mesh, and extranet connectivity can also be provisioned)	Service provider*
Draft Martini (AToM)	Site-to-site, point-to-point	Layer 2; allows the point-to-point transport of Layer 2 traffic over an MPLS backbone. Providers can consolidate legacy and IP/MPLS network infrastructure	Service provider
VPLS/IPLS	Site-to-site, multipoint	Layer 2 (Ethernet or IP only) transport; enables full-mesh connectivity	Service provider
GRE	Site-to-site; point-to-point	Layer 3; transports legacy protocols and IP over an IP backbone	Service provider or enterprise
IEEE 802.1Q Tunneling (Q-in-Q)	Site-to-site	Layer 2; segregates customer Ethernet traffic by adding an extra 802.1Q tag to the beginning of the Ethernet VLAN header	Service provider
IPsec	Site-to-site or remote access; point-to-point tunnels; usually used across the public Internet	Layer 3; encrypts or authenticates IP traffic between security gateways or hosts	Service provider or enterprise
Layer 2 Tunneling Version 2 (L2TPv2)	Remote access	Layer 2; can encapsulate and tunnel Point-to-Point Protocol (PPP) over an IP backbone	Service provider or enterprise
L2TPv3	Remote access; site-to-site	Layer 2; encapsulates Layer 2 protocols over a point-to-point IP connection	Service provider
SSL VPN (WebVPN)	Remote access	Layer 4-7; no client software required, so users can deploy dynamically	Enterprise or service provider

*Presumes a shared network service; some very large enterprises build their own private MPLS networks

BGP/MPLS IP VPNs, based on IETF RFC 4364 (formerly RFC 2547bis), is a Layer 3 VPN technology typically provisioned by service providers in which the PE devices maintain separate routing and forwarding tables for each VPN. Customer routes are advertised between PE devices using the Multiprotocol Border Gateway Protocol (MP-BGP), and customer address space and routes are distinguished using BGP attributes.

An alternative to the BGP/MPLS VPN is the Virtual Router VPN, based on an IETF draft called "Network based IP VPN Architecture using Virtual Routers" (**cisco.com/packet/182_4b1**). Here, completely separate logical routers are maintained on the PE devices for each VPN. Each logical router maintains its own entirely separate routing protocol instances.

Remote Access VPNs

PROTOCOLS AND TECHNOLOGIES that enable remote access VPNs include IPsec, L2TP, and SSL/TLS. IPsec and (client-initiated) L2TP remote access VPNs require client software to be installed on remote user devices, and allow connectivity and access to central site resources similar to that experienced by users physically located at the central site.

SSL VPNs (called WebVPNs in Cisco vernacular) can be quickly deployed, because they require no installation and maintenance of special client-side software and they offer application-layer access control. This dynamic capability enables mobile users and users in disaster situations to access the network from any browser-enabled client device with Internet access.

SSL provides digital certificate–based authentication, integrity checking, and confidentiality. Transport-layer confidentiality is supported through secret key cryptography.

When using a clientless SSL VPN, the encrypted connection between the remote user and the VPN gateway happens via a Web connection at the application layer. This trait allows enterprises to set up granular rules for which applications a given user can access, depending on what type of connection is in use, the user's access rights, and so forth. This differs from the Layer 3 network tunnel of an IPsec VPN, which gives all authenticated users access to all resources unless specific access policies are set up. If additional resource access is required with SSL VPNs, however, dynamically downloaded client software can be installed on remote users' workstations, and will enable "full" network access similar to that provided by IPsec remote access VPNs.

For organizations that require both SSL VPNs and IPsec VPNs, Cisco VPN gateways, including the VPN 3000 Series Concentrator and the newer Cisco ASA 5500 Series Adaptive Security Appliance, will terminate both types. The ASA 5500, which was recently enhanced, combines these VPN services with firewall, intrusion prevention, and network antivirus services in a single appliance (see story, page 48).

Who Provisions?

CUSTOMER-PROVISIONED, site-to-site VPNs are configured on CE devices such as routers, firewalls, and VPN concentrators. In this case, tunnels are configured between CE devices in the VPN, and customer data traffic is sent over these tunnels.

In CE-based Layer 3 VPNs, PE devices do not participate in (and are unaware of) customer network routing. Rather, they forward customer traffic based on globally unique addressing. In this case, tunnels are configured between CE devices using protocols such as IPsec and GRE. This configuration is sometimes known as an *overlay VPN*. Examples of overlay VPNs include those built using Frame Relay or ATM virtual circuits, as well as those built using GRE or IPsec tunnels.

When the PE devices do participate in CE routing, the configuration is known as a *peer VPN*. In peer VPNs, routes are exchanged between CE devices and PE devices. Peer VPNs are provisioned by the service provider.

Attributes to Consider

WHEN COMPARING BOTH PROVIDER- and customer-provisioned site-to-site VPNs, consider these factors:

• Connectivity—Is point-to-point or multipoint connectivity inherent to the VPN? For example, BGP/MPLS IP VPN services are inherently multipoint (assuming a "typical" configuration), while IPsec and GRE VPNs are point-to-point technologies. These characteristics affect the ease of deploying a range of network topologies, such as full mesh, hub-and-spoke, and partial mesh. Point-to-point VPNs require building the topology out of multiple point-to-point tunnels, while multipoint VPNs are inherently meshed.

• Geographic reach—Is geographic reach limited to a service provider backbone, or can it be extended across the Internet? MPLS-based VPNs, including BGP/MPLS IP VPNs and Draft Martini VPNs, generally limit traffic to within the perimeter of one or more carrier's backbone networks. By contrast, Layer 3 IPsec and GRE networks can traverse private IP and public Internet links.

• Security—Transporting MPLS Layer 3 and L2TP-based VPN traffic over IPsec boosts their security from good to excellent. L2TPv3 on its own can employ tunnel authentication, and a cookie enables resistance to blind insertion attacks.

• Inherent multicast support—Can multicast traffic be

Further Reading

- *Comparing, Designing, and Deploying VPNs*
  cisco.com/packet/182_4b2
- *Troubleshooting Virtual Private Networks*
  cisco.com/packet/182_4b3
- Cisco ASA 5500 Series
  cisco.com/packet/182_4b4
- Cisco VPN 3000 Concentrators
  cisco.com/packet/182_4b5

natively supported across the VPN? The answer is "yes" with Layer 2 VPNs; BGP/MPLS IP VPNs and IPsec Layer 3 VPNs require the use of GRE tunnels or the deployment of technologies such as multicast VPNs (MVPNs) in the case of BGP/MPLS IP VPNs or virtual tunnel interfaces (VTIs) in the case of IPsec VPNs.

• Quality of service (QoS) support—How does the technology differentiate levels of service for voice, video, and data applications? MPLS networks typically depend on priority markings in the experimental (EXP) field in the MPLS shim header. Hard QoS guarantees, including both transmission quality and service availability, additionally require support. In IPsec, L2TP, or GRE VPNs, traffic differentiation relies on markings in the type of service (ToS) field of the outer IP header.

There are many different types of VPN protocols and technologies which can be broadly classified as either site-to-site or remote access VPNs. Most of these VPN types are supported in Cisco IOS Software running on Cisco routers. In enterprise (customer)-provisioned VPNs, users can choose to offload both IPsec and SSL VPN termination operations from the router to Cisco VPN gateway equipment, such as the ASA 5500 and VPN 3000 series.

VPN technologies have evolved to solve different problems. Site-to-site Layer 2 VPN technologies allow the tunneling of Layer 2 protocols between PE or CE devices, and enable service providers to consolidate legacy and IP/MPLS networks, as well as allowing them to deploy newer Ethernet MAN/WAN services. Site-to-site Layer 3 VPN technologies, in contrast, emphasize strong security and low relative cost in the case of IPsec; or any-to-any IP connectivity, simplified customer WAN routing, and QoS in the case of BGP/MPLS IP VPNs. Remote access VPN technologies such as IPsec and SSL allow secure access for mobile or remote users to an organization's data or other resources. **P**

MARK LEWIS, CCIE NO. 6280, is technical director of MJL Network Solutions. He is also an active participant in the IETF, a member of the IEEE, and a certified Cisco Systems instructor. Lewis is the author of *Comparing, Designing, and Deploying VPNs* and *Troubleshooting Virtual Private Networks,* both published by Cisco Press. He can be contacted at mark@mjlnet.com.

Deploying Cisco Security Agent

OUT OF THE LAB AND INTO PRODUCTION

by chad sullivan

The Cisco Security Agent threat protection software for server and desktop computing systems surpasses the functionality of conventional endpoint security solutions. It does so by analyzing behavior to identify and prevent malicious activity before it can occur, rather than relying exclusively upon signature matching to ferret out unwanted behavior. Cisco Security Agent's unique architecture correlates behavior occurring on the end systems by monitoring clues such as file and memory access, process behavior, Component Object Model (COM) object access, and access to shared libraries. This function provides significant returns to network implementers by preventing unwanted outages that result from worm and viral infections. Cisco Security Agent protects against continuous zero-day threats and known exploits and supplies a mechanism to enforce written security policies (see figure, page 12).

The most common challenges faced during the Cisco Security Agent implementation relate to computer support practices. It is important to a successful deployment that the implementation team tasked with both policy definition and product installation has a solid understanding of how the organization interacts with the systems on a daily basis. In addition, implementers should arm themselves with a thorough understanding of the product's capabilities and a well-defined scope of what the organization is attempting to accomplish. Often, an implementation team begins by testing a product in its own lab and on IT staff PCs. While this seems rational, it can provide a skewed informational baseline that leads to inaccuracies in project timeline scoping and product interaction and support requirements. Rather than starting out on the systems of very technically savvy users, it can be much more productive to install the agents on the systems of "everyday" users—especially users who either generate many support cases or are likely to use applications that interfere with base Cisco Security Agent policy. With the product in the hands of real-world users early on, you are more likely to understand the implementation and support costs. Implementation consists of three phases: *preparation, agent software installation,* and *policy development.*

Preparation Phase

TAKING THE NECESSARY time for proper planning leads to significant time savings down the road. During this phase, inventory the software in your computing environment and understand it in great detail, including items such as operating systems and patch revisions used; applications expected (Microsoft Office, BigFix Agents, personal firewalls); e-mail applications; antivirus, VPN, firewall, and other security applications in use; software/patch deployment mechanism (BigFix Agent, MS SMS, login script, manual); chat and file-sharing mechanisms allowed; and internally developed software.

Building a list of applications that are allowed and expected in your environment will help when tuning the Cisco Security Agent software to allow trusted and supported applications to function as desired. Additionally, the implementation staff should understand how each of these applications and operating systems are maintained and supported on a daily basis by the help desk and the support technicians.

Support procedures are the most often overlooked aspect of Cisco Security Agent implementation. Early in most implementations, the software is installed in test mode, where the software informs the administrators what Cisco Security Agent actions would have been taken without any impact to the local user. Having this information is advantageous but can also cause the administrator to overlook the

TAKING THE necessary time for proper planning will lead to significant time savings.

AVERTING ATTACKS The Cisco Security Agent is specifically designed to protect against new types of attacks where there is no known signature.

deployment impact on a key function: technical support. Technicians who assist users on a daily basis are not accustomed to interacting with a preventive piece of software such as Cisco Security Agent. For this reason, the technicians themselves or a representative familiar with their daily tasks must be included in the project from the beginning. It is important to understand how technicians interact with systems while performing their jobs on site or using remote-control software.

In addition to requiring connectivity, technicians need the ability to interact with the software agent during installation, upgrades, and troubleshooting. Because many Cisco Security Agent implementations assume that the end user should not be able to view or interact with the local Cisco Security Agent software, steps must be taken to ensure that tasks that would normally require interaction, such as software installation, either do not require it or that interaction can be obtained if necessary. A common way to address the need for temporary interaction is to allow technicians to log into the system using their own user ID, which triggers a local user-based state in Cisco Security Agent. This state temporarily allows the agent to become interactive until the technician completes the necessary procedures and logs off. Optionally, the administrator can allow end users to view and use an advanced user interface for endpoint control over some of the policies enforced on the local workstation, such as personal firewall rules. If that is not desired, you can elect to not deploy this interactive capability or to enable its use for a defined subset of users.

The need for remote connectivity support procedures might not surface during a short pilot. But it should be taken into account so that efficient user support mechanisms can be put in place. It is neither always required nor productive to send a technician to an end user's desk when the technician can take corrective action from afar. Remote support tools could include various software packages such as Microsoft Remote Desktop for temporarily controlling a system, Microsoft Management Consoles for viewing remote system settings and resources, and various options for performing tasks such as modifying the remote system registry or starting processes and services on remote systems.

None of these tools will function as expected until the Cisco Security Agent administrator has allowed these actions to occur when originating from remote systems, a task that would commonly be viewed as an attack vector used by many worm variants today. To override the base Cisco Security Agent functionality that prevents this type of remote connection, the administrator must allow the specific access by tuning the policy deployed to the protected systems.

A best practice for Cisco Security Agent in the preparation phase is to identify how systems are automatically patched as part of normal system support cycles. This automation also applies to the installation of new software and software upgrades for existing packages. Many organizations use agent-based software deployment mechanisms as well as built-in mechanisms such as Microsoft Windows Update to help ensure that software and applications can be distributed without manual intervention. Because these updates will occur as necessary and should not involve any user interaction, you must be certain to allow the installation applications the abilities required to complete the installation tasks, such as installing drivers and system files. Without tuning these services during the pilot period, administrators could find themselves in a situation where the Cisco Security Agent protection also prevents new software that includes security updates from being deployed.

Continued on page 15

Cisco Security Agent, Continued from page 12

Agent Software Installation

THIS PHASE OF CISCO SECURITY AGENT deployment can be the most straightforward and simple to complete. The issues that might arise are no different than those encountered in any other software product installation. You should perform the same pre-installation tests on the Cisco Security Agent just as you would test any other software product before installation. The only difference is that this software interacts with the local operating system and applications in a way that is not common to many other applications, such as word processing.

Agents are initially installed on endpoints via an executable install file. Methods of installation include locally executing the EXE file manually or via many other scripted and automated installation procedures, such as an enterprise software installation system. Before you can install an agent kit on a workstation, you must create the appropriate initial modules, policies, and rules that the agent will use. Then you must define the group and attach policies to it. From there, you will create the agent kit and define a few installation kit parameters.

Policy Development

POLICY DEVELOPMENT often takes the longest and often overlaps with other phases. It can overlap because you will begin to tune Cisco Security Agent software policies during the very early stages of testing both in the production environment and in the lab, and some of that work will spill over during the production implementation.

Policy tuning is where skilled technicians interpret event messages from the software agents throughout the deployment and decide how to best tune the policy. These individuals should be aware of or have the capacity to understand how operating systems function and also be intimately familiar with how the software in your deployment should work. Their purpose is to implement a policy that is in line with your organization's written security and acceptable-usage documents, as well as with any other guidelines that have been developed during the early stages of the project.

It is common for the individuals tasked with this portion of the project—those who are the most operating system- and application-aware—to focus more on keeping systems operational than secure. There is a balance where systems continue to be just as functional as they were before the implementation yet are unlikely to be impacted by the next viral outbreak. Keep the project goal in mind: to secure systems and data without a major impact on users or training requirements. The only effective security controls are ones that are silent and do not cause additional workload. Users will seek to circumvent your security controls if you make their daily jobs more difficult.

Further Reading

- Cisco Security Agent Deployment Best Practices Guide
  cisco.com/packet/182_4a1
- Understanding Cisco Security Agent Components and Installation
  cisco.com/packet/182_4a2

For policy and rule changes created on the Cisco Security Agent Management Center to take effect, Cisco Security Agent must have the ability to contact and communicate with the Management Center over the network. This communication path can use any transport between the agent and management station as long as there is end-to-end IP reachability for the duration of the connection. For the agent to request the update or transmit event log messages, the agent attempts to resolve the Cisco Security Agent Management Center's IP address using Domain Name System (DNS) or any other local resolution means available, such as the local hosts file. This information must be correct to facilitate both a successful connection and to verify the certificate used in the Secure Sockets Layer (SSL) communication. SSL is used for management-to-agent interaction to ensure an authenticated and encrypted communication of the updates and event transmissions. Within an enterprise, name resolution is not typically an issue; however, if you have systems that will roam around disparate networks, you must be certain that the machine resolves the correct address and that the Cisco Security Agent Management Center server is reachable from those locations.

Solving a Complex Problem

CISCO SECURITY AGENT is a complex product that can accurately enforce your written security policies and prevent the next zero-day exploit from wreaking havoc throughout your organization. It is complex only because the problem it is attempting to solve is also very complex. This doesn't mean, though, that the product implementation is also complex. With proper preparation, you will not only be able to effectively roll out and support your Cisco Security Agent software, but your support team will better understand the software and systems they support. This is a beneficial byproduct of every Cisco Security Agent implementation that is reason enough to get this software out of the lab and into production. P

CHAD SULLIVAN, CCIE No. 6493, is a security consultant and co-owner of Priveon, Inc., a security and networking consulting firm. Prior to founding Priveon, he was a security consulting systems engineer at Cisco. Author of the Cisco Press book, *Cisco Security Agent*, Sullivan is currently authoring *Advanced Host Intrusion Prevention with Cisco Security Agent*. He can be reached at chad.sullivan@priveon.com.

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.

Configuration

Displaying Extensions in the Corporate Directory

We use the Corporate Directory Service for all Cisco IP phones. Our Active Directory contains a 7-character string, xxxyyyy, often with trailing spaces in the field. When an end user looks up a name in the Corporate Directory from a phone, they receive the full DID number. Users must use the EditDial button to remove the first three digits, because all of our internal extensions are four digits.

To strip digits within the Corporate Directory Service, edit the *xmldirectory.asp* file:

1. Before editing the original *xmldirectory.asp*, make a backup copy of C:\Ciscowebs\IPPhoneServices\CCMCIP\xmldirectory.asp.
2. Open the original in a text editor, such as Wordpad.
3. Cut and paste the following text to the Function section (near the beginning):

```
function rightTrim(sString) {
        while
(sString.substring(sString.length-
1, sString.length) == ' '){
        sString = sString.sub-
string(0,sString.length-1);}
return sString;}

function Right(str, n){
        if (n <= 0)
        return "";
        else if (n >
String(str).length)
        return str;
        else {
        var iLen =
String(str).length;
        return String(str).sub-
string(iLen, iLen - n);}}
```

This defines two javascript functions. The first function strips all trailing spaces from the string. The second function keeps a specified number of characters on the right side of the string and strips the remainder.

Resetting the Factory Default for IP Phones

It is often necessary to reset the factory default for a Cisco IP phone configuration. To do this, power the IP phone off and on. During the reboot, press and hold the pound key (#) during the entire reboot. The IP phone will display "Reset key sequence detected." Enter 123456789*0#, followed by 2, to erase the entire configuration. Use this procedure to completely remove a CTL file when moving an IP phone from a Cisco CallManager cluster configured in Mixed Mode, to a Cisco CallManager cluster configured in NonSecure Mode.

MASSIMO CUCCHI, *Dimension Data, Milan, Italy*

4. Add the text below to the following section and change 4 to the number of digits you need.

```
var telnum = Right(rightTrim
users[i].TelephoneNumber), 4);

for (var i = currentStart-1; (i <
listCount) && (count <
maxListSize); i++)
                {
        count ++;
        var fullname =
users[i].LastName + ", " +
users[i].FirstName;
        if (fullname.length >
63){fullname =
fullname.slice(0,63);}
                fullname = full-
name.replace(/&/g, "&amp;");

Response.Write("<DirectoryEntry>\n"
);

Response.Write("<Name>" + fullname
+ "</Name>\n");

Response.Write("<Telephone>" +
users[i].TelephoneNumber +
"</Telephone>\n");
Response.Write("</DirectoryEntry>\n
");
```

5. Change the line in bold above to:

```
Response.Write("<Telephone>" + tel-
num + "</Telephone>\n");
```

Thanks to Chris Martin for helping create this javascript.

RON SMITH
Florida Community College
Jacksonville, Florida, USA

Determining Modules Installed in a Router

To get a quick overview of which modules are installed in your router, use this command:

```
router#show diag | include IC|NM-|Serial|FRU
        PCB Serial Number         : FOC09162SSN
        Chassis Serial Number     : FTX0926A00R
        Product (FRU) Number      : CISCO3845
        PCB Serial Number         : FOC09182ET4
        Product (FRU) Number      : CISCO3845-MB
        Chassis Serial Number     : FTX0926A00R
        WIC Slot 0:
        VIC2 - BRI-NT/TE Voice daughter card (2 port)
        PCB Serial Number         : FOC091537HG
        Product (FRU) Number      : VIC2-2BRI-NT/TE=
        WIC Slot 1:
        VIC2 - BRI-NT/TE Voice daughter card (2 port)
        PCB Serial Number         : FOC0910103X
        Product (FRU) Number      : VIC2-2BRI-NT/TE=
        WIC Slot 2:
        PCB Serial Number         : FOC092150QX
        Product (FRU) Number      : HWIC-4ESW
        NM-1T3/E3 (clear/subrate) Port adapter, 1 port
        PCB Serial Number         : FOC09140Y4S
        Product (FRU) Number      : NM-1T3/E3=
```

MICHAEL KUNZE, *SYCOR GmbH, Göttingen, Germany*

Troubleshooting

Troubleshooting the Committed Access Rate on Routers Running WCCP

After implementing Committed Access Rate (CAR) on a router's packet over SONET (POS) OC-3/STM1 interface for a specific IP pool, the traffic was not being restricted on that IP pool. The IP pool used more bandwidth than was restricted in the CAR on the POS interface. The IP pool was attached to the router LAN. The issue was that the IP pool traffic was going through the Web Cache Communication Protocol (WCCP) cache servers, so the restriction was not being implemented on the traffic. After excluding the IP pool from the WCCP access list on the router, the IP pool did not receive more bandwidth than was defined in the rate limit.

Configuration example (before):

```
!
interface GigabitEthernet0/1
 ip address 10.10.10.1
255.255.255.192
 ip wccp web-cache redirect out
!
Interface pos 2/1
 Ip address 172.16.10.1
255.255.255.252
 Rate-limit input access-group 101
2048000 256000 256000 confirm-action
transmit exceed-action drop
 Rate-limit input access-group 101
2048000 256000 256000 confirm-action
transmit exceed-action drop
```

```
 ip wccp redirect exclude in
 ip wccp web-cache redirect out
!
!
Remark access list for the abc
client for 2 MB Bandwidth
access-list 101 permit ip
10.10.10.128 0.0.0.7 any
access-list 101 permit ip any
10.10.10.128 0.0.0.7
!
Remark WCCP Access list
access-list 114 permit ip
10.10.10.0 0.0.0.255 any
!
!
ip wccp version 1
ip wccp web-cache redirect-list 114
```

With the above configuration, the IP pool in access list 101 used more than 2 Mbit/s bandwidth. After excluding the IP pool from the WCCP access list (see below), no more than 2 Mbit/s bandwidth is consumed.

```
interface GigabitEthernet0/1
 ip address 10.10.10.1
255.255.255.192
 ip wccp web-cache redirect out

Interface pos 2/1
 Ip address 172.16.10.1
255.255.255.252
 Rate-limit input access-group 101
2048000 256000 256000 confirm-action
transmit exceed-action drop
 Rate-limit input access-group 101
2048000 256000 256000 confirm-action
transmit exceed-action drop
 ip wccp redirect exclude in
 ip wccp web-cache redirect out
!
!
```

POP QUIZ

Level: CCNA

IP ROUTES

1. How do you configure a static route on a Cisco router?
2. What is split horizon?
3. What is administrative distance (AD)?
4. How do distance vector routing protocols keep track of changes to the internetwork?
5. What is convergence?

ANSWERS: SEE PAGE 82.

Source: CCNA Flash Cards and Exam Practice Pack

Cable Diagnostics on the Cisco Catalyst 6500 Platform

This command allows you can run basic cable diagnostic tests on a Cisco Catalyst 6500 Series Switch:

```
Test cable-diagnostics tdr 13/5 (catOS) - minimum version 7.6
Test cable-diagnostics tdr interface g13/5 (IOS) - minimum version
12.2(17d)SXB
Catalyst> (enable) show port tdr 13/5
TDR test last run on Wed Sep 28 2005, 16:00:59
```

Port	Speed	Local pair	Pair length	Remote pair	Pair status
13/5	auto	Pair A	15 +/- 2 meters	N/A	Open
		Pair B	16 +/- 2 meters	N/A	Open
		Pair C	15 +/- 2 meters	N/A	Open
		Pair D	15 +/- 2 meters	N/A	Open

This is only supported on newer blades (WS-X6148-GE-TX, WS-X6548-GE-TX, etc.)

AURELIO DESIMONE, *CCIE No. 10267, Refco Group Inc., Chicago, Illinois, USA*

```
Remark access list for the abc
client for 2 MB Bandwidth
access-list 101 permit ip
10.10.10.128 0.0.0.7 any
access-list 101 permit ip any
10.10.10.128 0.0.0.7
!
Remark WCCP Access list
access-list 114 deny ip
10.10.10.128 0.0.0.7 any
access-list 114 permit ip
10.10.10.0 0.0.0.255 any
!
!
ip wccp version 1
ip wccp web-cache redirect-list 114
```

ZEESHAN AHMED
WorldCALL Multimedia
Lahore, Pakistan

Packet Shaper Connections and Slow Internet

My company recently installed a Packeeter 6500 bandwidth and packet shaper device behind our Cisco PIX 525 Firewall and Cisco Catalyst 2950 Switch. A few hours later we noticed the Internet connections were slow. We realized our problem was that the bandwidth device was configured to the maximum 100 full duplex and our Cisco network devices were using gigabit. In order for the device to work properly the ports must also be configured to 100 full duplex.

Cisco 2950 Switch:

```
c2950(config)#set port speed
<mod_num/port_num> 100
```

Cisco PIX 525 Firewall:

```
pix525(config)# interface ethernet1
100full
```

The device is now working fine.

ABU EMRAN ABU BAKAR
International Islamic University Malaysia
Kuala Lumpur, Malaysia

techtips

CALCULATE BANDWIDTH USAGE FOR VOICE NETWORKS

This document explains voice codec bandwidth calculations and features to modify or conserve bandwidth when designing and troubleshooting voice over IP (VoIP) networks.

cisco.com/packet/182_4e1

CONFIGURE VIRTUAL LANS ON WLAN CONTROLLERS

View a sample configuration for VLANs on wireless LAN (WLAN) controllers and the Cisco Catalyst switch that is associated with the controller. It is assumed that a working DHCP server provides IP addresses to the access points that are registered to the controller.

cisco.com/packet/182_4e2

CREATE CALL DETAIL RECORDS IN CISCO CALLMANAGER

This document describes a potential reason why the Cisco CallManager cluster fails to create Call Detail Records (CDRs), and provides a solution in a Cisco CallManager environment.

cisco.com/packet/182_4e3

TROUBLESHOOT DAEMON MANAGER ERROR MESSAGES

Identify and resolve an "Unable to Locate DLL" error message in CiscoWorks Daemon Manager.

cisco.com/packet/182_4e4

ELIMINATE CMS SHUTDOWNS

Learn how buffer overruns can potentially shut down a Configuration Management Service (CMS) process on the Distributor Admin. Workstation (AW). This tip provides a workaround and software fix in a Cisco IP Contact Center (IPCC) Enterprise environment.

cisco.com/packet/182_4e5

NAC Implementation Basics

PREPARING YOUR ENVIRONMENT FOR ENDPOINT SECURITY WITH NETWORK ADMISSION CONTROL by thomas howard

Left unmanaged, remote and mobile devices can expose your corporate network to viruses and other infections. Client devices can pick up malicious code when users disconnect from the corporate network, use an "untrusted" network such as the public Internet, and then reconnect.

Limiting the potential damage of such threats requires that enterprises confirm the identity and security posture of the client device, or host, as it attempts to gain access to the network. The Cisco Network Admission Control (NAC) framework enables this function. NAC builds upon traditional authentication, authorization, and accounting (AAA) network services, including the IEEE 802.1X authentication protocol and Remote Authentication Dial-in User Service (RADIUS), to ensure that the hardware and software configurations of a network host comply with organizational security policies. A network access device (NAD), such as a Cisco router, switch, wireless access point, or VPN concentrator, plays the role of authenticator between the host and a backend AAA server. In the Cisco NAC framework, that AAA server is a Cisco Access Control Server (ACS).

Given the risk of damage from viruses, worms, and malware, knowing not only who, but what, is attempting to access the network is extremely important. Using an extensive security policy built from group identities, operating system (OS) information, client security rule sets, antivirus updates, and patch status, network administrators can mitigate threats from both humans and machines.

Security Policy Definition

AT A HIGH LEVEL, your security policy should include what, where, when, how, and for whom network access will be permitted or denied (see table page 22). Such a policy will determine whether or not network access is granted based on location and access method, who is authenticated, what posture is required for the access method, and when these rules apply. After the high-level policy is expressed, go back and fill in the policy details such as minimal OS hot fixes, application versions, antivirus signature file versions, access control lists (ACLs), and virtual LAN (VLAN) assignments.

In larger organizations, the exact details of different parts of the policy will be owned by different departments and individuals. Depending on how cross-functional your organization is, you might find technical implementation less challenging than the political challenges of coming to an agreement on a single network access policy, authorized personnel to change policy, responsibility for implementation, and the methods for implementation.

Agentless Hosts

NAC IS A STRAIGHTFORWARD process when all hosts run the Cisco Trust Agent, which is client software that awaits a network challenge, then collects security posture information from local NAC-compatible applications and reports it to a posture validation server. The Cisco Trust Agent has an optional IEEE 802.1X supplicant; you can also use an 802.1X supplicant already installed on the host, or use a supplicant from a third-party vendor.

> With proper planning and phased implementations, Cisco Network Admission Control can provide an endpoint security architecture that mitigates threats from known vulnerabilities.

But what happens when there is no such host software present? This is often the case with network printers, IP phones, devices with embedded or hardened operating systems, and hosts with personal firewalls enabled. Hosts without the Cisco Trust Agent or an 802.1X supplicant cannot respond to a NAC challenge. Within the NAC framework, you must

understand which of these devices you have, in what quantity, and how you plan to handle them as exceptional cases. For Cisco IP phones, you can use Cisco Discovery Protocol in Cisco IOS Software to identify the device. You can also create whitelists of allowed devices and blacklists of restricted devices that identify devices by their MAC or IP address. The ACS compares the address with the preconfigured exception lists and determines the appropriate authorization group.

This function can be used with wildcards; in other words, rather than maintaining very lengthy lists of full addresses, you grant or deny access based on a partial address called an organizationally unique identifier (OUI). This approach is easier to manage but less specific and therefore less secure. For example, an outsider could find a printer OUI, configure his laptop with that OUI, and gain the printer's profile and access rights. While the OUI method is still vulnerable to MAC spoofing, it is one more hurdle for anyone trying to simply plug in to your network. If you prefer not to use static MAC address lists, there is an audit server option that uses vulnerability assessment technologies to verify a device type. Regardless of the method you choose, be aware that agentless hosts will be an issue. It is best to inventory the various types of agentless devices you anticipate seeing on your network and develop a strategy to handle them.

Isolation and Quarantine

THE GOAL OF NAC is to separate compliant, or healthy, users and hosts from noncompliant, or quarantined, users and hosts. Isolation of a noncompliant host is critical to mitigating risk to other hosts. The closer the network enforcement is to host, the better isolation you can provide to keep the unhealthy host from affecting healthy ones. Use of a network quarantine enables the noncompliant host to be brought into compliance so that it can be productive again. It is unproductive to deny all network access and force computer users into a helpless state. Instead, design your quarantine for self-help with URL redirection to a site containing mandatory software or automated updates with antivirus and patch management servers.

Automated Patch and Remediation

IDEALLY, MOST HOSTS are maintaining their compliance status well in advance of any network quarantines through the use of an automated patch management system. Only when this update method fails or an unauthorized user or host attempts to gain access is a quarantine or complete denial of access justified. Exactly how restrictive the quarantine is depends on your method for remediation. Assuming you want the host to self-remediate, you must understand which servers and network protocols will be used in the remediation process and permit access to them during quarantine.

The integration of an automated patch management system is critical to any successful NAC deployment. Without one, users will be left with little or no understanding of why they cannot connect and how to fix their posture. Depending on your software distribution strategy, patch management might constitute a visit by a support tech bearing a pack of CDs. Alternatively, you might set up a Website with all required software available for download. Or, you might turn to an automated patching agent and server from a third party. The most effective patch management will integrate with NAC enforcement by triggering patch

SETTING RULES Starting at a high level, build a security policy that addresses what, where, when, and how network access will be permitted and for whom.

Sample Security Policy for NAC

Where	Who	What	When	How
All	Employees	Windows XP SP2 Antivirus Anti-Spyware Patch	Always	802.1X: Employee VLAN
All	Employees	Noncompliant	Always	802.1X: Quarantine VLAN
LAN	Printers	MAC Address	Always	MAC-Auth-Bypass: Print Servers Only
WLAN	Guests	Don't Care	7am-7pm	Guest Hotspot/Internet Only
All	Others	Don't Care	Always	Deny All

downloads immediately upon quarantine, notify the user what is happening, and alert the network when finished so the host can attempt reauthentication. This minimizes the window of unproductive time that the host and its user will spend in quarantine.

Phased Deployment

NAC REQUIRES ADMINISTRATIVE and technical collaboration between the various groups responsible for desktops, servers, network operations, antivirus, client security, patch management, and support. At a minimum, before deployment you should know whether you want to check identity, posture or both; which application posture attributes you want to check; and whether the Cisco ACS make all access decisions alone or whether some will be distributed to application-specific servers. This knowledge helps you to scope the collaboration effort and determine which pieces must be in place before a pilot can begin.

Roll out NAC functionality in phases, starting with small, pilot environments. Pilots verify that the essential components and features are working smoothly. Many organizations prefer to start with a simple NAC L3 IP implementation on a router or VPN concentrator to check remote office and VPN user connection attempts (see "Posture Assessment Trigger Methods").

Another recommendation is to begin all pilots with NAC in a monitoring-only state, which requires enabling NAC such that all policy decisions grant full access. Getting started this way ensures that the basic mechanics of the process are working—that all clients have the Cisco Trust Agent and are communicating the expected posture information to the ACS, and that the ACS is talking to the other posture validation servers. This will allow you to determine the success of agent rollouts to hosts, try different policy rule sets, and begin to baseline scaling requirements and implications from timer settings and logs.

Inevitably, you will also encounter a variety of network devices or access scenarios that you hadn't anticipated. Reviewing NAC authentication results during the monitoring period will reveal them and allow you to craft alternative policies or network access arrangements. The logging information generated from these transactions will be valuable for developing troubleshooting tools and procedures for the support desk and effective reports about network health. Posture status from the logs will also validate how quickly your managed hosts are picking up new patch baselines and highlight those that aren't.

After achieving a level of proficiency from NAC monitoring-only mode, you will eventually feel more confident about enforcing restrictions for noncompliant hosts. Simply change the default network access permission on the network access devices and change the session-based authorization policies on the ACS. As the hosts reauthenticate throughout the day, they will each be migrated to an enforced NAC state.

With each NAC pilot, all of the security teams involved will gain more confidence with the collaborative security approach. Eventually, you will be able to deploy to your entire network and handle the NAC L2 IP and NAC L2 802.1X and increase your number and type of posture policies. Depending on the current state and size of your network, it may take a year or more to enable NAC on every edge device in your network.

With proper planning and phased implementations, Cisco Network Admission Control can provide an endpoint security architecture that mitigates threats from known vulnerabilities. Start with a simple security policy and continue taking incremental steps to include all desired features and components. New security threats continually evolve—and so must your security policy. P

Posture Assessment Trigger Methods

NAC L3 IP: Challenge for AAA credentials takes place in a Layer 3 device, such as a router or VPN concentrator, using the EAP-over-UDP protocol.

NAC L2 802.1X: Challenge takes place at a Layer 2 connection point (switch port or wireless access point) using the IEEE 802.1X with the EAP-FAST protocol.

NAC L2 IP: Challenge takes place at a Layer 2 switch port via IP. Host is challenged using EAP-over-UDP when it is assigned a Dynamic Host Configuration Protocol (DHCP) address or upon the first Address Resolution Protocol (ARP) request.

Further Reading

- NAC Deployment Guide
  cisco.com/packet/182_4f1
- Cisco NAC
  cisco.com/go/nac

THOMAS HOWARD is a solution engineer in Cisco's Security Technology Group. One of the NAC solution architects, he regularly educates and helps customers with their NAC deployments and assists with NAC Program partner software integration. He can be reached at thomas@cisco.com.

Boosting Streaming Media

HOW ENTERPRISES CAN ENHANCE AUDIO-VIDEO BROADCAST PERFORMANCE by silvano da ros

Streaming media is analogous to conventional analog television broadcasting, in that you see and hear content concurrently with its arrival from a cable or antenna feed. Where digital media files once required you to download an entire audio-video file before viewing it on your computer, newer streaming-media file formats allow media players to present content right as you download it, packet by packet. • Cisco Application Networking Services (ANS), previously called content networking, enable enterprises to accelerate the performance of broadcast audio and video streaming-media applications. Acceleration is especially important when delivering rich media to remote locations over comparatively bandwidth-constrained WANs. • Activating certain intelligent network services in your router network accelerates the delivery of your streaming media. For example, if your clients support IP Multicast, you can enable this feature in routers to prevent broadcast-based applications from flooding your network. Also, using quality of service (QoS)—classifying traffic and marking IP Type of Service (ToS) and Differentiated Services Code Point (DSCP) bits in the packet header for priority—helps ensure that real-time streaming packets take priority over bulk file transfers within the network. Finally, enabling the Resource Reservation Protocol (RSVP) in your routers, provided that your streaming applications support it, allows receivers to signal their desired resource requirements per streaming flow.

Streaming Media is supported in several protocols and formats. For real-time applications, TCP provides the transport for control traffic and UDP does so for data traffic. Unlike TCP, UDP does not use packet buffering for reordering out-of-sequence packets and retransmitting missing packets, reducing delay. Real-time applications prefer dropped packets to packet buffering to avoid latency.

Moving up the network stack, the most popular application-layer protocols that Cisco streaming media features use are Real-time Transport Protocol (RTP)—the UDP-based application-layer protocol for streaming audio and visual data; Real Time Streaming Protocol (RTSP)—the TCP-based application layer protocol for providing the control of streaming flows, such as play, pause, fast-forward, and rewind; and Microsoft Media Services (MMS)—the control and streaming application layers for Windows streaming media content.

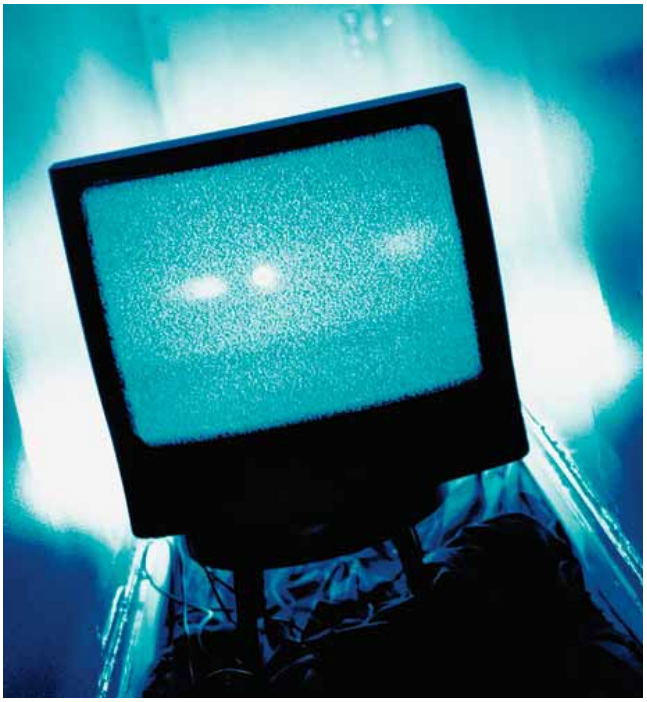

GETTY IMAGES

BY ACTIVATING intelligent services in the network, you can accelerate delivery of streaming media.

CISCO CONTENT NETWORKING ARCHITECTURE

- Cisco Wide-Area Application Engine (WAE) hardware platform—formerly called the Cisco Content Engine (CE) for running acceleration and redirection services
- Cisco Application and Content Networking Software (ACNS) for content distribution to local caches. ACNS runs on the WAEs and supports standard and third-party streaming server software
- Cisco IP/TV Program Manager for scheduling live and rebroadcast events
- Cisco's IP/TV Broadcast server (or a third-party streaming data converter) for content distribution to an origin server, which stores the content and fulfills service requests
- Optionally, Cisco Content Routers for routing client requests to the most appropriate WAEs
- Optionally, Cisco Content Distribution Manager (CDM) for centralized network-wide management tasks, such as orchestrating the distribution of live content and prepositioning on-demand streaming content

The streaming formats that Cisco streaming media devices can serve include Apple QuickTime, MPEG4, Real Networks Media, and Windows Media.

Here's how the components of the Cisco streaming media content networking architecture work together: Source media is input and encoded digitally using capture cards on the Cisco's IP/TV Broadcast server or third-party streaming data converter. Third-party streaming data converters include Apple QuickTime Broadcaster, Real Network RealProducer, and Microsoft Windows Media Encoder (WME).

The packet stream is then fed to an on-demand/live origin server, where it is either relayed to requesting clients live or encoded into a streaming container file and stored for future rebroadcasts and on-demand requests. Origin servers can include Apple QuickTime Streaming Server, Real Networks Helix Universal Server, and Microsoft Windows Media Server.

In Figure 1, the Cisco IP/TV Broadcast server behaves both as an on-demand/live origin server and a streaming data converter, requiring the IP/TV Program Manager for scheduling live and rebroadcast events. By contrast, third-party on-demand/live origin servers mentioned previously perform the scheduling function directly.

If you are using Cisco IP/TV with CDM, the IP/TV Program Manager interacts with CDM for scheduling live and rebroadcast streams.

Streaming Video on Demand

A VIDEO ON DEMAND (VOD) is a multimedia presentation that users can request and view at their convenience. VoD files are encapsulated in standard MPEG4 format or in the specific streaming container file format of the vendor software. Streaming file formats differ from nonstreaming formats in that they contain information about how to stream the file over the IP network. For example, Real Networks and Microsoft Windows streaming media files contain packetized data, which

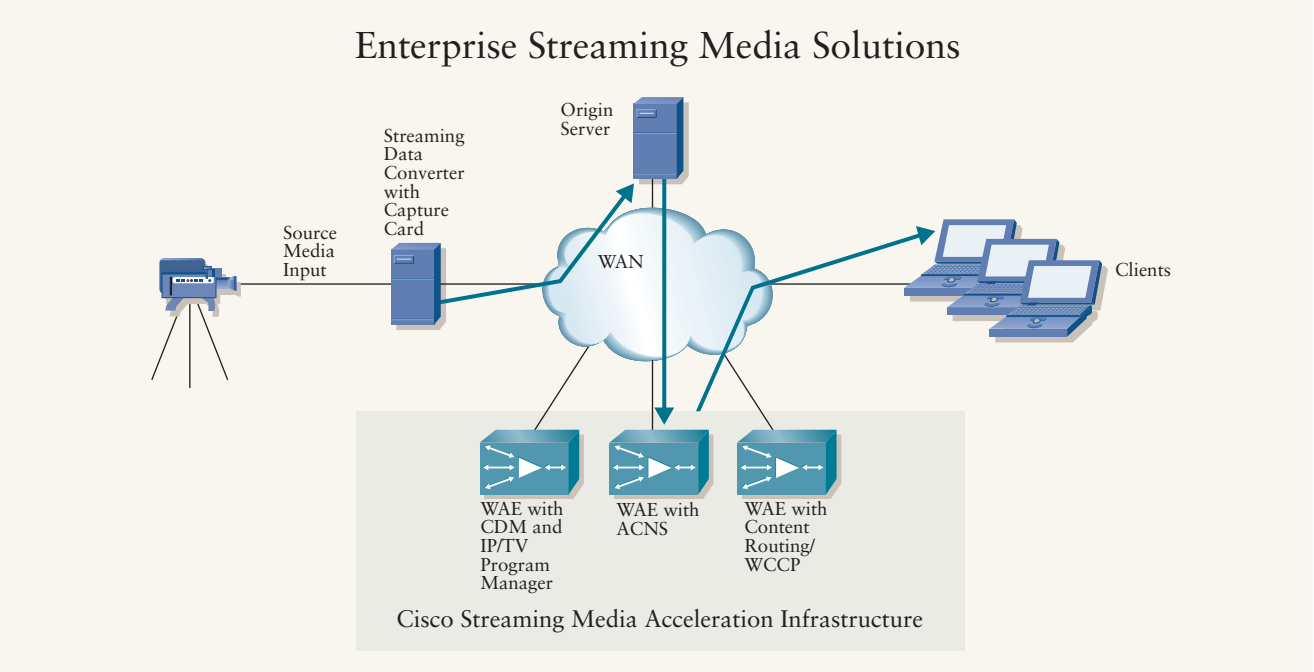

FIGURE 1 Cisco streaming-media acceleration features that reside in ACN enable you to enhance the performance of both Cisco IP/TV and third-party streaming software.

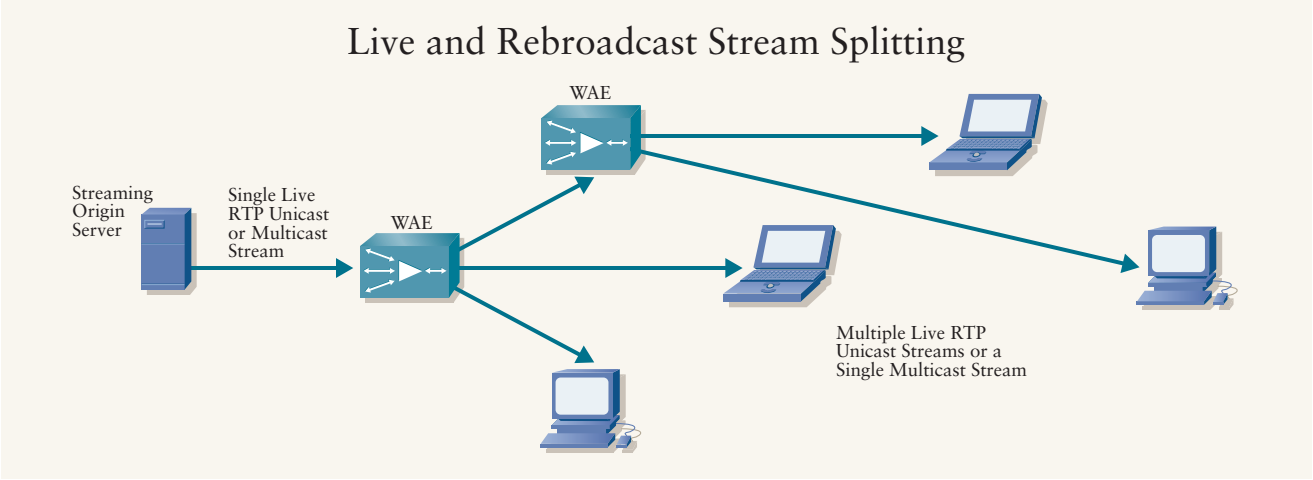

FIGURE 2 Live splitting conserves bandwidth and network resources by reducing the number of duplicate streams on the network.

interleaves the audio and video tracks together so that they are synchronized when the client downloads the stream of packets. Alternatively, MPEG4 and Apple QuickTime files contain separate hint tracks to tell the server how to synchronously packetize the data during real-time delivery.

Cisco WAEs support VoD streaming individually or as a part of an ACNS solution that you administer with a CDM. You can configure your individual WAEs with conventional caching to locally deliver Windows and Real media files to requesting clients. The WAEs can cache this content from internal servers or from anonymous servers on the Internet. As a complete ACNS solution, you can preposition Real Networks, Microsoft Windows, MPEG4, and Apple QuickTime files to WAEs throughout your network using Cisco CDM. Prepositioning involves distributing VoDs to WAEs throughout your network, prior to client requests, using distribution trees. You create the trees using CDM and ACNS; ACNS software supports both unicast and multicast VoD file distribution. The benefit of prepositioning is that you can distribute corporate streaming media content after hours, when branch office bandwidth utilization is lowest.

Routing Client Requests

CLIENTS ARE REDIRECTED to the appropriate local Cisco WAE appliance automatically in one of three ways: by the Web Cache Communication Protocol (WCCP) running in the local router, by a proxy running on the client device, or by a Cisco WAE functioning in "content router" mode at headquarters and supporting subdomains for redirection.

For its part, WCCP optimizes content delivery between clients and WAEs by transparently redirecting client requests to appropriate WAEs. WCCP does not inspect the URL or HTTP request to classify traffic for redirection. Instead, a WCCP-enabled router inspects packets on incoming or outgoing interfaces and matches them against service groups that you configure for WCCP inspection. WCCP defines service groups by port numbers.

WCCP version 2 is required for streaming media request redirection. For example, WCCPv2 currently supports redirection of standard RTSP, MMS-over-UDP, MMS-over-TCP, and Microsoft's implementation of RTSP over UDP. It also supports multiple redundant routers, improved security, Layer 2 redirection, and redirection of applications other than HTTP over TCP port 80.

As an alternative to WCCP, you can also configure Cisco content switches to redirect client requests to your WAEs. The advantage of content switch redirection is that you can configure the switch to classify traffic by criteria such as URLs, file names, and extensions. Another option is to manually configure the Internet connection settings of each workstation in your network to point to the WAE in closest proximity to the particular client (for example, the one located in the same branch office as the client). If you are using CDM administration, the complexity of this manual approach can be offset by using dynamic Proxy Auto-Configuration (PAC). Dynamic PAC uses coverage zones to determine the best-suited WAE for client requests and dynamically adjusts a client's proxy settings with the WAE URL or IP address.

A coverage zone is similar to a static route, in that it contains the association between a client subnet and the IP address of the WAE that serves that subnet. As a result, you must create a coverage zone for every subnet in your network that uses your streaming content. Also available in a CDM-administered

Continued on page 81

SILVANO DA ROS is a network consultant and a Cisco Press author. Previously a systems engineer at Cisco Systems, he worked with enterprise organizations on emerging network solutions.

Static Routing

TROUBLESHOOTING REDUNDANT PATH NETWORKS WITH SELF-TUNNELS

by franco conti and luca piorico piorgo

Using static routes in a network topology with redundant paths, a LAN interface failure on a remote site active router can cause two LANs to become isolated even though Hot Standby Router Protocol (HSRP) is configured on the routers. Features in Cisco IOS Software Release 12.3 and later, such as Object Tracking for Reliable Static Routing Backup, or tunneling between an ISP and remote sites might overcome this issue, but if you are running an earlier IOS version in which these features are unavailable, you can use a pure static routing strategy.

What Happens When Failure Occurs?

IN FIGURE 1 (PAGE 30), ISP1 and ISP2 are ISP routers, where ISP1 is the HSRP active router with the HSRP track command configured on its WAN interface (Serial0/0). According to ISP policies, only static routes to remote customer sites can be configured on those devices (no tunnel is allowed either). CUST1 and CUST2 are the routers in the customer remote site, with CUST1 configured as the active device.

If a failure occurs on the Ethernet interface of CUST, the following static route is configured on ISP1:

```
ip route 10.26.248.0 255.255.255.0 ser0/0
```

Under normal conditions, all the traffic destined to 10.26.248.0/24 flows through ISP1. When the CUST1 interface failure takes place, ISP1, unaware of the failure, is still the active ISP router and keeps sending the packets to CUST1 via its WAN link, which has remained active. Because no mechanism exists to reroute traffic destined to 10.26.248.0/24 from ISP1 to ISP2, the two sites become isolated. To restore the correct routing operation without the intervention of the ISP we should, for example, turn CUST1 off. This causes the WAN link to ISP1 to go down and consequently, according to the HSRP track command, ISP2 becomes the new active router.

Troubleshooting with Self-Tunnels

TO SOLVE THIS ISSUE without manual intervention the concept of "self-tunnels" is introduced (Figure 2, page 30). A self-tunnel is a tunnel that is sourced from and destined to the same router. Add the following configuration lines to CUST1:

```
interface Loopback6000
ip address 10.26.247.1 255.255.255.252
! see note
!
interface Tunnel1000
ip address 10.26.247.9 255.255.255.252
backup interface ser0/0
tunnel source 10.26.247.2
tunnel destination 10.26.247.2
no keep-alive
!
interface Tunnel2000
ip address 10.26.247.13 255.255.255.252
backup interface Loopback6000
tunnel source 192.168.60.61
tunnel destination 192.168.60.61
no keep-alive
!
Interface FastEthernet0/0
Ip address 10.26.248.202 255.255.255.0
Standby 1 ip 10.26.248.201
Standby 1 preempt
Standby 1 priority 105
Standby 1 track Serial0/0

Interface Serial0/0
Ip address 172.16.1.2 255.255.255.252 (omitted)

ip route 192.168.60.60 255.255.255.252
FastEthernet0/0 ! This subnet is used only
to change Tunnel2000 interface status.

ip route 10.26.249.0 255.255.255.0
172.16.1.1
```

In order for this mechanism to work properly, Loopback6000 must not belong to one of the major networks that CUST1 is able to route (that is, neither a default route nor a static route such as 10.26.247.0 with subnet mask < 30 should be

present. For example, if the route 10.26.0.0/16 is present in the routing table, Tunnel1000 will always be active, preventing the mechanism from working).

Under normal circumstances, Tunnel2000 is active when the static route to 192.168.60.60/30 is present in the routing table (that is true provided that Fa0/0 is active), and the backup interface command of T2000 causes Loopback6000 to remain down. With Loopback6000 down, Tunnel1000 is also down (because the loopback is both the source and the destination of the tunnel) and Serial0/0 is active (because it is the backup interface of Tunnel1000).

In the case of a FastEthernet0/0 failure the static route is withdrawn from the routing table and T2000 goes down, causing Loopback6000 to become

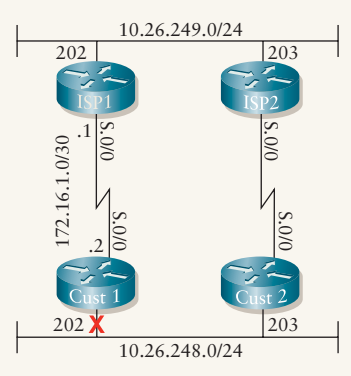

FIGURE 1 Customer sites become isolated from each other when the CUST1 Ethernet interface fails.

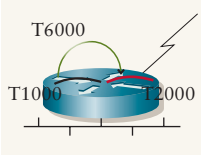

FIGURE 2 The self-tunnel is sourced from and destined to the same router.

active. But T1000 also becomes active, and finally, thanks to the backup interface command, Ser0/0 goes down.

After implementing this mechanism, the serial interface of ISP1 also goes down, triggering the HSRP status change (active > standby) by means of the **track** command configured on that serial. Therefore, the inbound traffic is forwarded through ISP2, which has finally become active.

This is a basic scenario. In a more complex environment, you might need to implement further networking strategies along with self-tunnels. For example, you can use recursive lookup in a network topology using a dynamic protocol under CUST1-2 to prevent a similar issue. P

FRANCO CONTI AND LUCA PIORICO PIORGO, both network engineers for Tele Sistemi Ferroviari in Rome, Italy, have more than ten years experience in the industry. They can be reached at f.conti@tsf.it and l.piorico@tsf.it, respectively.

Data Center Consolidation Tips

MIGRATING TOWARD CONSOLIDATION? HERE ARE SOME CONSIDERATIONS AND DESIGN EXAMPLES. by zeeshan naseh

Enterprise data centers are complex environments that host mission-critical servers, databases, and data storage devices. As such, it goes without saying that data center designs must address high availability, scalability, and redundancy.

There are several distinct approaches and methodologies to building a data center with these traits. Being successful requires that you keep application behavior in mind during the planning and design phase. This means considering issues such as application flows, scalability, redundancy, load balancing, migration, management, and security. For example, if you are consolidating multiple data centers into one, take inventory of the number of servers and other devices you are migrating so you can appropriately size the new facility to accommodate the consolidated resources. A recommended approach is to build the new facility with rack space that accommodates scaling by five- to tenfold to allow for growth.

It's also important to understand the various application flows to be supported—which devices each application traverses and which protocols they use—so that each application can be supported properly. For example, if an application receives 10,000 requests per second, it cannot be supported successfully from a single server and will thus require load balancing across multiple servers.

Finally, building a data center in a modular fashion that allows physical devices to provide virtualized services across a number of business units and applications is integral to building scalability into your data center design.

ACHILLI BIGLIARDI PHOTOGRAPHY

Why Consolidate Data Centers?

KEY BUSINESS DRIVERS BEHIND data center consolidation include improved cost control, increased operational efficiency, and more effective resource utilization. Following are some of the primary motivations behind consolidated data centers:

- Network infrastructure cost savings—The reduction in the number of network services and enterprise devices such as switches or routers yield direct cost savings. For example, a consolidated data center design enables the enterprise to virtually segment a server load-balancing device to service multiple applications or business units instead of operating a separate pair of load-balanced servers for every business unit. Similarly, a firewall services module deployed in a virtualized fashion in an aggregation switch allows the one physical firewall to function as multiple logical firewall instances.
- Standardization and centralization of resources—Often in data center consolidation, the number of distributed server farms is reduced, and servers are relocated to existing or new facilities. These new consolidated data centers are interconnected, and the operational, support, and design best practices standardized. Likewise, application, server, and storage resources are centralized. This standardization and centralization yields much lower total cost of ownership (TCO) than before the consolidation.

In a different take on centralization, if certain functions are collapsed into services modules in a routing/switching device, such as the Cisco Catalyst 6500 Series Switch, the management of those separate functions is consolidated, as well. For example, if a pair of redundant Catalyst 6500 switches is configured with a Content Switching Module (CSM) for load balancing, an SSL Services Module (SSLM), and a Firewall Services Module (FWSM), the management of these devices can be centralized to some extent. The FWSM does require separate management using the Cisco VPN/Security Management Solution (CiscoWorks VMS). However, the status of all modules can be polled from the Catalyst 6500

Continued on page 33

THE MOST common trend affecting data centers worldwide today is consolidation.

Data Centers, Continued from page 31

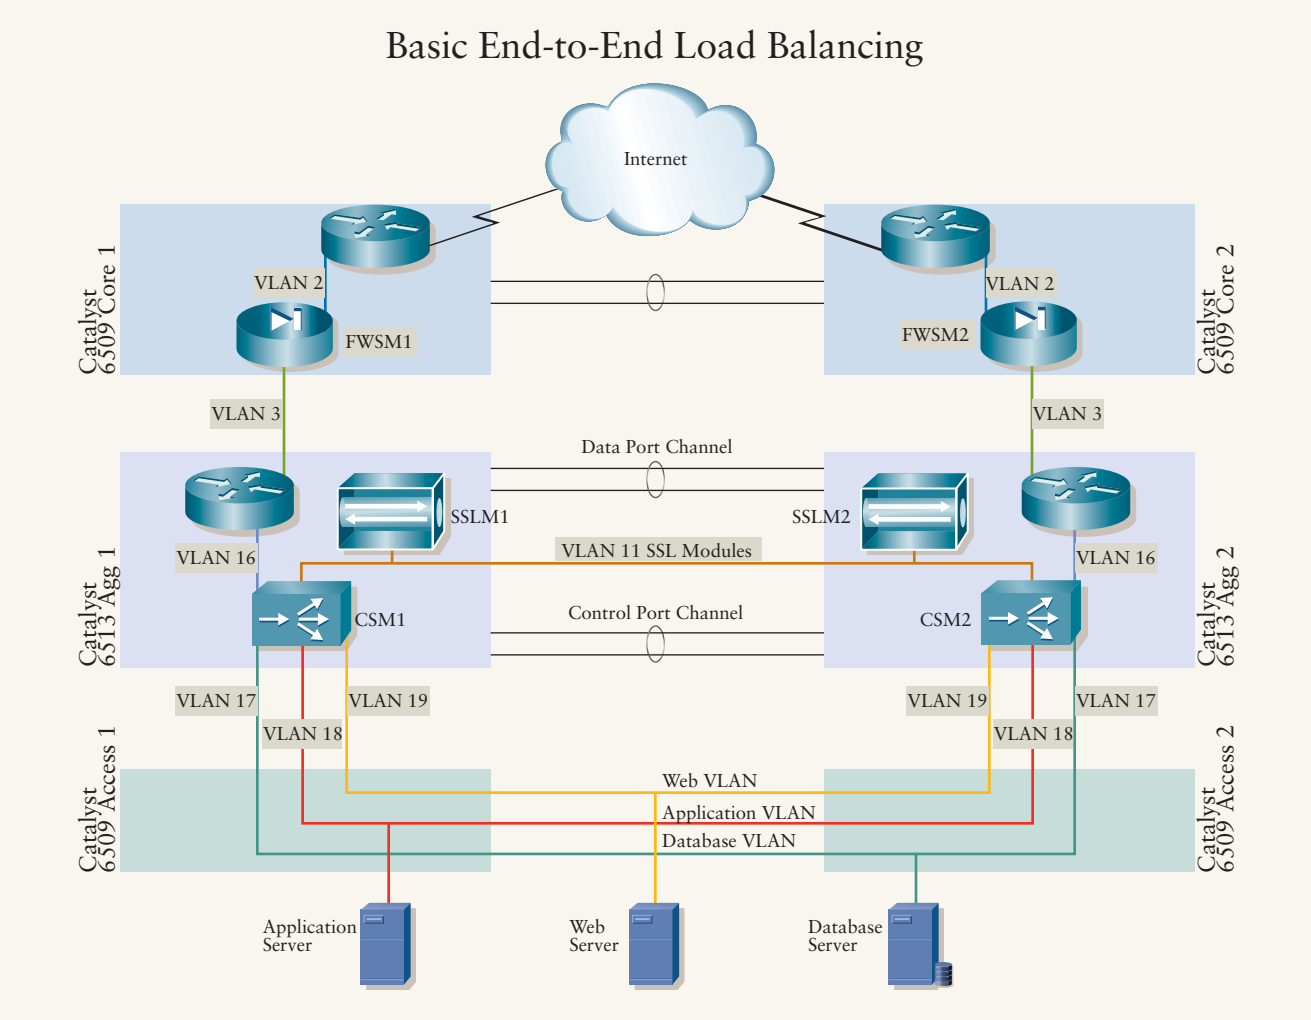

PROTECTING PERFORMANCE Firewall services in the core first check incoming server requests for malicious activity. Then, server load balancing in the aggregation layer allows traffic to be load balanced from client to Web server, from Web server to application server, and, finally, from application server to database server for user authentication and authorization.

Supervisor Engine. In the case of the CSM, the Supervisor Engine can also poll the status of the physical servers.

Consolidated data center designs also substantially reduce cabling that is otherwise needed to connect the appliances with the infrastructure. Service modules do require power from the Catalyst 6500 Series Switch, but it is significantly less than the power required by separate appliances.

• Virtualization of network resources—As the network resources are consolidated, deploying virtualization on a single module or across all the service modules is simpler. Virtualization basically means to logically split the network resource, such as a firewall, into multiple entities dedicated to different departments or business units.

• Scalability—Consolidated data center designs use the core-, aggregation-, and access-layer model. Network services such as load balancing, firewall, and SSL offload are provided in the aggregation layer. This design approach makes the data center modular. Each aggregation switch pair, together with multiple access switches that connect to it, forms an aggregation module. The data center can be scaled by deploying more aggregation modules, each of them connecting to the core switch pair.

Such designs use the high-bandwidth service modules on the Catalyst 6500 Switch to provide server load balancing and firewall services in a scalable, centralized fashion. Deployment in a pair of switches in the distribution or aggregation layer allows them to be used by all the servers connected to the access switches belonging to that aggregation switch pair. In instances where separate physical devices are required by the application or security team, use of an appliance is desired.

Reviewing the data center design option described in this article and in three alternative designs provided at Packet Online (**cisco.com/packet/182_5b1**) will help you to understand how a modular, virtual, service-based data center design helps you achieve the benefits described.

Sample Data Center Design: End-to-End Load Balancing

BASIC END-TO-END LOAD BALANCING IS a fairly common design implementation where the FWSMs or appliance-based firewalls reside in the core network layer, and the CSMs and SSLMs are implemented in the aggregation layer. The figure on page 33 displays the logical topology, not the physical connectivity representation.

The strategy behind this design is to enable load balancing along the entire traffic flow. In the case of a client request designed for Web servers, the first session comes in from the Internet (HTTP port 80) and traverses the firewall module or appliance, running in Layer 3 mode, to make sure that the request is not malicious. The FWSM (or appliance) constitutes the security perimeter between the untrusted Internet and the trusted data center, and the aggregation and access layers are considered trusted zones.

From there, the request is forwarded to the Catalyst 6500 aggregation switch pair, destined for a virtual IP (VIP) address that resides on the CSM, which is deployed in routed mode. Using virtual LANs (VLANs), the CSM can load balance connections to all back-end resources. When the request reaches the virtual server on the CSM, the CSM load balances it across the appropriate Web servers, which can initiate another session to the application servers, if necessary. In this design, traffic from the Web server to the application server will also be load balanced. The application server now needs to communicate with the database server to obtain user authorizations and credentials, and the CSM load balances those requests, as well.

The key importance of this design is that it enables load balancing in all directions. The IP connectivity in this design can be summarized as follows:

- The Multilayer Switch Feature Card (MSFC) router services module on the Catalyst 6509-Core-1 and Catalyst 6509-Core-2 connects to the Internet. This link is direct or through edge routers.
- Toward the inside, the MSFC connects to the FWSM over VLAN 2.
- The FWSM connects to the MSFC in the aggregation switches over VLAN 3.
- The aggregation MSFC connects to the CSM on VLAN 16.
- All load-balanced server VLANs (17 to 19) exist on the CSM.
- VLAN 11 is the SSLM VLAN on the CSM.
- Connectivity between aggregation and access is at Layer 2 (switched).

High Availability Best Practice

FOR HIGH AVAILABILITY, all components are deployed in a redundant manner: two core switches, two firewalls, two aggregation switches, two CSMs, and redundant paths to the aggregation switches. Cisco Hot Standby Routing Protocol (HSRP) support in the Catalyst 6500 Series Switch enables stateful failover of routing/switching from one switch to the other. In addition, the CSM and FWSM are stateful devices, so a shift from one switch to another retains all the CSM and FWSM state information for seamless failover. In the figure, the CSM's default gateway is the HSRP group IP on the aggregation-layer Catalyst 6513 MSFC on VLAN 16.

Design 1 Configuration

FOLLOWING IS THE PARTIAL CSM configuration template used for this design:

```
!
module ContentSwitchingModule 3
 vlan 16 client
  ip address 10.16.1.12 255.255.255.0
  gateway 10.16.1.1
  alias 10.16.1.11 255.255.255.0
!
 vlan 11 server
  ip address 10.11.1.2 255.255.255.0
  alias 10.11.1.1 255.255.255.0
!
 vlan 17 server
  ip address 10.17.1.2 255.255.255.0
  alias 10.17.1.1 255.255.255.0
!
 vlan 18 server
  ip address 10.18.1.2 255.255.255.0
  alias 10.18.1.1 255.255.255.0
!
 vlan 19 server
  ip address 10.19.1.2 255.255.255.0
  alias 10.19.1.1 255.255.255.0
!
!
 serverfarm ROUTE
  no nat server
  no nat client
  predictor forward
!
 vserver ROUTE
  virtual 0.0.0.0 0.0.0.0 any
  serverfarm ROUTE
  persistent rebalance
  inservice
!
```

The following configuration example shows a partial configuration from the MSFC in the aggregation switch:

```
MSFC SVI
!
interface Vlan16
  ip address 10.16.1.2 255.255.255.0
  standby 16 ip 10.16.1.1
  standby 16 priority 150
```

Other Considerations

IN THIS BASIC DESIGN, the security perimeter created by the firewall in the core protects the Internet-to-Web server edge only; this security is not possible between the Web, application, and database tiers. In the aggregation layer, however, some security using VLAN tags on the CSM to segregate traffic is possible.

Also, some extra configuration apart from the load-balancing configurations is needed on the CSM for direct access to servers and non-load-balanced, server-initiated sessions. Because the MSFC is directly connected to the CSM on VLAN 16, a capability called route health injection (RHI) can be used. RHI allows the CSM to advertise the availability of a VIP address throughout the network. Multiple CSM devices with identical VIP addresses and services can exist throughout the network. This allows one CSM to override the server load-balancing services of others if those services on the other devices should become unavailable.

In this design, both application traffic and management traffic are load balanced. If there is a crash in the inline CSM, management access to the physical servers is thus cut off. The same condition applies to any server-initiated flows: syslog messages, upgrade downloads, or Internet-launched software patches will all traverse the load balancer. Ideally, downloads and management traffic should bypass the load balancer. Otherwise, if you want to send traffic to a specific server to determine management information, load balancing might send your request to the wrong server.

The three additional data center design options available at **cisco.com/packet/182_5b1** describe how to build upon the basic configuration in this article to layer on tiered security and to prevent management traffic from being load balanced. **P**

ZEESHAN NASEH, CCIE NO. 6838, is a technical leader in Cisco's Advanced Technologies Services Group and author of the data center book *Designing Content Switching Solutions* (ISBN 1-58705213x). He can be reached at znaseh@cisco.com.

PEOPLE REACHING PEOPLE

New Unified Communications product portfolio changes everything

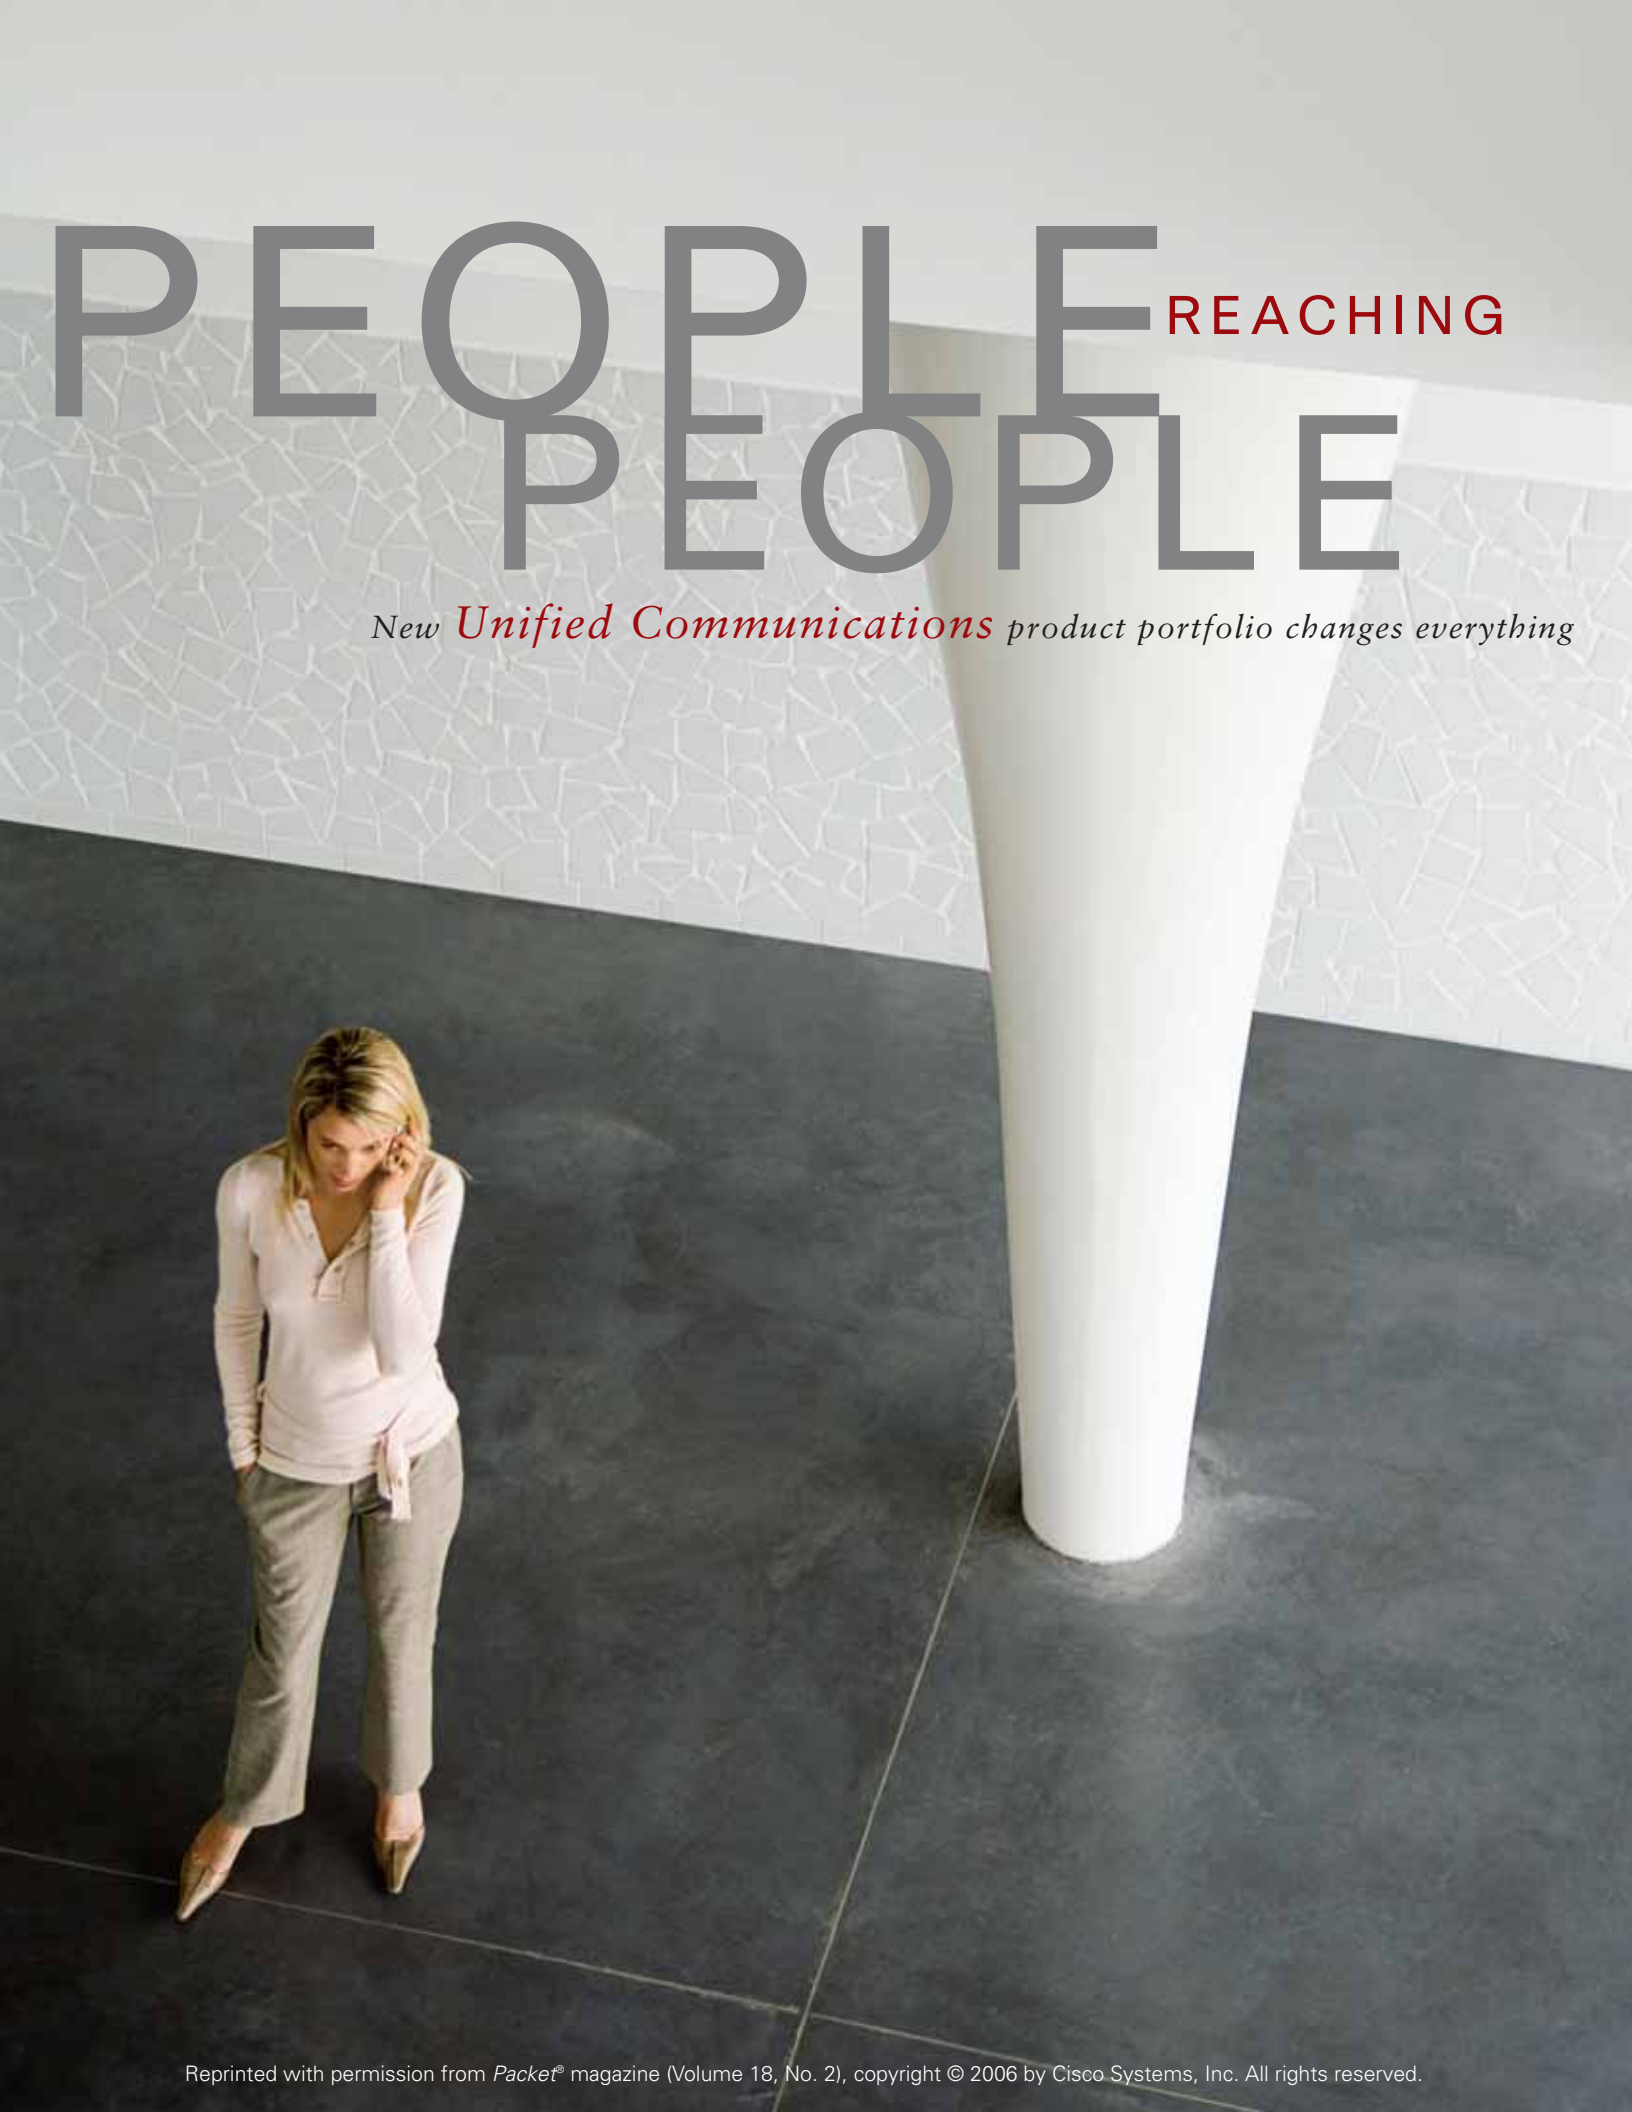

about the way businesses communicate.

by RHONDA RAIDER

Communications technology is poised on the brink of a transition as significant as the one from analog to digital, and from digital to IP. "The first hurdle in unified communications was tuning the network to handle performance-sensitive applications such as voice and video," says Rick Moran, vice president for IP communications in the Product and Technology Marketing Organization at Cisco. "Cisco Unified Communications heralds the next phase, which is to improve the way that people work with applications and with other people using presence services and open standards."

Introduced in March, Cisco Unified Communications is an integrated, comprehensive business communications system that includes more than 30 new products, enhanced management and administration, and simplified pricing and service offerings. Major enhancements include:

- Presence and preference information that helps employees reach the right person, the first time, on the available device that will provide the richest communications experience.
- Greater openness through Session Initiation Protocol (SIP) support, enabling developers to introduce new SIP-compliant endpoints and applications.
- More options for reaching mobile employees, including dual-mode phone support and Mobile Connect for single-number reach.

GETTY IMAGES

Overcoming Complexity

Why the change, and why now? "The myriad communications technologies that have evolved to support a global, mobile workforce have introduced too much complexity," says Vickie McGovern, director of marketing for unified communications in Cisco's Product and Technology Marketing Organization. The average business person uses more than six communications devices and more than four communications applications, such as e-mail and various instant messaging (IM) clients, according to a September 2005 survey by Sage Research. And Gartner predicts that by 2007, more than 66 percent of the overall workforce will use mobile or wireless computing devices. "If you want to reach a co-worker right away, how do you know which device to use?" McGovern asks.

In fact, the tedious process of trying to reach co-workers and partners on one device after another—a concept that McGovern calls "human middleware"—is a glaring omission in today's business process integration (BPI) efforts. "Neglecting to integrate communications into business processes delays decision-making, which negatively affects the top and bottom lines," she says. Inefficient communications hamper strategic initiatives such as lean/just-in-time manufacturing, supply-chain optimization, customer relationship management (CRM), and Six Sigma. It can have disastrous results in perishable goods industries. And it can impede the development of next-generation contact center environments that depend on the ability to quickly escalate customer inquiries, route them to the right expert across the distributed contact center, and resolve them quickly using the optimum mix of media.

The goal of Cisco Unified Communications is to use the power of the network to transform communications into an asset rather than a liability to business processes. "Cisco Unified Communications integrates communications with business processes so that organizations can be more responsive to customer needs and enhance business and employee productivity," says Don Proctor, senior vice president of the Cisco Voice Technology Group. "It can directly improve the top and bottom lines."

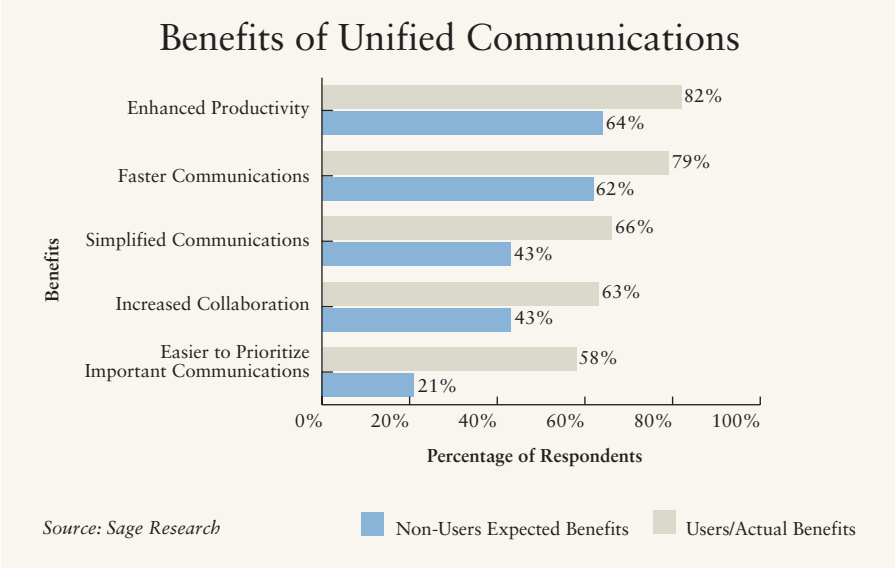

UNIFIED COMMUNICATIONS Actual benefits exceed expected benefits for an increasingly mobile workforce.

Top-Line Impact of Unified Communications

Say a customer calls to place a large order and the account executive is away from her desk. In most businesses today, the sale is delayed until the account rep checks voice mail and is able to connect with the customer—who by this time might also be away from the phone.

Cisco Unified Communications provides new capabilities that can speed up revenue realization. If the account rep has a dual-mode phone, he or she can receive the call via GSM when out of range of the office's 802.11 network, making it unnecessary for callers to be aware of and dial two different numbers (see sidebar, "IP Phones: What's New," page 39). Or the account rep does have separate cellular and office phones, he or she can use the Cisco Mobile Connect service, enabled by Cisco Unified Mobility Manager, to either have multiple phones ring at the same time or have all calls routed to a particular device. The Cisco Mobile Connect service even enables the rep to transfer a call in progress on her cell phone to her office phone when she arrives at the office, without interrupting the call.

Now suppose the customer on that call requests a product modification to clinch the sale. Cisco Unified Communications can help facilitate an ad-hoc meeting. The rep walks into an office and launches a new desktop application, Cisco Unified Personal Communicator, which shows presence and preference information that the Cisco Unified Presence Server has aggregated from other elements of the Cisco Unified Communications system (see sidebar, "Cisco Unified Personal Communicator").

The rep quickly identifies product team members who are currently available by phone, and simply drags their names into a window to launch an immediate collaborative session on Cisco Unified MeetingPlace Express, combining voice and video conferencing and Web collaboration. (Employees using the new Cisco Unified IP Phone 7985, a personal desktop video phone, can escalate from a voice-only session to voice and video at any time during the call.) They view the agreement and work out any details in real time, enabling the account rep to close the sale in the shortest possible time.

Presence and Presence

Cisco Unified Communications uses the pervasive reach of an intelligent network so that people reach people, not just unattended communications devices. Cisco Unified Presence Server aggregates and publishes presence information from all SIP-enabled devices and applications in the network, including Cisco Unified CallManager, Cisco Unity Voice Mail, Cisco Unified MeetingPlace, and new models of Cisco Unified IP phones.

Cisco Unified Personal Communicator

Available for both Microsoft Windows XP and Macintosh OS X operating systems, Cisco Unified Personal Communicator integrates directories, presence information, and Cisco Unified MeetingPlace into a single Web-based interface, helping employees communicate more effectively and be more productive. To communicate with another person on the network, a user simply drags that person's name into a window, and the system reaches out according to the presence and preference information available on the network—dialing the appropriate phone, sending an instant message, or calling the Cisco Unity voice-mail box if no other options are available.

IP Phones: What's New

Enhanced Mobility: Cisco Unified Communications supports dual-mode phones from Nokia, which combine the functionality of a wireless IP phone and a cellular phone. They connect via IEEE 802.11b/g when the signal is available and otherwise via GSM, enabling single business-number reach. Another boon for mobility, the Mobile Connect service enables employees to transfer calls between their cell phones and Cisco IP phones, or vice versa. "Recently my cell phone rang when I was at the Munich airport," says Brian Dal Bello, director of Cisco Unified CallManager product marketing. "The call was from an account executive in the US who had dialed my Dallas office. I was able to resolve her issue immediately instead of waiting until I heard a voice mail many hours later." Had the call continued after Dal Bello arrived at a Cisco office, he could have transferred the call to any Cisco IP phone, avoiding unnecessary per-minute charges from the cellular service provider.

Gigabit Ethernet Phones: Several new Cisco Unified IP phones can connect to computers with Gigabit Ethernet connections: the Cisco Unified IP Phone 7961G-GE (gray-scale with six buttons), 7941G-GE (gray-scale with four buttons), and 7971G-GE (color with touch-sensitive display).

Enhanced SIP Support: When used with Cisco Unified CallManager 5.0, Cisco Unified IP Phones 7971G, 7970G, 7961G, 7941G, and 7911G support enhanced SIP services, such as presence.

Unicode Support: Cisco Unified CallManager 5.0 now offers Unicode support, enabling users in Japan, Korea, and China to see their own alphabets on the display of Cisco Unified IP phones. Support for Arabic is planned.

Presence information appears on SIP-enabled devices and applications, including Cisco Unified IP phones and Cisco Unified Personal Communicator. An employee who views a directory or call log on either interface sees beside each entry whether the caller is on the phone, off the phone, or has left a message such as "Back in five minutes," or "Call my cell phone."

"Most people still think of presence as, 'Are you there?'" says Joe Burton, director of engineering for the Cisco Voice Technology Group. "Cisco extends the concept by also tracking the users' device capabilities, enabling the richest communications experience possible." For example, by communicating with the Cisco Unified CallManager, Cisco Unified Presence Server knows whether an individual in a Cisco Unified MeetingPlace conference also has video capabilities.

Users can supplement or override presence information by entering preferences in Cisco Unified Personal Communicator—for example, that they prefer e-mail to voice mail, or do not want to be disturbed under any circumstances until 10 a.m.

"Preference features in Cisco Unified Personal Communicator enable employees to opt in or opt out of a particular communications mode," says McGovern. "The concept is sometimes called, 'Find me, follow me, hide me.'"

SIP for Presence and Openness

Cisco Unified Communications takes advantage of standards such as SIP, SIMPLE, SOAP, and Ajax to enable interoperability among different types of communications devices—voice, video, IM, cellular—from different vendors.

An IETF standard, SIP defines how disparate devices set up a session, share real-time media streams, and terminate the session. SIP makes Cisco Unified Communications an open platform. People who use SIP-compliant voice, video, or IM clients from different vendors can communicate, and developers can freely add new capabilities. "The idea of SIP is to connect people to people, shielding those people from the details of how device is connected to device," says McGovern. In the future, for example, an IM message sent to a person who only has a cell phone available might be translated to speech for delivery.

Cisco has incorporated line-side SIP support into Cisco Unified CallManager 5.0, Cisco Unified CallManager Express 3.4, Cisco Unified Survivable Remote Site Telephony 3.4, and certain Cisco Unified IP phones and video endpoints. SIP is implemented natively in Cisco Unified CallManager as a base protocol. "Companies that use Cisco Unified CallManager can decide whether to use SIP alone, SCCP alone, or both," says Hadden-Boyd. Today SCCP has a richer feature set, but SIP will surpass it as vendors take advantage of the standard to add new functions. Hadden-Boyd expects that many customers will begin with SCCP and then gradually add SIP as they deploy SIP-enabled voice and video endpoints and applications, some with features not even imagined today.

Hadden-Boyd notes that Cisco Unified Communications includes SIP support not only on Cisco Unified CallManager, but also Cisco Unified CallManager Express and Cisco

New Appliance Model Option for Cisco Unified CallManager

Cisco is introducing a new appliance model, in which Cisco Unified CallManager comes preloaded on a server. Customers won't need to touch the underlying Linux operating system. For organizations that prefer an open operating system model, Cisco Unified CallManager is still available in a Windows operating system version, providing full visibility into the operating system and the option to customize certain aspects of call processing.

Survivable Remote Site Telephony (SRST). "The advantage of implementing SIP in SRST is that the call-processing system works the same way even if a WAN link between a centralized Cisco Unified CallManager server and a remote office goes down," she says. "Users can continue to use the same devices and applications."

Federated Presence

SIP support in Cisco Unified Presence Server also enables it to share presence information with other SIP-compliant presence servers, enabling a "federated presence." "If one department in a company has deployed a SIP-compliant presence server such as Microsoft Live Communications Server (LCS) or Lotus SameTime, and another department deploys Cisco Unified Presence Server, employees in all departments can use their respective client devices and applications to view presence information for all employees," says Burton.

The Network Knows

Multiple vendors are introducing unified communications clients, some approaching it from the desktop layer, others from the application layer. Only Cisco has approached presence from the network layer. One advantage is that the network is privy to presence information not only from the call-processing component (Cisco Unified CallManager) but also the applications layer (Cisco Unity Voice Mail and Cisco Unified MeetingPlace), creating a more complete picture of availability.

"Network-based presence information is also more accurate than calendar-based presence information because the calendar's owner might not have attended a scheduled meeting, for example, or double-booked just in case one canceled," adds McGovern.

Another advantage of the network-based approach is that the Cisco Unified Communications solution inherits all the security built into the IP network, avoiding the need for security investments specific to the voice system.

Holistic Management

Cisco Unified Communications improves IT staff productivity with two new tools for integrated network and system management. "The more types of devices and applications that we add to the IP network, the more important it is to have a holistic management system that addresses both the applications the underlying infrastructure of switches and routers," says Moran.

Cisco Unified Operations Manager provides a visual, dashboard-like interface for IT to monitor the real-time operational status of all components in the Cisco Unified Communications system, including Cisco Unified CallManager, Cisco Unified IP phones, Cisco Unity Messaging, Cisco Unified MeetingPlace, and underlying Cisco routers and switches. It shows the real-time relationship among solution components, such as which Cisco Unified CallManager server is associated with a particular set of IP phones.

"I can find out the status of any phone in the company in about a minute simply by entering its extension," says Mike DeDecker, network administrator for Warner Pacific Insurance Services, which uses the Cisco Unified Communications management software.

The new Cisco Unified Service Monitor provides real-time reporting of voice quality anywhere in the network. It captures and analyzes statistics about the voice session and reports them as a Mean Opinion Score (MOS) every 60 seconds.

Taking Fuller Advantage of the Network

"The guiding principle in Cisco Unified Communications is to take advantage of the network's awareness and capabilities for more effective communication and streamlined business processes," says Hadden-Boyd. "Cisco has already taken the intelligence and security of the network and applied it to the call-processing layer. With Cisco Unified Communications we are using SIP, SIMPLE, and other standards to extend network intelligence and security to the application layer." P

Further Reading

- Cisco Unified Communications system
  cisco.com/go/unified
- Cisco Unified Presence Server video
  cisco.com/packet/182_6a1
- Podcast with Cisco Distinguished Engineer Cullen Jennings
  cisco.com/packet/182_6a2

THE NETWORK Lifecycle

BEST PRACTICES FOR ORGANIZATIONS DEPLOYING ADVANCED TECHNOLOGIES

by joanna holmes

FORWARD THINKERS and early technology adopters, rejoice. Cisco Lifecycle Services promises to take much of the pain out of deploying new advanced technologies such as network security or IP communications.

The Cisco Lifecycle Services approach is a framework of best practices that outlines the critical activities required to deploy and operate an advanced technology. This detailed methodology is based on six lifecycle phases—prepare, plan, design, implement, operate, and optimize. Galvanizing this approach is Cisco's strategy to ensure that its customers have access to highly qualified channel partners to see them through each phase.

Karl Meulema, Cisco's vice president of services marketing and channels, is credited with expanding on the industry-recognized concept of incorporating technology deployment and operations into a lifecycle framework. To give life to this vision, Meulema developed sound methodologies that can be reproduced and scaled by enabling partners and then collaboratively supporting customers.

Packet: How do you define Cisco Lifecycle Services?

KM: Lifecycle Services is an accurate description of the activities that an end customer needs to be successful with Cisco technologies. In essence, it's about accelerating the learning curve around our technologies. It is not a service offering. I can't emphasize that strongly enough. It's a collection of best practices for a given technology, developed through our own experience and validated through that of our partners. Lifecycle Services is a very specific, documented set of best practices.

Packet: So how does it help Cisco customers?

KM: For any technology that a company such as Cisco

PHILIPPE LARDY

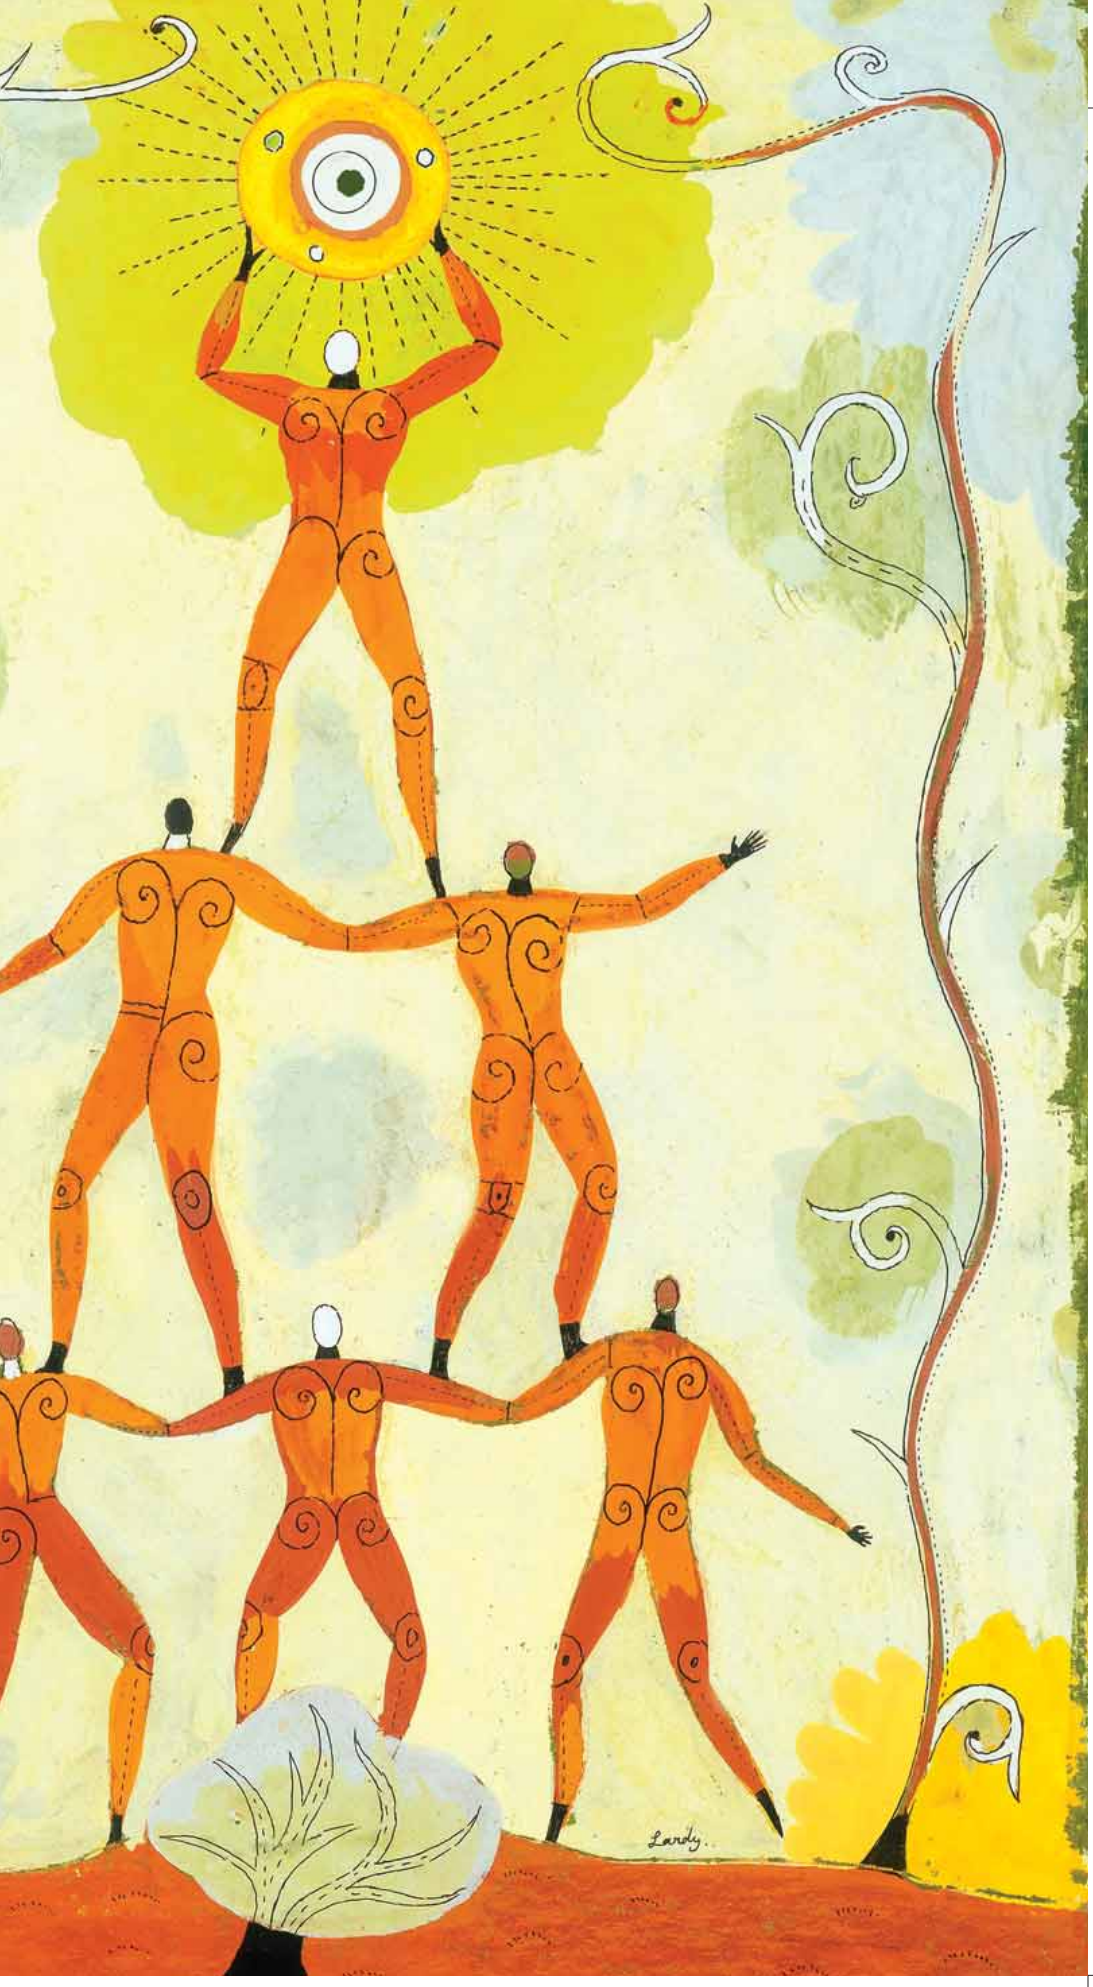

might introduce, if that technology is relatively new to the outside world, then there is limited or no experience in the marketplace to help make its deployment successful. As a technology matures, it commoditizes. People have taken every conceivable path to a potential end zone. What works and what doesn't work become general knowledge, and things become normalized in the process.

But if you're a customer just starting out with a new technology and the knowledge base isn't there yet, trying to map your own course to that end zone can waste a lot of cycles, money, and time.

Packet: And that's where Lifecycle Services comes in?

KM: Exactly. The goal is to shortcut that learning curve by building up our knowledge base very early in the product lifecycle, and then documenting it in a way that we can share it with others. So rather than expecting our customers and partners to figure it out by themselves, we share our knowledge of what actually works, through every stage of the network or solution lifecycle. We impart this knowledge to our partners in such a way that they become an extension of Cisco in supporting the customer.

If customers properly implement the methodologies and descriptions associated with a given lifecycle, we expect that they will be successful, because those methodologies and descriptions have led to previous successes. I think that's a key component, the fact that our foundation is not based on theory, but on proven best practices.

Packet: Will these methodologies be available for each new advanced technology that Cisco introduces?

KM: Yes. Every time we come out with a new advanced technology—or a major revision of an advanced tech-

nology—we will create lifecycle services to accompany it. And Cisco committed to come out with new advanced technologies like clockwork, so we'll continue to have many technologies that are not yet mature—and for those technologies, the importance of Lifecycle Services is enormous.

Packet: What about Cisco's channel partners? What's their role in these activities?

KM: The way we make Lifecycle Services a successful program for Cisco is through taking these activities and combining them with our go-to-market strategy, which is highly channel-driven. Between those two elements, we make sure we understand what's required to make Cisco partners capable and profitable, and what we need to add to augment the partners' capabilities in terms of our own services portfolio.

We've added Lifecycle Services to our Channel Partner Program as one of the key specialization criteria. So our strategy to ensure that our customers have access to capable, profitable partners is fully aligned with Lifecycle Services. And now, the ways in which our channel partners make margin and get discounts from us is directly tied to their demonstrable capabilities in delivering services that address the lifecycle as we define it.

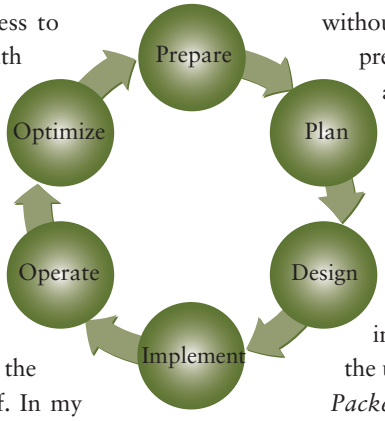

Packet: Who actually provides the services for customers? The partners?

KM: It could be the customers themselves, the partners, Cisco, or any combination thereof. In my organization we always say that we're about accelerating customer success with our technologies. Lifecycle Services is one of the best proof points for accelerating customer success, because—by its definition—we are precisely describing the customer success that our technology can encompass, regardless of who delivers it.

Customers can take on part of these activities themselves, although, realistically, few of them will. We very seldom see Cisco users actively engaging in the preparing, planning, designing, implementing, operating, and optimizing stages of their own networks and network solutions. For most customers, it's difficult to develop and maintain that skill set as a core competency.

Regardless, the "who" is not so important when we talk about service delivery. What is important is that between our partners, customers, and Cisco, we perform all the prescribed activities in the right sequence. Only then can we be confident that the solution will do what it's supposed to do, when it is supposed to do it.

The value that our customers are looking for is not that hot new technology—it's a business return for their solution. So while they're planning their return on investment, it's critical that the solution actually meets expectations.

In our experience, the customers who have been most successful in realizing strong returns on their investments are those who budgeted for the appropriate services. Often, we see buyers trying to contain the overall price of a solution by removing services from their final bid. The hidden cost of a failed or semi-failed deployment is far greater than the savings on these services, which are so critical for success.

Customers might feel that a solution can be successful without upfront assessments of the network, a comprehensive architecture, or ongoing optimization activities to ensure that the solution not only works, but continues evolving as the customers' business processes evolve. The problem with this approach is that they end up with a solution that doesn't do what they had expected of their investment throughout its lifecycle. In most of these instances, an array of hidden costs far outstrips the upfront savings on services.

Packet: If Lifecycle Services means helping Cisco users achieve successful deployments for new advanced technologies, what first step should those users take?

KM: That's an easy one. I'd like our users, when they think about implementing a solution, to think about the full solution, end to end. When they think this way, they include services in their earliest ideas about what it would take to make their technology deployment successful.

It's not just about a lot of very high-tech, very innovative products. It's also about planning for the right services upfront and having the solution delivered on time, meeting customers' requirements, and meeting the ultimate business needs. P

Current Supported Technologies

Cisco provides Lifecycle Services for Cisco Unified Communications, security, wireless LAN, and routing and switching.

To learn more about Cisco Lifecycle Services, visit **cisco.com/go/lifecycle**.
To learn more about Cisco services in general, visit **cisco.com/go/services**.

the connected

STATE, LOCAL, AND PROVINCIAL governments around the world are using wireless, IP telephony, and other network solutions to improve safety, enhance services, and foster educational excellence. Examples of such solutions dot these pages—and many more abound. Take Arlington County, Virginia, which delivers actionable information, applications, and voice services to public safety responders

IP PHONES with integrated displays in each classroom reduce the time teachers spend on routine tasks such as reporting attendance and ordering supplies.

VIDEO ARRAIGNMENTS with inmates housed in remote facilities save court staff time and reduce transport and security costs.

FROM THEIR squad cars, officers can access data such as warrants and file reports and timesheets, without returning to the police station.

LOCATION-BASED technology finds Bob, and he uses his mobile phone to e-mail urgent sales data to corporate headquarters.

AARON MESHON

Network technologies build and strengthen communities everywhere.

community

in the field using wireless connectivity, and broadcasts real-time training to staff in ten fire stations. Many of these connected communities save US$200,000+ in telecom costs annually, gain countless hours in productivity, and improve communications across the board. And all are improving their community quality of life. P

IPTV technology in helicopters transmits live video footage of major incidents to police stations in the city and neighboring jurisdictions.

INTERACTIVE, real-time video and video on demand increase staff efficiency and enhance student curricula.

PARKING METERS capture video of illegally parked cars with a timestamp; citizens who receive parking citations can view the video on the Internet.

AT A CITY HOTSPOT, building inspectors wirelessly submit reports and retrieve forms and data on their laptops, expediting the approval process.

WIRELESS
Hotspots

winning the security game

THE FINAL PIECES IN A FULLY DEPLOYABLE SELF-DEFENDING NETWORK

by janet kreiling

So many network endpoints; so many internal network components. So many applications with critical business depending on them. The time in which a threat can damage a network and a business has gone from days to seconds. Meanwhile, the amount of vigilance required to keep threats at bay has increased exponentially. The advantage, as in any good defense, comes from outsmarting and out-strategizing your opponent. In a technical game such as network security, it also comes from solid underlying technology.

As threats have grown more numerous and dangerous, technologies for detecting and defeating them have evolved too, moving from firewalls, antivirus programs, and spam detectors acting separately to integrated, collaborative, and adaptive activities that can respond virtually in real time.

"Our goal has been a completely deployable, Self-Defending Network that performs two broad tasks—automatically containing and controlling threats, and protecting connectivity," says Jeff Platon, vice president of security and application market management at Cisco. "This goal has now been achieved."

For a few years, Cisco has been introducing the pieces of a Self-Defending Network categorized by three levels of deterrence: to identify, prevent, and adapt to threats. The bottom level consists of products that detect and mitigate harmful activity at any endpoint and at multiple places within the network. These products are *integrated* into or with network components, so layers of network elements can deploy defenses and enforce policies. At the next level, the products *collaborate* with one another: an intrusion prevention device can tell an access control list (ACL) to deny access to a threat, for example. At the top level, products and technologies are *adaptive*. Some can dig into any packet and determine its points of origin and destination, its payload, and other data to assess the content or behavior of traffic flows or individual communications. Intruders are detected and thwarted at the gate, before they can get inside.

KEVIN TWOMEY

If you think of these products and technologies as the building blocks that form a Self-Defending Network, the Cisco Security Management Suite, introduced in February 2006, is the keystone. It's the final, central piece that holds everything together, according to Platon. It consists of two management tools that encompass operational control and police management across a network: the Cisco Security Manager and an enhanced version of the Cisco Security Monitoring, Analysis, and Response System (MARS). The two are interlinked, so that information from one can be used by the other. Their collaboration helps identify threats centrally and protect the network globally.

Also in February, Cisco announced new foundation blocks—reinforced protection for endpoints in the Cisco ASA 5500 Series Adaptive Security Appliance that make it the broadest, deepest anti-threat system available, as well as one capable of establishing up to 5,000 Secure Sockets Layer (SSL) or IP Security (IPsec) VPNs.

"We have certainly not finished developing security products," Platon says, "but for the first time, instead of just talking about the Self-Defending Network as a strategy, we can invite customers to deploy one."

SECURITY MANAGEMENT

NETWORK MANAGERS CAN configure firewall, VPN, and intrusion prevention system (IPS) policies on Cisco firewalls, routers, and appliances throughout the network using just one system: the Cisco Security Manager. And they can do it simply.

For example, configuring limits on a traffic type such as instant messaging across the entire network takes just a few clicks. This configuration can then be deployed on Cisco ASAs, IOS Software routers, PIX firewalls, Catalyst switches, or any Cisco system affected, even home-based routers used by teleworkers. The Cisco Security Manager separates policies from the devices enforcing them, so policies can be shared across the network without managers translating them for each device. This "policy abstraction and sharing" capability boosts productivity, ensures uniformity, and enforces compliance.

"Customers will no longer have to choose different point products and then find they don't talk to each other," says Calvin Chai, product marketing manager for the Cisco Security Management Suite. "In addition, network managers can use

the Cisco Security Manager to push configurations and policies to a large number of distributed destinations, such as retail stores. If you don't want employees at your stores to use the computers for peer-to-peer file sharing, you can define that policy one time centrally and then the Cisco Security Manager will ensure that it is configured on all the relevant devices. If a device is down for maintenance or power is out at that store, when the device starts up again, it asks for updates. It's not just a traditional push model."

While Cisco Security Manager configures the network, version 4.2 of the Cisco Security MARS monitors and analyzes it. Cisco Security MARS correlates the collection and reporting of security incidents across the entire network, often in real time, and recommends precision removal of threats. It shows network personnel visually where the affected point is in the network and the best point of mitigation.

For example, Cisco Security MARS might detect an unusual pattern of requests for data through a browser from a PDA. It shows the PDA and the VPN linking it to the network, the router it goes through, and whether the pattern suggests a threat or not. Upon perceiving a threat, it may recommend shutting down the device, VPN, or an interior network path, whatever is the best method to contain the offending element. Cisco Security MARS can also analyze the huge amounts of data it collects for regulatory and compliance reporting.

Each of these systems is valuable in itself, but they derive considerable additional benefits when used together.

"If a network manager sees an event on a firewall log on Cisco Security MARS and wants to know what caused it, he or she can select that log to see the network behavior involved, click on an icon, and automatically launch Cisco Security Manager to see the rule on the device that caused that event to be logged," says Chai. "The manager can also order an immediate fix. The ability to drill down across applications very efficiently saves time in troubleshooting and improves network availability significantly."

ANTI-EVERYTHING PROTECTION, CLEAN AND SIMPLE

One of Cisco's goals throughout its security product development has been simplicity of use, and version 7.1 of the Cisco ASA 5500 Series Adaptive Security Appliance combines in a single component both secure VPN connectivity and broad endpoint and network threat defense. A choice of four editions—the Anti-X, VPN, IPS, and firewall—lets users customize an ASA deployment to their security requirements. According to Joel McFarland, senior manager in the Security Technology Group at Cisco, the VPN edition is appropriate for enterprises seeking to protect links to remote workers; the Anti-X edition, for those primarily interested in protecting endpoints from network threats; the IPS edition, for those concerned more about servers and other interior network elements; and the firewall edition, for people who want the market-proven capabilities of the Cisco PIX Firewall in a next-generation appliance.

The Anti-X software, which is built into the Cisco Security and Control Security Services Module (CSC-SSM), enables the ASA 5500 Series system to protect all endpoints, whether they are parts of VPNs or not, from all types of threats: malware aimed at the endpoint itself, such as viruses, spam, phishing, and keyloggers as well as inappropriate URLs and offensive content; and malware aimed at the network, such as Trojan horses, worms, and spyware. Thus, one system can prevent employees from giving away company information to phishers or spyware inadvertently, downloading contaminated files or programs, receiving spam (with an extremely low rate of false positives), and accessing non-work-related Websites. All told, the CSC-SSM can protect information and resources as well as clients and identities; fend off hackers and denial-of-service attacks; guard e-mail, Web browsing, and file exchanges; and provide verification of communications for regulatory purposes.

Anti-X in Depth

"The Anti-X protection represents a major increase in gateway security, satisfying a very big customer need," says Joel McFarland of the Cisco Security Technology Group. "Cisco has combined Trend Micro's InterScan security suite with several of its own technologies, such as the PIX firewall and its IPS, to create a product unique in the market for its breadth and depth. For example, the Anti-X protection contains tens of thousands of signatures, rather than the several hundred typically available with antivirus products, and can counter threats that are still 'in the wild' as well as more established ones." Searching for threats against endpoints, it applies store-and-forward techniques to file-based communications such as e-mail, holding them until it verifies that they are clean, McFarland explains. Searching for threats against the network, it inspects traffic flows and even some individual packets. Granular policy controls can deny access at Levels 3 or 4 to individual endpoints, to endpoints at certain locations, at certain times of day or days of the week, or grant access to only certain parts of corporate resources.

Seeing the Network—Any Way You Want

The Cisco Adaptive Security Device Manager gives managers various views of the network: device, policy, and topology. The device view shows the security devices and what policies they can implement; the policy view, policies that can be customized to specific business needs; the topology view, various levels of network maps from which managers can implement policies at the different levels. For example, from the policy level, the network manager might select a policy regarding clientless access to certain databases and implement it on all Cisco ASA 5500 Series systems, which set up this type of VPN securely.

UNIVERSAL SECURE ACCESS

THE CISCO ASA 5500 SERIES VPN edition supports both remote users and site-to-site links among corporate locations. Remote users can establish either SSL VPNs through dynamically downloaded client software or IPsec VPNs through pre-installed client software. The ASA also supports clientless VPNs over the Internet, which is especially helpful for remote users who are not using corporate-owned devices for communications, such as business partners, or employees using public Internet terminals or home PCs. Version 7.1 brings to 2,500 the number of SSL and IPsec VPNs that can be supported simultaneously in large enterprises, five times the number previously possible (support for 5,000 simultaneous sessions will be available soon). Built-in load balancing for multiple systems enables companies to support tens of thousands of concurrent users, and the single system eliminates equipment previously needed for the different types of VPNs for better manageability and lower costs.

The depth and amount of features available for remote access VPNs are constantly expanding, according to Pete Davis, product line manager for remote access VPNs at Cisco. Version 7.1 includes the Cisco Secure Desktop, which guards SSL endpoints. Whether the user is working from a corporate-managed PC, a personal device, or a public terminal, before the ASA 5500 Series system grants SSL VPN access, the Cisco Secure Desktop checks for antivirus software on the unit and for certain types of malware such as a keystroke logger. If, after the session begins, the Cisco Secure Desktop detects a malicious program, it prompts the user to stop the session. If the endpoint is corporate-owned, the Cisco Secure Desktop can implement access policies for it. After the session is finished, the system sanitizes it using an algorithm developed by the US Department of Defense.

The ability to set up and secure clientless VPNs, Davis points out, "is very useful when enterprises want to permit an employee using a noncompany-owned device or an outsider, such as a partner, to access a subset of Web corporate resources. The technology is highly compatible with Web browsers, so it's easy for outsiders to connect to the network." The system prevents data sent to a clientless user from being left on a remote system by checking to see whether the device is corporate-owned, if there's a keystroke logger running on it (depending on policy instructions, the system can deny access), encrypting data, and finally, wiping the device clean of corporate data when the session ends.

The Cisco clientless VPN technology also includes support for accessing legacy applications through the Citrix Metaframe application server, adds Davis. Cisco's support does not require downloading any additional software, so sessions begin faster and there's less risk of conflict with other software on the PC.

Another new feature on version 7.1 of the Cisco ASA 5500 Series is safe caching for remote Web users. Remote users benefit from caching when working with Web applications to speed up their interactions, but if they're coming in over the Web through an untrusted device, the Web browser might convey intruders into the cache and from there into the network. Now, they can keep the feature on while communicating through a VPN and rely on the Cisco Secure Desktop to eliminate dangers in data intended for caching.

In addition to the ASA 5500 Series, SSL VPN and Cisco Secure Desktop capabilities are also available on Cisco 800, 1800, 2800, and 3800 series Integrated Services Routers and the Cisco 7200 Series and 7301 IOS routers, which support from 10 to 150 simultaneous sessions. Enterprises using the ASA 5500 Series VPN edition can purchase the Anti-X module (CSC-SSM) for complete endpoint protection.

"The overall goals of the Self-Defending Network are to align security provisions with business requirements, use IT investments to take the offensive against threats to the business, reduce the complexity of the IT environment overall, and to achieve visibility and control over intruders," says Platon. "Rather than just a defensive posture, excellent security can help business processes to become more effective by removing barriers to internal and external communications and allowing employees better access to corporate applications and data. And in the process, it can save costs in protecting the business and in doing business." P

Further Reading

- Security Management Suite
  cisco.com/go/security_management
- Cisco ASA 5500 Series
  cisco.com/go/asa
- Cisco Adaptive Security Device Manager
  cisco.com/go/asdm

Switching into Overdrive

REAL-TIME APPLICATIONS ARE DRIVING NEXT-GENERATION WIRING CLOSETS.

by joanna holmes

Changing characteristics at the network edge are driving new requirements for wiring closet switches. A shift in applications and security needs over the last several years has many network managers wringing their hands, trying to placate their users while solving a spate of new problems with limited resources • "Many organizations devote a great deal of time to designing their core and distribution networks," says Marie Hattar, director of enterprise switching in the Product and Technology Marketing Organization at Cisco. "But the current industry trend shows services and applications moving closer to the network edge. Our users need to start focusing on the wiring closet as a strategic point on the network." • "Typically, people apply higher-level services in the network core, where a greater concentration of outgoing traffic exists," says Fred Weiller, a marketing manager for enterprise switching at Cisco. "Now, however, applications increasingly reside at the network edge, and they're moving from a client/server model to a more distributed, any-to-any model." New kinds of applications such as IP telephony and Cisco Unified Communications are part of the changing landscape (see story on page 36). To address the requirements these applications introduce, network managers need to scrutinize how they deploy their services. The goal is to think strategically about consistent services across the entire network.

"We're driving more and more real-time applications from the wiring closet, and if reliability, security, and applications performance are not there, that affects a company's ability to do business," says Hattar. "There's a need for intelligence in the closet, because it cannot continue to be underserved in the new computing environment."

According to Weiller, the wiring closet is changing because today's networks were largely designed for client/server-based data applications. They are data centric and were built for delay-tolerant client/server applications, with clients at the edge making requests to servers in the core or data center. And that, says Weiller, has led to the current centralized architecture and hierarchical format of these networks. They offer strong performance, intelligence, traffic management, and security capabilities in the core, some of which are mirrored in the distribution layers—but the wiring closets show neglect.

GETTY IMAGES

REAL-TIME APPLICATIONS like VoIP and network security are driving more traffic to the network edge, bringing the need for more intelligence.

There's a need for intelligence in the closet, because it cannot continue to be underserved in the new computing environment.

MARIE HATTAR, DIRECTOR, CISCO PRODUCT AND TECHNOLOGY MARKETING

"There are many more ways to communicate than there were five years ago," Weiller adds. "Devices in the wiring closet now include desktop video cameras, surveillance cameras, IP telephony, and collaborative applications, in wireless or wired environments."

The requirements that all these new applications place on the network call for some rethinking. A service such as IP telephony for instance, requires very high availability. And like other real-time applications, it also requires very consistent quality of service (QoS). For that reason, Weiller says, making sure you can delilver applications performance and protect strategic applications are also key to the successful deployment of real-time applications at the edge.

Security, too, is a vital factor as enterprises deploy more types of autonomous devices and applications in campus LANs. A new device within the campus network can provide a back door to the information system. And in this newly collaborative enterprise network, users' laptops pose a constant threat of spreading viruses and other malware to the rest of the network. "As part of the redefined requirements," Weiller says, "enforcing the security structure at the network edge includes making sure you know who's connecting, and ensuring that all your users are running the right kind of security tools and software."

Services in the Wiring Closet

FROM THE CISCO perspective, wiring-closet switching capabilities for next-generation networks fall into four categories: *business continuity, integrated security, applications performance*, and *simplified operations*.

BUSINESS CONTINUITY

"Deterministic business continuity is the goal of this high availability feature," says Weiller of Cisco's routed access capability. Simply put, routed access is the ability to apply routing starting from the desktop, the first touch point of the network. Routed access enables very high availability by providing ultra-quick recovery—less than 200 milliseconds—in the event of a failure. It works transparently, so end users are never aware of network hiccups, even for real-time applications. From the network manager's perspective, routed access offers simplified troubleshooting and network implementation, because it runs on only one control protocol—Open Shortest Path First (OSPF) or Enhanced Interior Gateway Routing Protocol (EIGRP).

"A very small percentage of our customers has deployed routing at the access layer," says Weiller. "Many don't realize that the Cisco Catalyst switches they have in their wiring closets today support this Layer 3 capability—they just need to turn it on."

One Cisco customer that recently began using routed access is San Antonio Water System (SAWS), a water company serving about one million people in San Antonio, Texas. SAWS had

Product Profile: Cisco Catalyst 4500 Series

The Cisco Catalyst 4500 Series is the industry's most widely deployed modular platform. For the wiring closet, it offers a high-performance platform with a comprehensive set of security and high-availability features to ensure a safe, nonstop networking foundation. Innovative security features such as 802.1X, Network Admission Control (NAC), man-in-the-middle attack mitigation, and security anomaly detection help to ensure that security threats are stopped at the network edge.

The 4500 architecture carries a raft of high availability hardware and software features, such as redundant power supplies and Control Plane Policing to prevent CPU overload during denial-of-service (DoS) attacks. Dual supervisors with nonstop forwarding and stateful switchover deliver 50-ms failover, ensuring seamless recovery for time-sensitive applications such as voice. This feature set ensures that all ports have 100 percent redundancy in case of a key component failure.

For customers planning to converge their networks, Cisco Catalyst 4500 switches provide convergence features such as 15.4-watt Power over Ethernet (PoE) on all ports, QoS to differentiate time-sensitive traffic, and multicast features for video to the desktop. The 4500 is a highly flexible platform with deployment and operational simplicity, and a secure, feature-rich switching solution designed for lasting investments.

been using the Hot Standby Router Protocol (HSRP) as a failover system, with fairly traditional Layer 2 access and Layer 3 distribution layers. The access switches were dual-homed into distribution switches, with load bearing on only one link.

"With more bandwidth-intensive applications and more applications at the desktop, our utilization on our uplink ports was getting pretty high," says Darrin Gannaway, senior network engineer at SAWS. "Now we're running EIGRP and letting it load balance across both links simultaneously. That was one of the benefits we saw right away."

By moving to a routed access layer, SAWS was able to establish a common set of protocols that would run from end to end, from core to campus. With the extra intelligence introduced by routing, Gannaway and his team enabled faster recovery time and delivered more intelligent traffic management for SAWS real-time applications.

INTEGRATED SECURITY

The wiring closet must act as a line of defense against anyone capable of causing network threats. Among the new services that today's wiring closets need to support are perimeter defense, identity-based trust and identity management, connectivity services, and secure management. Historically, these services have often arrived on the network via separate appliances. The new trend, however, calls for *integrated* security to improve management and simplify network operations.

When SAWS deployed routed access at the edge of its network, Gannaway was pleased with his newly enabled security capabilities. "In the past, when we were at Layer 2, even unwanted traffic would go all the way up to the distribution layer before it was filtered. We can employ better security controls now, because with the routed access model, I can run ACLs [access control lists] at the access layer, right at the entrance point. There's just more inherent security with having the control right there on the port—and [at the same time], the routed model helps reduce our congestion, too."

APPLICATIONS PERFORMANCE

Other vital services that help make up the toolset for this new access layer are multicast and QoS capabilities. Layer 3 intelligence provided by Cisco Catalyst switches can ensure that applications, no matter how delay sensitive or real time in nature, always get the throughput they need.

In the SAWS network, one of the most critical business applications is a database system called Maximo. "It's very bandwidth-intensive," says Gannaway, "and it doesn't like to share bandwidth with any other application." Advanced QoS capabilities were necessary to ensure that Maximo's mission-critical traffic and real-time IP telephony traffic never experienced delays. Furthermore, with the new routed access model, the switches run EIGRP, which supports unequal cost load balancing. "Sometimes I don't want all the traffic being equally divided," says Gannaway. With EIGRP, he can customize it as he chooses.

SIMPLIFIED OPERATIONS

In addition to providing the rich services and strong defenses needed within the next-generation wiring closet, consolidating on Cisco Catalyst platforms from the core to the wiring closet is vital for consistent services and lower operational costs. Networks are highly sophisticated, but new technologies, products, and processes help simplify network operations. Some touchstones of simplicity to aim for include configuration automation, standardization on fewer device platforms and protocols, and intelligent power management.

For Gannaway, network simplification was a welcome bonus of migration to routed access. "Now everything in our environment, whether it's a router or a Layer 3 switch, relies on routed technology." He also estimates that the time his team now spends on configuring and maintaining roughly 300 network devices has been reduced by 25 to 30 percent.

The Big Picture

"WHEN YOU PUT ALL THESE CAPABILITIES together, you get a service that's always on, with no threats permeating the information system," summarizes Weiller. "Then you combine that with simplification of operation, and you arrive at a very solid foundation that not only improves handling of the applications you're already using, but also enables immediate and future deployment of new applications."

In support of the switching capabilities needed for new wiring-closet requirements, Cisco provides three product lines:

- The Cisco Catalyst 3560 and 3750 series provide new baseline services.
- The Cisco Catalyst 4500 Series provides higher-grade services and higher availability.
- The Cisco Catalyst 6500 Series provides the highest levels of service, offering data center-grade switching services for the wiring closet. P

Network Assessment

A good first step in preparing your wiring closet for next-generation services is a general network assessment. The Cisco Discovery Tool (**cisco.com/go/partner-discovery**), a PC application, can create a record of all connected network devices, including product platforms and operating system versions. This application is available from Cisco account managers or channel partners.

Cisco IP over DWDM

DELIVERING NETWORK CONVERGENCE FOR CAPACITY GROWTH AND OPERATIONAL EFFICIENCIES

by errol roberts, ori gerstel, and tony sarathchandra

Two major events are driving the growth demands in core capacity: network convergence and traffic expansion, predominantly from video—the new "killer application"—in its various forms including IPTV, business video, managed video, and Triple Play. Within global service provider markets, the trend is for legacy services such as Frame Relay and ATM switched services to migrate to Multiprotocol Label Switching (MPLS)-based services (e.g., pseudowires and Carrier Ethernet). Meanwhile, several service providers are converging both their business and residential access network into a single, consolidated next-generation network, demanding a higher bandwidth capacity core.

The trend toward network convergence gives service providers and large enterprises an opportunity to reduce the complexity of their network by consolidating optical and IP/MPLS network equipment. Core (long haul) and metro optical today are primarily fully amortized assets built on legacy SONET/SDH protocols. In the early 1990s, SONET and SDH became the de facto technology for efficiently and reliably transporting voice and data using time-division multiplexing (TDM).

SONET/SDH performs three critical functions: it grooms traffic from slower links, such as T1/E1 and T3/E3, onto higher-speed links; it protects against failures and restores a traffic link or path in the event of a failure; and it provides mechanisms for centralized administration, monitoring, and control.

DWDM Transport

DENSE WAVELENGTH-DIVISION multiplexing (DWDM) was developed to dramatically increase the information-carrying capacity of fiber networks by multiplexing many wavelengths over one physical fiber for long haul and metro networks. But service providers and large enterprises were put in the position of having to invest in multiple layers of network equipment to meet the increasing IP traffic growth. Three required layers of equipment include an IP/MPLS router, a SONET/SDH add/drop multiplexing (ADM) transport, and then a DWDM infrastructure. This setup is inefficient to operate, complex to provision and maintain, and requires considerable operational complexity to configure, manage, and troubleshoot.

Today, not surprisingly, both service providers and large enterprises are exploring ways to not only reduce their network capital and operating expenditures but to accomplish new service rollouts efficiently. Any new solution must meet the current business and technical requirements of core and metro networks.

Because service provider regional and core points of presence (POPs) are typically many hundreds of kilometers apart, router interfaces need strong error detection and correction mechanisms. Technical needs for metro networks are a little less stringent, because distances between locations are shorter, typically less than 100 km. Service providers also want any core solution to integrate into their existing Operations, Administration, Maintenance, and Provisioning (OAM&P) business processes, so optical networks must conform to the ITU G.709 standard.

With the current generation of core and aggregation routers terminating multiple Gigabit Ethernet and 10 Gigabit Ethernet of IP traffic, the grooming formerly provided by SONET/SDH is not as widely required. Given the trend toward higher capacity interfaces such as Gigabit Ethernet, the operational support and traffic protection can now be performed by routers and their associated interfaces through alternative mechanisms such as G.709 framing and Fast ReRoute (FRR). In metro networks, operators take advantage of Equal Cost Multipath (ECMP)

DWDM IS THE unifying technology that consolidates IP applications over a common transport infrastructure.

GETTY IMAGES

load balancing or EtherChannel on the router/switch platform. So, there is now an opportunity to reliably and efficiently carry IP traffic directly over the DWDM transport, without the need for an additional layer of network elements.

Cisco IP over DWDM Solution

IN DECEMBER 2005, CISCO announced an IP over DWDM (IPoDWDM) solution, which takes advantage of integration between the IP and DWDM network layer elements to help manage expenses, improve network reliability, and increase time to market and speed to service. This announcement included an industry first for core networks: integrated, tunable four-port 10 Gigabit Ethernet and one-port 40-Gbit/s DWDM interfaces for the Cisco CRS-1 Carrier Routing System.

The Cisco CRS-1 1-Port OC-768c/STM-256c Tunable WDMPOS Interface Module uses a modulation scheme that provides up to 40 Gbit/s of data throughput across existing 10-Gbit/s fiber infrastructures. The 4-Port 10 Gigabit Ethernet Tunable WDMPHY Interface Module also enables the Cisco CRS-1 to interconnect directly with existing 10-Gbit/s certified DWDM systems and fiber plants while providing OAM&P capabilities akin to SONET/SDH at typical Ethernet price points.

Both modules are fully tunable across the ITU C band with 50-GHz spacing and support high-gain Enhanced Forward Error Correction (EFEC), extending reach up to 1,000 km without the additional signal regeneration required by typical core applications.

For metro networks, Cisco offers a range of integrated pluggable optical interface modules for Cisco Catalyst switches and Cisco 7600 Series routers. Cisco DWDM XENPAK supports 10GBASE Ethernet, and the Cisco DWDM Gigabit Interface Converter (GBIC) supports Gigabit Ethernet. Each cost-effective, high availability interface module supports hot insertion and can be ordered in one of 32 fixed wavelengths.

These integrated, ITU-compliant interfaces allow Cisco routers and switches to interconnect directly with existing DWDM systems installed at the edge of metro and core networks without the need for additional costly transponder equipment at these locations.

Rounding out the solution is the Cisco ONS 15454 Multiservice Transport Platform (MSTP). MSTP supports point-to-point or ring DWDM topology, and integrates photonic switching via Reconfigurable Optical Add/Drop Multiplexer (ROADM) functionality. Its highly automated capabilities have made this platform a leader in the industry for ease of planning, deployment, and operation.

Reduced Capital Costs

CISCO'S IPoDWDM SOLUTION can eliminate a significant amount of costly network equipment in a POP. Historically, optical connections within POPs have used standard short-reach 1,310nm optics (or "gray light"). Long-reach or "colored" optical signals enter the POP typically through 10-Gbit/s SONET/SDH circuits multiplexed via DWDM onto a physical fiber.

Traditionally, these signals are demultiplexed and fed into racks of transponders that convert them from their current optical form to electrical to short-reach optical (this process is called an OEO conversion). The signals are then fed into SONET/SDH cross-connect equipment or a patch panel before being terminated on a router or switched through to another POP location.

With transponders integrated directly into the Cisco CRS-1, and with MSTP integrated ROADM functionality, additional external transponders and optical cross-connect

Integrated or Segmented Management?

The Cisco IPoDWDM solution supports different network management models for service providers, depending on their preference. Because current service providers' operational organizations are often segmented into IP and transport, OAM&P tools available on the CRS-1 and MSTP provide for both a unified or segmented administration mode.

The transport network operators can access and control the optical network configuration, and the IP operations staff can manage and control the Layer 2 and 3 network configuration on the CRS-1 system. The tools are interoperable with third-party management systems or service provider-developed operational support systems through the support of TL-1, CORBA, and Simple Network Management Protocol (SNMP) for northbound interfaces. Communication is via Extensible Markup Language (XML), direct command-line interface (CLI), and HTTP/S.

Similarly for metro networks, service providers can choose to have IP network operators oversee the passive optical network with integrated pluggable optics on Cisco routers and switches. If they require a more flexible metro optical network, they can use MSTP, and transport network operators can monitor and manage both MSTP and the DWDM pluggable optics.

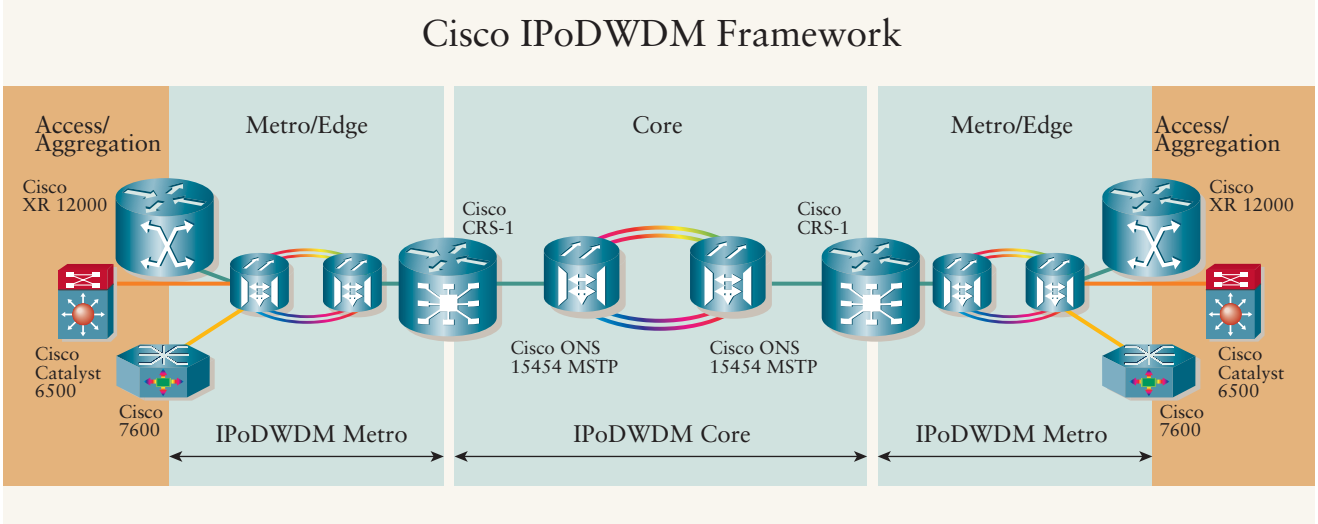

OPTICAL ADVANTAGE The Cisco IPoDWDM solution integrates long-haul transponders within the CRS-1 and provides managed wavelengths across the ROADM for scalable bandwidth, reduced complexity, and faster time to market with services.

or patch panel equipment are eliminated. This approach reduces capital outlay by as much as 50 percent in some cases and decreases network complexity.

Lowered Operating Expenses

SIMPLIFYING NETWORKS USING a Cisco IPoDWDM solution can also lower network operating expenses for service providers and large enterprises. Eliminating standalone transponders decreases power, cooling, and rack space requirements in a central office or POP. Reducing the number of active network components also decreases potential network failure points, which boosts network reliability and lowers maintenance costs.

Service providers can further simplify core operations by using an integrated management approach across the IP and transport networks (for more on management, see the sidebar opposite page). Cisco has developed network operation lifecycle tools that help facilitate planning, installation, and operation of metro and core networks. One aspect to these tools is that they combine wavelength provisioning from the DWDM port on the router through the optical transport layer. Traditionally, changes to traffic patterns must be engineered in advance, and in-service changes such as adding new filters in the POP need to be done manually and typically disrupt services. Using MSTP, service providers and large enterprises can remotely add, drop, or change available fiber wavelengths to respond to changing network conditions or services.

The ability to reconfigure the optical network remotely eliminates truck rolls, which, in turn increases service flexibility and decreases operations costs. It also eliminates error-prone manual operations that can cause disruptions in existing services.

Improved Network Resilience

THE IPoDWDM FUNCTIONALITY in the Cisco CRS-1 conforms to the ITU G.709 Optical Transport Network (OTN) standard, fully supporting the management and performance management aspect of typical SONET/SDH connections while interfacing directly into the DWDM layer. With this enhanced visibility into the DWDM layer, the Cisco CRS-1 can potentially proactively monitor all optical paths end to end, and detect gradually degrading wavelength paths caused by environmental conditions, hardware degradation, or other factors.

By monitoring the Forward Error Correction (FEC) code, the Cisco CRS-1 will identify when the rate of corrected errors justifies preemptive action at the IP/MPLS layer. When that threshold is reached, the router will trigger reconfiguration behavior such as FRR, an MPLS function, to switch over to a backup path before the primary path breaks down, avoiding any traffic interruption.

This mechanism is comparable to or better than traditional SONET/SDH-based failover mechanisms, because it meets or exceeds the 50-ms threshold while using less framing overhead compared to SONET/SDH.

In the scenario described above, the FEC extends beyond FEC defined in the G.709 standard (GFEC). The router interface can be provisioned to use an EFEC. EFEC allows the DWDM signal on Cisco CRS-1 Carrier Routing System interfaces to reach up to distances of about 1,000 km (621 miles) without optical regeneration—an industry first.

For metro applications, Cisco ITU pluggable optical interface modules support Digital Optical Monitoring (DOM), an industry-wide standard that defines a digital interface to access and monitor the operating parameters of transceivers. With DOM, metro network operators can perform in-service

transceiver monitoring and troubleshooting operations, ensuring a more available yet highly cost-effective network.

In addition to the performance monitoring characteristics of the ITU service interfaces provided by DOM and G.709, an intelligent photonic layer based on MSTP will provide end-to-end wavelength power management and visibility.

IPoDWDM in Practice

SERVICE PROVIDERS AND SOME large enterprises are already using Cisco IPoDWDM solutions to simplify their networks and reduce expenses, improve network reliability, and increase speed, scalability, and flexibility for IP-rich media applications. Comcast Cable, for example, is the largest US provider of cable services, and is expanding its cable operations to deliver digital services, provide faster Internet and IP-enabled phone service and innovative programming.

"Comcast's and Cisco's shared vision of integrated 10-Gbit/s and 40-Gbit/s DWDM interfaces was one of the key reasons we selected the CRS-1," according to Vik Saxena, Comcast Director of IP Architecture. "As we move to transport all services over IP, including broadcast video and video on demand, cost-effectiveness, scalability, reliability, and service flexibility become critical. By meeting these important needs, Cisco's CRS-1 is an ideal platform on which to base our service growth." P

Further Reading

- Cisco IPoDWDM Solution for IP NGN
  cisco.com/packet/182_7b1
- Cisco CRS-1 Carrier Routing System
  cisco.com/packet/182_7b2
- Cisco ONS 15454 Multiservice Transport Platform
  cisco.com/packet/182_7b3

ERROL ROBERTS is a Cisco Distinguished Systems Engineer with a focus on optical and data center technologies. He can be reached at eroberts@cisco.com.

ORI GERSTEL has been leading optical networking architecture efforts for 15 years and currently manages the advanced technology team in Cisco's Carrier Routing Business Unit. He can be reached at ogerstel@cisco.com.

TONY SARATHCHANDRA is a product marketing manager with the service provider routing and switching team at Cisco. His current focus is in the core routing arena encompassing the Cisco CRS-1 platform. He can be reached at tsarathc@cisco.com.

Service Control Benefits ISP

JAPAN'S PLALA NETWORKS DELIVERS SECURITY AND "PEACE OF MIND" TO ALL ITS BROADBAND CUSTOMERS.

Too often in the Internet world, the people who need protection most are those least likely to have it. People who consider computers intuitive have no trouble finding and loading antivirus, antispyware, URL filtering, and other safeguards. But many people do not know how and therefore do not do it—even when the software is made easily available by their ISP. • That was the experience of Plala Networks, headquartered in Tokyo. Its goal for the three-quarters of its 2 million subscribers who have broadband service is to enable them to use the Internet "securely and with peace of mind." Initially Plala offered e-mail virus checking and encryption as optional services, but found that these services were adopted only by customers who understood to a certain degree the dangers of the Internet. Novices without that understanding typically did not opt for the protection that these services provided.

According to Katsumi Nagata, director of Plala Networks and general manager of System Development and Network Engineering, "There is an increasing number of people who do not know what a virus is or what they should do to prevent children from accessing pornographic, violent, or other inappropriate sites. However, we would like such people to use the Internet more securely. To accomplish this goal, a potential solution is useless unless it is made the default or offered free of charge. We wanted to make a system in which our least-experienced users can be protected from the outset."

Plala's solution was to eventually install Cisco Service Control Engine (SCE) 2020 systems at all regional access aggregation points. On this platform, the company makes available its "Net Barrier Basic" service, which includes URL and packet filtering. These two features help protect users from inappropriate Websites and content and assist in detecting security breaches. The platform yields other benefits as well. Plala can also limit applications that consume large amounts of bandwidth, so all subscribers find high-speed Internet access readily available.

In 2003 Plala Networks, which has the highest customer satisfaction rating among Japanese broadband providers, began working toward providing all subscribers with a dramatically enhanced quality experience. Plala's first step was to optimize its network by managing peer-to-peer (P2P) traffic using the Cisco SCE 1010 Service Control Engine (previously

ACHILLE BIGLIARDI PHOTOGRAPHY

GOT HER COVERED
URL and packet filtering help Plala keep harmful content and security breaches away from its subscribers.

known as the SE 1000). Plala used the SCE initially to discover what types of traffic were traveling over its network and tying up bandwidth—finding a good deal of congestion from P2P file-sharing services—and then to impose limits on those applications. Everyone is affected when a few users consume more than their fair share of resources for P2P downloads. "With the Cisco Service Control Solution, we were able to create an environment that allows everyone to share resources fairly," says Shinya Adachi, manager of the Network Engineering Department at Plala.

Application-Level Packet Control with SCE 2020

THE SCE 2020, Cisco's flagship model for its service control solutions, uses deep packet inspection to identify protocols for voice over IP, video, gaming, instant messaging, P2P file sharing, and others—any and all applications found in a typical broadband provider's network.

"The system can identify over 900 protocols, anywhere from simple, port-based ones to the more advanced protocols that require Layer 7 information, such as P2P application traffic and voice over IP," says Moti Beharav, Cisco theater manager for SCE business in Japan. Because many new protocols use well-known ports for masking or use port-hopping, Beharav adds, it's not always possible to know what application is actually being used without access to the Layer 7 information. "We've seen applications such as P2P and video on port 80 originally used for HTTP, for example, and a lot of other applications on ports that weren't intended for them."

The SCE 2020 controls and limits bandwidth by applying policy rules to application traffic on a per-subscriber basis if required. For example, it can limit the bandwidth used for file sharing while giving priority to voice over IP, video, or gaming—any application that's sensitive to delays, Beharav explains. It can also redirect traffic that violates the programmed policies.

"The Cisco SCE 2020 can perform judgments at an application level, making it possible to block traffic that is spoofing the port number," says Yasuhiro Horike, manager of the System Development Department at Plala. "With this ability, we can protect our network from performance degradations. High performance is a vital contributor to our reputation for leading customer satisfaction."

A Safe, Well-Lit Internet Experience

NEXT, PLALA TURNED TO the matter of providing better security for all its subscribers, and in 2005 instituted clientless access to its Internet service, simultaneously upgrading to the SCE 2020. In its network, a Cisco SCE 2020 is located at each point where regional access is aggregated and is connected to a management server (see figure).

Now, as Adachi explains, "Our IP addresses change each

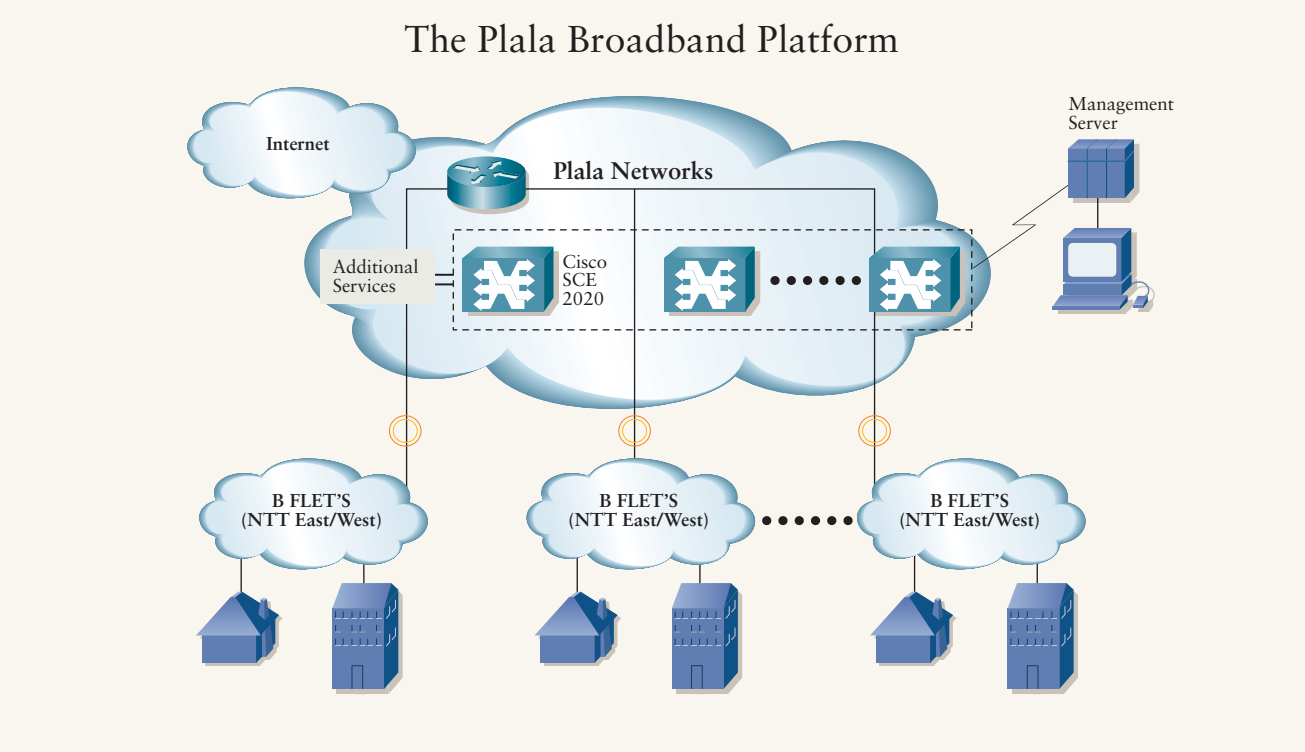

NET BARRIER BASIC Plala's default security service depends on installing the Cisco SCE 2020 at access aggregation points, linked to a management server. Together the two negotiate ever-changing IP addresses crucial to provisioning the service.

time a user connects. So the management server linked to the SCE performs the role of identifying each customer, even when the IP address changes, and sends the customer's settings to the SCE 2020 so it will provide the appropriate services."

The Cisco SCE 2020 can serve subscribers in this network configuration because it is subscriber-aware. It is able to obtain real-time information on the IP address associated with a given subscriber at any given time and then take action or enforce policies for each subscriber and the application he or she is using. This capability allows the SCE 2020 to deliver Plala's Net Barrier Basic service, which includes URL and packet filtering, even though users have not installed any protective client software on their PCs—and to provide these features as a default service, free of charge.

The URL filtering enables Plala to block access to pornographic sites or sites with inappropriate content. Packet filtering enables the company to detect behavior that indicates the presence of worms and viruses and take steps to seal them out of the network. The filtering is transparent and does not affect network performance. "Since we first used the Cisco SCE 1010 for bandwidth control, we have valued its ability to deliver extremely high throughput," says Horike. "Now that we've moved to the SCE 2020, we do not experience any reductions in performance even when using URL filtering and packet filtering." The purpose-built system can track and manage 2 million concurrent unidirectional flows, and the overall solution is designed to scale even higher.

The system gives Plala Networks a platform for adding future security and other services.

"We would like to provide security measures that allow us to logically determine responses for any discovery of a security problem or vulnerability," according to Nagata. "We would also like to offer a system that provides even greater protection by warning customers of infection by worms or viruses and notifying them of how to respond to the situation. By providing more secure services, we would like to let people use the Internet with a greater sense of peace of mind."

In step with Plala's goal, the Cisco SCE 2020 can identify patterns and behaviors that indicate worms and viruses. "If it finds a behavior—say a massive amount of outgoing e-mail traffic—that falls into the category of a worm," Beharav explains, "the provider can redirect the traffic to warn the customer of the infection and ask him or her to download the correct patch from a security Website. In the meantime, the system protects the rest of the network."

Commitment to Security

PLALA'S COMMITMENT TO security is evident from its receiving both the Information Security Management System certification from BRS, an international accreditation body, and the Privacy Mark certification, a Japanese standard for protecting personal information. These awards require the company to continually review all operational and system aspects of internal security.

Customers are happy with its service, and Plala finds itself with a platform for growth. "In a way, Net Barrier Basic is a uniform service," Nagata says. "The main theme was to provide the service by default and free of charge, but some users want a more advanced service. For example, with URL filtering, some people want greater customization and basically wish to block violence and pornography but still allow certain items. The system that we have created will provide the necessary foundation for delivering this level of customized service." P

Solid Platform for the Future

Deploying the Cisco SCE 2020 has yielded several business benefits for Plala Networks:

A consistent broadband experience. In controlling consumption of bandwidth by P2P file sharing, Plala has created a network in which users can fairly share resources.

Predictable, safe broadband access. With its Net Barrier Basic service, intrusions and worms can be blocked at the application level, and URL filtering denies access to inappropriate Websites.

Optimized support services. As fewer users need technical support to get rid of viruses and other intruders and to block harmful content, support staff is freed up to deal with other issues.

A foundation for customized services. The combination of the management server and the Cisco SCE 2020 enables Plala to identify each user, providing the foundation to deliver services tailored for individual customers.

Further Reading

- Cisco SCE 2000 Series Service Control Engine
  cisco.com/packet/182_8b1
- Cisco SCE 2000 Series White Papers
  cisco.com/packet/182_8b2

IMS Migration Primer for MSOs

HOW TO TRANSITION TO NEW MULTIMEDIA AND FIXED MOBILE APPLICATIONS

by jonathan rosenberg

Evidenced by the joint ventures between major cable providers and mobile carriers, multiple system operators (MSOs) are serious about adding mobile services to their cable bundles. Maintaining parity with telcos (with strong mobile offerings) and changes in consumer media consumption patterns have brought attention to the technology alternatives for multimedia services over cable infrastructures. More and more subscribers want their mobile phones to serve as their primary phones and want the ability to access Internet, video, and other multimedia services from their handsets, PCs, or TVs.

Talk of "triple play" services—data, voice, and video—has evolved to "quadruple play" with the addition of wireless services. But it would be more accurate to describe the vision in terms of delivering "triple play on the move." With that vision comes the desire to introduce blended services such as video phone, multimedia chat, and gaming to myriad devices.

To remain competitive, MSOs must provide a variety of new multimedia services in the near term, and in a faster, more efficient way. Examples of upcoming services include "one-number" phone service, unified messages (combining e-mail and voicemail), dual-mode cellular Wi-Fi handsets, video on demand (VoD) streaming to cell phones, and cross-device functions such as viewing caller ID using a set-top box (STB).

MSOs need an infrastructure that supports a wide range of devices. Services must support presence (the ability to know if a person is available for communications) and identity management and personalized services. Furthermore, high-quality voice and video applications require built-in quality of service (QoS) capabilities in the network. Even more challenging, MSOs must ensure that they also continue to provide their existing primary-line voice services without interruption.

As with any technology that spans multiple vendors, network architectures, and media, standards must play a vital role in meeting the challenges faced by MSOs. There is a functionality gap between today's PacketCable 1.x specification and what is needed. In response, CableLabs and a consortium of participating industry players have evolved the Network Call Signaling (NCS) standard and developed PacketCable 2.0.

PacketCable 1.x defines a centralized architecture based on a variant of the Media Gateway Control Protocol (MGCP), called NCS. Call management servers (CMSs or softswitches), media gateways, and media terminal adapter (MTA) endpoints are used to mimic the functionality of the PSTN. This standard supports only MTA endpoints (for connecting legacy phones and fax machines) and does not go beyond POTS over IP at the application and service levels. PacketCable 2.0, by contrast, will include Session Initiation Protocol (SIP) to achieve a distributed multimedia signaling architecture with presence and identity capabilities. A broad range of SIP endpoints are supported within a PacketCable 2.0 environment, including SIP MTAs, SIP phones and soft phones, STBs, video phones, PCs and game consoles, and dual-mode cellular Wi-Fi handsets.

> IMS allows MSOs to take full advantage of their existing IP core network and avoid overloading the cable infrastructure.

SIP, however, is a protocol and not a network architecture. A consortium of players in the mobile market, the 3rd Generation Partnership Project (3GPP), developed the IP Multimedia Subsystem (IMS) specification—a standardized network architecture for SIP-based services. It has been adopted by other standards bodies, major carriers, and equipment vendors alike. While IMS originated in the

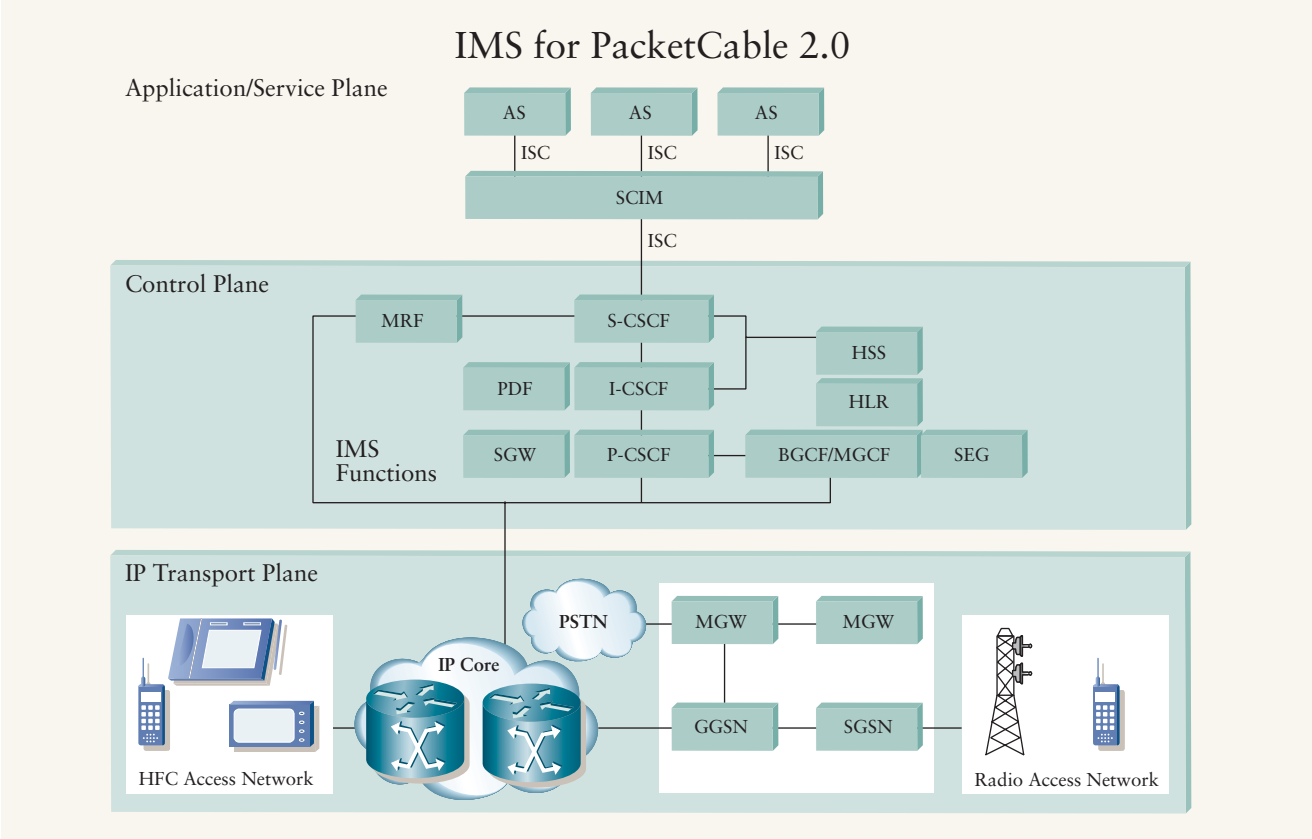

THE IMS PROMISE By introducing a common subscriber data model and standardized interfaces for application servers, IMS has the elements to create a flexible platform for launching new services.

mobile market, it uses SIP to support services across virtually any access technology.

IMS has evolved so that multiple access technologies can be blended: Wi-Fi, WiMAX, DSL, broadband cable access, and even enterprise-level T1. The IMS architecture includes an agnostic control plane that can work over cable or mobile networks, and serves to bridge packet, circuit, wired, and wireless worlds. A layered approach decouples the network infrastructure from services with a standardized, horizontal approach, and enables a flexible platform that makes it easier for providers to respond rapidly to new service requirements and the need to better address service personalization (e.g., self-subscription, buddy lists) and control (e.g., QoS, class of service, charging, security, content filtering).

Because of the momentum of SIP and IMS, it is logical that MSOs and CableLabs would decide to adopt SIP and IMS as a foundation for next-generation networking services as specified in the evolving PacketCable 2.0 standard. PacketCable Multimedia (PCMM) will be the mechanism for implementing QoS for SIP services as part of PacketCable 2.0. Outside of PacketCable 2.0, PCMM can also provide QoS for non-SIP or non-IMS multimedia applications.

Migrating to IMS/SIP Multimedia Services

PROTECTING CURRENT INVESTMENTS and leveraging lessons already learned with PacketCable 1.x requires a gradual, evolutionary path to get to a PacketCable 2.0 infrastructure. Such a path can be realized by using current business needs to define and drive the addition of incremental capabilities to existing networks. The evolution to an IMS architecture will involve incorporating the three core IMS control functions (see figure). The *Serving Call Session Control Function (S-CSCF)*, also referred to as the home proxy or subscriber proxy function, manages access to the subscriber database and uses the information stored in that database to invoke features and applications in response to subscriber requests. The *Interrogating Call Session Control Function (I-CSCF)* controls the boundary to the network and is responsible for routing requests to the right S-CSCF. The *Proxy Call Session Control Function (P-CSCF)* acts as an interface to clients, secures the link to the client, and facilitates roaming.

These IMS control functions sit on top of the IP transport plane. For GSM mobile communications, the IP layer includes the Gateway GPRS Support Node (GGSN) and Serving GPRS Support Node (SGSN) mobile access routers. On the cable

side, this layer includes the HFC access network components. In the application/service layer of the network, an IMS architecture introduces the IMS Service Control (ISC) interface for connecting the S-CSCF to a Service Capability Interaction Manager (SCIM). The SCIM performs feature interaction management, and connects to application servers using the same ISC interface.

The IMS is an important part of a service provider's network, providing a control plane for SIP-based services. However, IMS is only part of the story. A complete IP-based next-generation network is composed of three distinct layers: the *network layer,* which includes the entirety of the IP network; the *service control layer* (of which IMS is a part); and the *applications* that reside on top. The control layer needs to provide a link between the applications and the underlying IP network for both SIP and non-SIP applications. The vast majority of IP-based applications today are not SIP-based, and a layer of service control is needed for these, too. To this end, MSOs require a Service Exchange Framework (SEF) that supports a transparent migration to IMS, with full support for IMS and non-IMS endpoints. With a robust service exchange solution for IMS, cable MSOs will be able to support multiple applications on a common infrastructure that supports both SIP and PacketCable 1.x, while offering subscribers a customized service experience based on real-time state information and profile preferences.

The importance of IMS for future services, and the importance of growing and protecting existing PacketCable 1.x services, calls for a phased introduction of IMS. Consider the following example as one path an MSO might take.

Phase 1: PSTN Bypass

BY MAKING BETTER USE OF MEDIA gateway controllers (MGCs), MSOs can move more of their long-distance voice service off the PSTN and onto the IP network, which reduces the dependencies on competing carriers and reduces costs simultaneously. First, the softswitch must be decomposed into a subscriber-facing CMS and a PSTN-facing MGC. By separating the softswitch into these components, the network can be more easily scaled for better overall network efficiencies.

Once PSTN and subscriber control functions are separated, MSOs can then introduce a combined I-CSCF and Breakout Gateway Control Function (BGCF). (BGCF is the interface for interconnecting IMS with legacy networks.) This IMS element allows PSTN interconnects to be shared by multiple CMSs. CMSs can be added as needed, allowing the network to scale with increases in subscribers. PSTN interconnects can be added as traffic requires. Calls between subscribers can stay on-net, routed to the correct terminating CMS by the I-CSCF function. This configuration offloads calls from the PSTN.

Phase 2: Add SIP-Based Services

PHASE 2 INTRODUCES NEW REVENUE streams for IMS-enabled services, which can include capabilities such as voice dialing, caller ID on a STB, and click to dial. New SIP-based services can now be rapidly introduced and delivered by connecting new application servers to the CMSs. Multiple application servers from multiple vendors can all interconnect to CMSs over the IMS ISC interface, a SIP-based interface for use with application servers. Here, a true blending has occurred. The ISC interface from IMS and PacketCable 2.0 is used, but to provide features to existing PacketCable 1.x endpoints. This allows a minimally disruptive transition with no forklift upgrades and no replacement of endpoints.

Phase 3: Business/Commercial Voice

IN THIS PHASE, AN MSO might focus on the business needs related to expanding the commercial subscriber base. MSOs can build on the IMS environment by adding SIP endpoints (e.g., the Linksys SPA9000, a residential standalone box that provides a small, integrated IT PBX solution with support for up to 16 lines, or SIP MTAs). To interface to these SIP endpoints, the P-CSCF is introduced, which provides PCMM QoS features that ensure business-grade voice services. The P-CSCF connects to the underlying CMTSs and policy servers to provide QoS and security functions. This phase also uses the in-place CMSs, allowing them to control the SIP endpoints. Through the ISC interface, the SIP endpoints also gain access to the new applications deployed from the previous step. With the completion of Phase 3, an MSO has introduced a significant portion of an overall IMS architecture.

Phase 4: Fixed Mobile Convergence

THE MOVE TOWARD FIXED MOBILE convergence involves support for dual-mode handsets and the introduction of two servers. The dual-mode devices can communicate over the cellular network, or act as a new endpoint on the IP network. The Home Subscriber Server (HSS) manages subscriber data uniformly between cellular and IP. The Handoff Server runs on top of the ISC interface and provides a seamless experience when subscribers move from cellular to a Wi-Fi network.

The CMS remains the functional center of the network, but with the introduction of the HSS, has added the Cx and Sh interfaces defined by the IMS, taking it a step further to becoming a complete S-CSCF. By continuing to take advantage of the CMS in each phase, MSOs accomplish a truly evolutionary move to IMS. **P**

DR. JONATHAN ROSENBERG is a Cisco Fellow in the Routing and Service Provider Technology Group at Cisco. He can be reached at jdrosen@cisco.com.

beyond
speeds+feeds

Storage Switch Redefines Scalability

NEW CISCO MDS 9513 PACKS THE INDUSTRY'S HIGHEST PORT DENSITY WITH UP TO 10 GBIT/S by gene knauer

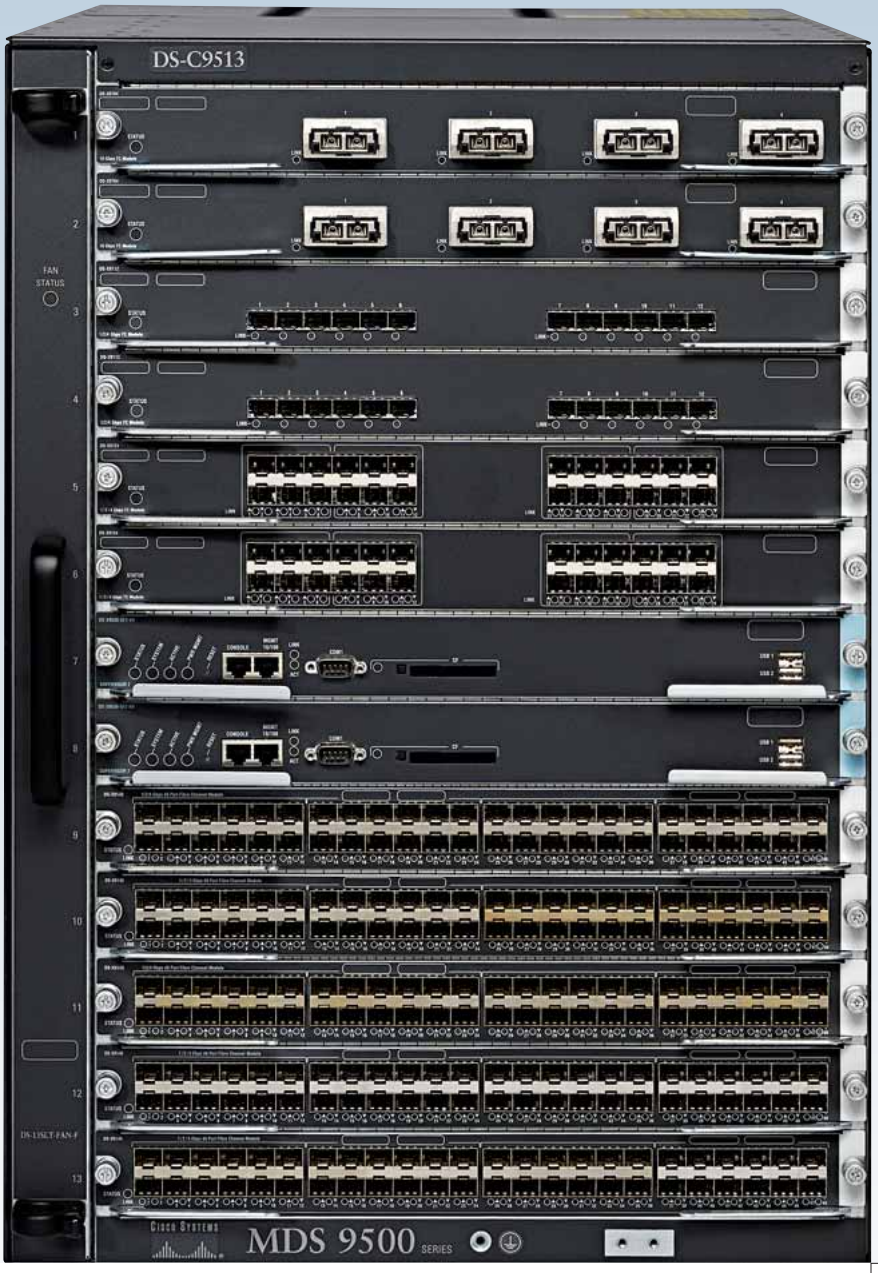

SCALABILITY and consolidation are bywords for the Cisco MDS 9513 Multilayer Director Switch. Built for enterprise and service provider data centers, this mighty platform can scale up to 528 ports with a choice of 1, 2, 4, and 10 Gbit/s of Fibre Channel—the highest speeds for Fibre Channel-over-IP (FCIP), SCSI-over-IP (iSCSI), and Fibre Connection (FICON) all in one system.

"The Cisco MDS 9513 is a great core switch for consolidating multiple SANs [storage-area networks] into one infrastructure," says Alison Conigliaro-Hubbard, senior manager of storage networking at Cisco.

The modules used in the first-generation Cisco MDS 9000 switches—for different types of

Continued on page 76

STORAGE STUD
The MDS 9513 sports 10-Gbit/s speeds on different blades in the same switch, and a choice of storage transport.

newproduct dispatches

Switching

Cisco ME 6500 Series Ethernet Switch

The new Cisco ME 6500 Series Ethernet Switch is optimized to meet the stringent performance, reliability, and quality-of-service (QoS) requirements for delivering voice, video, and data services to consumers as well as providing VPN services for Ethernet-to-the-home (ETTH), Ethernet-to-the-business (ETTB), and DSLAM aggregation deployments. The Cisco ME 6524 Switch extends advanced Multiprotocol Label Switching (MPLS), QoS, multicast, and IPv6 features into Carrier Ethernet access and aggregation networks, for both fiber and copper deployments. The ME 6524 is available in two configurations: one with 24 Gigabit Ethernet Small Form-Factor Pluggable (SFP) downlinks and eight Gigabit Ethernet SFP uplinks, and one with 24 Ethernet 10/100/1000 downlinks and eight Gigabit Ethernet SFP uplinks.

cisco.com/packet/182_npd1

Cisco Catalyst Blade Switch 3030 for Dell

Designed for the Dell PowerEdge Blade Server Enclosure, the Cisco Catalyst Blade Switch 3030 for Dell helps ensure consistency across a Cisco network for services such as security and resiliency. Security features include extended access control lists (ACLs), QoS, and Network Admission Control (NAC). Resiliency features for high availability allow uninterrupted network performance. A graphical device manager simplifies switch configuration and supports CiscoWorks management software. The switch provides two ports for 10/100/1000BASE-T Ethernet connections and four SFP-based ports for 100BASE-T and 1000BASE-SX fiber connectors. Up to four Cisco Catalyst Blade Switch 3030 units can be installed in a single blade server chassis.

cisco.com/go/cbs3030

CISCO ME 6500 SERIES ETHERNET SWITCH

SPOTLIGHT ON

Cisco Security IntelliShield Alert Manager Service

The Cisco Security IntelliShield Alert Manager Service is a web-based security threat and vulnerability alerting service that notifies individuals within an organization, in advance, about emerging threats that potentially impact the specific technologies for which they are responsible. This advanced notification allows a proactive response and helps to avoid the impact of the most crippling new and emerging security threats. Cisco Security Intellishield Alert Manager combines a comprehensive database of thousands of threats and vulnerabilities with client-customized "smart filters" and notifications as well as a workflow management system. The service encompasses four primary components:

A **portal** provides a secure and customizable interface for receiving information on the specific networks, systems, and applications used by the organization. The portal can push notifications to other devices, applications, and vulnerability management systems.

The **back-end intelligence engine** collects threat data and filters it through a rigorous verification, editing, and publishing process. Cisco Security IntelliShield Alert Manager experts review and analyze each threat to confirm the characteristics and produce information in a standardized, easy-to-understand format. Each threat is objectively rated on urgency, credibility of source, and severity of exploit, helping users make response decisions more quickly.

The **historical database** provides an extensive collection of past threat and vulnerability data to help users identify threats against specific products or track alerts by type.

A **built-in workflow system** helps users track their remediation efforts for identified vulnerabilities, and allows IT managers to see the assignment and current status of all remediation efforts.

cisco.com/go/intellishield

Security and VPNs

Cisco Security Manager

The Cisco Security Manager software allows administrators to centrally provision device configurations and security policies for firewalls, VPNs, and intrusion prevention systems. The software provides distinct views at the device, policy, and topology levels to meet different operational preferences and needs. Other features include policy sharing for efficiency and policy distribution for greater network resilience. Cisco Security Manager is effective for managing small networks with fewer than ten devices, but also scales to efficiently manage large-scale networks with thousands of devices. Scalability is achieved through intelligent, policy-based management techniques that can simplify administration. For more on Cisco Security Manager, see page 48.

cisco.com/go/csmanager

Cisco ASA 5500 Series Adaptive Security Appliance: New Services Module and Software Version

The Cisco ASA 5500 Series Adaptive Security Appliance is a modular platform that provides security and VPN services in enterprise networks. The new Content Security and Control Security Services Module (CSC-SSM) delivers comprehensive threat protection and content control services. These services include antivirus, anti-spyware, file blocking, anti-spam, anti-phishing, URL blocking and filtering, and content filtering. Version 7.1 of the Cisco ASA 5500 Series software delivers new features for Secure Sockets Layer (SSL) VPNs including customizable user portals; Web content transformation for compatibility with Java, ActiveX, complex HTML, and JavaScript; and scalability up to 5,000 concurrent users. For more on the ASA 5500 Series module and version 7.1 software, see page 48.

cisco.com/go/asa

Voice and Video

Cisco Unified Personal Communicator

The Cisco Unified Personal Communicator software transparently integrates many communications applications and services into a single desktop computer application. It provides quick and easy access to tools for voice and video calls, Web conferencing, call management, directories, voice messages, and presence information. These tools help users communicate effectively and work more productively. See related article on page 36.

cisco.com/packet/182_npd2

Cisco AS5400XM and AS5350XM Universal Gateways: New High-Density Packet Voice/Fax Feature Card

The new High-Density Packet Voice/Fax Feature card for Cisco AS5400XM and Cisco AS5350XM Universal Gateways provides scalability from 64 to 384 channels. The Cisco AS5400XM now offers

CISCO AS5400XM HIGH-DENSITY PACKET VOICE/FAX FEATURE CARD

increased voice call capacity up to a full channelized T3 (CT3) with all codecs, and the Cisco AS5350XM offers 12-E1, 16-T1, and CT3 configurations using the G.711 codec. The High-Density Packet Voice/Fax Feature Card introduces support for the Adaptive Multi-Rate Narrow Band (AMR-NB) codec, media authentication and encryption using Secure Real-Time Transport Protocol (SRTP), and modem relay.

cisco.com/packet/182_npd3
cisco.com/packet/182_npd4

Cisco Unified Presence Server

The Cisco Unified Presence Server software collects information about a person's current availability and communications capabilities to help users connect with colleagues more efficiently. Cisco Unified Presence Server provides the presence information used by Cisco Unified Personal Communicator and Cisco IP Phone Messenger. It also supports Microsoft Office Communicator functions for click-to-dial and phone monitoring. See related article on page 36.

cisco.com/packet/182_npd5

Cisco Unified Customer Interaction Analyzer

The Cisco Unified Customer Interaction Analyzer is a managed service that helps contact centers analyze each interaction that occurs between customers and service representatives. Customer interactions are captured on site and transferred to a dedicated, secure hosted facility for storage, retrieval, and evaluation. Through a Web-based interface, evaluation results are presented on tailored reports or dashboards that help contact center managers with employee coaching and call scripts. See related article on page 36.

cisco.com/packet/182_npd6

Cisco Unified IP Phones 7900 Series: New Video Phone Model and Enhanced SIP Feature and Gigabit Ethernet Support

The Cisco Unified IP Phone 7985G is a personal desktop video phone that enables instant, face-to-face communications. It incorporates the

newproduct dispatches

camera, LCD screen, speaker, keypad, and handset into a single, easy-to-use unit. The new Cisco Unified IP Phone 7971G-GE (color touch-screen display), 7961G-GE (grayscale with six programmable buttons), and 7941G-GE (grayscale with two programmable buttons) include a Gigabit Ethernet port for integration to a PC or desktop server. When used with Cisco Unified CallManager 5.0, the Cisco Unified IP Phones 7971G-GE, 7970G, 7961G, 7941G, and 7911G now support enhanced Session Initiation Protocol (SIP) services. See related article on page 36.

cisco.com/packet/182_npd7

Cisco Unified Wireless IP Phone 7920 Multi-Charger

The Cisco Unified Wireless IP Phone 7920 Multi-Charger can charge up to six Cisco Unified Wireless IP Phone 7920s and six batteries simultaneously. The multi-charger is ideal for environments where employees work in multiple shifts and need to keep their phones charged throughout the day. The unit has a flexible design for placing on a desk or mounting on the wall.

cisco.com/packet/182_npd8

Cisco 7800 Series Media Convergence Server: Call-Processing Appliance on New Models

New models of Cisco 7800 Series Media Convergence Servers (MCS) provide a call-processing appliance with preinstalled Cisco Unified CallManager 5.0 software. The server appliance is fully operational upon startup, requiring entry of only a few parameters such as IP address and domain. The new models offer choices for single or dual processors and can serve as many as 7,500 IP phones per appliance and up to 30,000 IP phones per cluster.

cisco.com/packet/182_npd9

Network Management

CiscoWorks Interface Configuration Manager

The CiscoWorks Interface Configuration Manager software provides graphical configuration and reporting capabilities to help manage access ports on Cisco switches, reducing the time and effort needed to roll out consistent bulk configurations across the network. The software provides tools to simplify configuration of Layer 2 Network Admission Control (L2 NAC) in conjunction with the CiscoWorks LAN management solution, reducing operator errors and increasing network availability. It provides a L2 NAC Readiness Report to help optimize network research and planning. NAC technologies such as L2 802.1x, L2 IP, AAA fail open, MAC exceptions, PBACLS, and RADIUS server configuration are supported. The configurations can be saved, copied, and edited, freeing up time for network administrators. Customized device and port groups can be created and re-used in multiple configurations.

cisco.com/go/cwicm

CISCO UNIFIED WIRELESS IP PHONE 7920 MULTI-CHARGER

Cisco Router and Security Device Manager Version 2.3

The Cisco Router and Security Device Manager (SDM) is a Web-based device-management tool for Cisco IOS Software-based routers. SDM version 2.3 delivers comprehensive support for simplified management and real-time performance monitoring of concurrent services on Cisco Integrated Services Routers. Key features include integrated IOS WebVPN management for secure remote-access connectivity, NetFlow and Network-Based Application Recognition (NBAR) traffic statistics for performance monitoring, real-time threat alerts, easier and more intelligent signature updates for intrusion prevention systems (IPS), VPN design guide for intelligent selection of appropriate VPN technology, and router IOS image management.

cisco.com/go/sdm

Cisco Unified Operations Manager

The Cisco Unified Operations Manager software provides a real-time, service-level status view for each element in the Cisco Unified Communications solution, including the underlying network infrastructure. The application remotely polls and collects data from the monitored devices, and also provides diagnostic capabilities for faster trouble isolation and resolution. See related article on page 36.

cisco.com/packet/182_npd10

Cisco Unified Service Monitor

The Cisco Unified Service Monitor software provides a method to monitor and evaluate voice quality in Cisco Unified Communications solutions. It delivers near-real-time notification when the voice quality of a call

does not meet a user-defined quality threshold. The Cisco Unified Service Monitor analyzes voice data collected by the associated Cisco 1040 Sensor, which is installed in campus and remote locations to monitor call quality. See related article on page 36.

cisco.com/packet/182_npd11

Networked Home

Linksys Wireless-G Broadband Router and Wireless-G PC Card with RangeBooster

The Linksys Wireless-G Broadband Router with RangeBooster (WRT54GR) provides an Internet-sharing router, a four-port LAN switch, and a Wireless-G access point that supports greater range, throughput, and coverage area. The associated Linksys Wireless-G PC Card (WPC54GR) provides high-speed Wireless-G networking for a notebook computer with RangeBooster technology. The card also interoperates with standard Wireless-G and Wireless-B networks, and for security supports Wi-Fi Protected Access and up to 128-bit encryption.

cisco.com/packet/182_npd12
cisco.com/packet/182_npd13

Linksys Wireless-G Media Storage Link Router with SpeedBooster

The Linksys Wireless-G Media Storage Link Router with SpeedBooster (WRTSL54GS) combines an Internet-sharing router, Wireless-G access point, and network storage link. The router enables users to connect USB disk drives directly to create a secure home network with networked storage. A built-in media server streams music, video, and photos located on the attached storage device to any Universal Plug and Play (UPnP)-compatible media adapter. Storage devices attached to the WRTSL54GS can also be set up for public download via a Web browser or password-protected user accounts.

cisco.com/packet/182_npd14

Linksys Compact Wireless-G Internet Video Camera

The Linksys Compact Wireless-G Internet Video Camera (WVC54GC) wirelessly connects to a home or small office network and distributes a live video stream via the Internet for video monitoring applications. The camera can be mounted on a wall or placed in its desktop stand. When the security mode is activated, the WVC54GC will send e-mail alerts with video clips to three e-mail addresses upon detecting motion within the camera's field of view.

cisco.com/packet/182_npd15

Linksys Wireless-G Music Bridge

The Linksys Wireless-G Music Bridge (WMB54G) allows users to wirelessly stream Yahoo! Music services and other music stored on a PC to a home entertainment center. The Wireless-G Music Bridge connects to a home stereo using standard consumer electronic cables. It also connects to a home network via Wireless-G networking or standard 10/100 Ethernet cabling.

cisco.com/packet/182_npd16

Cisco IOS Software

WebVPN Software

Cisco IOS WebVPN is the industry's first router-based solution offering Secure Sockets Layer (SSL) VPN remote-access connectivity integrated with security and routing features on a converged data, voice, and wireless platform. Using SSL VPN, companies can securely and transparently extend their networks to any Internet-enabled location. Cisco IOS WebVPN supports clientless access to applications such as HTML-based intranet content, e-mail, network file shares, and Citrix; and the Cisco SSL VPN Client, enabling full network access remotely to virtually any application. Cisco Secure Desktop, as part of WebVPN, offers data theft prevention even on non-corporate devices. Cisco Router and Security Device Manager (SDM) eases WebVPN deployment and performs real-time monitoring and management of SSL VPN sessions.

cisco.com/packet/182_npd17

Cisco RSVP Agent

The Cisco RSVP Agent feature uses the network to deliver Call Admission Control (CAC) and quality of service (QoS) with Cisco Unified CallManager deployments. When Cisco RSVP Agent is active, users experience superior call QoS and reliability. Resource Reservation Protocol (RSVP) is based on IETF standards to secure and reserve bandwidth for calls across the WAN. With CAC, the network can accept or reject a call based on bandwidth and policy considerations. Cisco RSVP Agent is supported on the Cisco 2600XM, 2691, 2800, 3700, and 3800 Series Integrated Services Routers.

cisco.com/packet/182_npd18

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 February and April 2006. For real-time announcements of the most recently released products, see "News Archive, News Releases by Date" at **cisco.com/dlls/index.shtml.**

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

Cisco MDS 9513, Continued from page 71

transport, speed, and storage management applications—are fully compatible with the MDS 9513. And the new 4- and 10-Gbit/s modules for the MDS 9513 are backwards compatible with the previous model switches.

"In the past, replacing director switches required a major forklift upgrade, but that's not necessary with the MDS architecture," says Dan Hersey, technical marketing engineer in Cisco's Data Center, Switching, and Security Technology Group. "The modules are forwards and backwards compatible. No other vendor can say that."

In the first-generation MDS switches, Cisco designed the 32-port card of its MDS 9509 to take advantage of oversubscription, which occurs when the overall bandwidth available to a switch is less than the aggregate bandwidth of all ingress switch ports. This design gives users the flexibility of having multiple ports vying for a limited amount of bandwidth for better overall utilization of each port.

Although 4 and 10 Gbit/s are available on the MDS 9513, most servers, storage subsystems, and applications run at 1 to 2 Gbit/s, so bandwidth must be carefully allocated to accommodate actual I/O throughput at each interface. The earlier 32-port modules in the MDS 9000 Series provided the ability to manage bursts of speed with a solution that allowed any device to burst at high data rates and still have performance available for the other ports. A new feature in the Cisco MDS 9513—Port Bandwidth Reservation—allows administrators to selectively assign line rates to particular interfaces, as needed, so oversubscription can be more granularly managed.

Developed to overcome the inefficiencies of isolated SAN islands, Cisco virtual SAN (VSAN) technology maps SAN islands onto a common physical infrastructure without merging control and management of the independent fabrics. Each port in the Cisco MDS 9513 is assigned to a different set of VSANs, and the VSAN management is centrally controlled. Individual VSAN management of services and policies can be fully or partially assigned to application administrators using role-based access control (RBAC) based on their respective applications.

"With the MDS 9513, you can consolidate up to hundreds of virtual SANs in one core director switch," says Conigliaro-Hubbard. "You can manage and provision them, perform diagnostics, and provide comprehensive security using the same storage management solution."

Find out more at **cisco.com/packet/182_9a1.** P

productreview

Cisco Aironet 1300 Series Wireless Bridge

THE FOLLOWING product review is excerpted from the Networking Professionals Connection Website and was submitted by Sankar Nair, a network engineer at General Datatech, Texas. For the full product review, visit **cisco.com/packet/182_9c1**.

DESIGNED FOR EASE OF INSTALLATION AND MAINTENANCE

WHY DID YOU choose the Cisco Aironet 1300?

My day-to-day work revolves around Cisco IP telephony and wireless. In the last year, I have completed a handful of outdoor wireless installations, especially with the Cisco Aironet 1300. What I like most about this product is the ease with which it can be deployed. Of course, it does take a fair amount experience with similar devices to properly install the equipment. The bridge enclosure meets the NEMA 4 standard, which helps the product survive the most extreme outdoor temperatures.

WAS THE PRODUCT easy to install?

A site survey is definitely a requirement before the installation. The site survey determines if there is clear line of sight between the locations, how much coax cable is needed, what kind of antenna/bridge combination is required for the solution, etc. Cisco provides an excellent tool called the Outdoor Bridge Range Calc Utility, which assists in finding the right bridge/antenna/coax cable combination to achieve the maximum range/throughput (theoretical). An average point-to-point installation takes about four to eight hours to physically mount the hardware and cabling (excluding configuration, testing, etc.).

WHAT LEVEL OF experience is needed to install the product?

Installation requires some level of prior experience with similar products. My past experience with the Aironet 350 Series and other vendor bridges has helped me to a great extent. Integrating these bridges into a wired network also involves configuration of VLANs, trunking, routing, etc., on the switch or router used in the network. Familiarity with configuration of Cisco routers/switches is a plus.

HOW HAVE YOU deployed the Aironet 1300 Series?

Although the 350 and 1300 series are installed more or less in a similar fashion, my first installation of the Aironet 1300 was very challenging. The client for whom I was doing the installation had four locations separated by about a half mile to one mile. Each location had two to four PCs running a ticketing system that talks to a centralized server located at corporate headquarters. The client was using 56k dialup access and wanted to move toward a T1 at one location with all the other locations sharing this T1. Running fiber between all the locations was out of the question due to the harsh outdoor environment where the networks were installed. Each location had a small scale house where the computers were installed, along with a switch or hub. To get clear line of sight between these locations, we had to install the antennas on top of cement silos, electric poles, etc., and run coax cable from the antenna to the bridge, which was installed on the inside of each building. Also one of the locations did not have clear line of sight with the hub location, so we had to use an intermediate location as a repeater site, which eventually did work out very well. The entire installation was completed in two days. The client was extremely satisfied with the ease of installation and the capabilities of the product. We did ten more installs last year spanning from point-to-point to point-to-multipoint installations throughout Texas. P

To submit a product review, visit **cisco.com/go/product_review**.

netproexpert

Configuring and Troubleshooting Cisco ASA 5500 and PIX 7.0

The Networking Professionals Connection is an online community for Cisco experts and networking colleagues. Following are excerpts from a recent Ask the Expert forum, "Configuring and Troubleshooting ASA 5500 and PIX 7.0," moderated by Cisco's Mynul Hoda. To view the full discussion, visit **cisco.com/packet/182_10a1**. To join other live online discussions, visit **cisco.com/discuss/networking**.

GOT A QUESTION?

Expert Mynul Hoda will answer your questions about configuring and troubleshooting Cisco ASA 5500 and PIX 7.0 in a live discussion forum June 5 through June 30, 2006. Join your networking peers!

cisco.com/go/askeexpert/packet

WHY IS IT NECESSARY to assign the PIX a management IP address in Transparent mode, without which it doesn't allow inside hosts to ping outside hosts? What is the significance of the IP address to the PIX in Transparent mode?

To understand the need for assigning IP addresses from the same subnet as the inside and outside interface, it is important to understand how the ASA works when configured for Transparent mode. In this mode, the ASA operates at Layer 2; therefore, the routing table is not examined to find out the outgoing interface for the traffic going through the ASA. Just like the switch, ASA needs to build the MAC table (CAM table) to forward the packet to the appropriate interface. However, the way the MAC is discovered on the ASA is very different than on the switch. When the ASA receives a packet and if the destination MAC is not in the Layer 2 forwarding table, then ASA doesn't flood the packet out both interfaces like the switch does, as that would be a security risk. Rather, it will send either an Address Resolution Protocol (ARP) or Internet Control Message Protocol (ICMP) message out both interfaces, trying to determine which interface the destination MAC address resides. ARP is used when destination host is on the same local segment, but ICMP is used if not in the same segment. In both cases, a source address is used and that address is the management interface IP address. If you do not assign this IP address, then ASA will not learn the MAC address and will not be able to pass any traffic.

DOES THE ASA forward the broadcast to the outside? How does the outside interface learn the MAC address and the IP address of the inside host?

No, it doesn't, except the ARP broadcast between the inside and outside host. When the outside host sends the ARP broadcast, ASA will receive it and build up its MAC table. The ASA will make the ARP request on behalf of the outside host based on the destination IP address. Once the host inside replies back, the ASA would know about the inside host MAC address. If the destination IP is not in the same subnet, the discovery of the MAC will take place using the ping packets. To make the ARP request or when sending ping packets, the ASA is required to have the IP address, as both types of packet will use management interface IP as the source, which is the reason you need this IP. Refer to **cisco.com/packet/182_10a2** for packet flow through Transparent FW.

WHEN DOING A cable-based failover and using a dedicated interface for stateful failover between two PIX firewalls, do the firewalls also monitor the stateful interface for failover purposes along with the other interfaces? Can I use a crossover to connect the stateful failover interfaces together?

Yes, the PIX boxes monitor the stateful failover interface too. And yes, you can use a crossover cable for the stateful link. The PIX documentation recommends this method. P

MYNUL HODA is a lead engineer in the High Touch Technical Support group at Cisco.

Streaming Media, Continued from page 27

environment is simplified hybrid routing. This option also uses coverage zones, but clients send requests directly to the CDM or content router. The CDM or content router then redirects client requests to the most appropriate WAE using HTTP redirection, based on coverage zones that you configure.

Splitting Live and Rebroadcast Streams

WAEs also offer stream-splitting services to clients requesting live and scheduled rebroadcast streams. With stream splitting, a WAE receives an incoming stream and fans it out to requesting clients (Figure 2, page 27). WAEs treat rebroadcast streams that you schedule on your streaming servers no differently than live streams.

With live splitting, the live stream always originates from the origin server as a single stream to the incoming interface of the WAE. Both front- and back-end streams can be either multicast or unicast.

With unicast-to-unicast delivery, the WAE receives the live or rebroadcast stream as unicast and splits it to clients as individual unicast streams. Each client receives a distinct stream causing potential duplicates on the network, depending on client location. For example, if two requesting clients are on the same subnet, two duplicate streams will flow to that subnet over the network.

Alternatively, in a unicast-to-multicast configuration, the WAE generates a multicast stream from the unicast input and sends it to a multicast group that you can configure on the WAE. Provided that you enable multicast on the downstream network, the routers will forward the multicast group traffic to its members without duplication. This delivery mechanism is known as *pull-splitting*, because client requests trigger the WAE to proactively "pull" the unicast stream from the origin server.

You can also configure multicast-to-multicast splitting. This is beneficial because you can conduct services in the WAE, such as content authentication, URL filtering, and Internet Content Adaptation Protocol (server offload) services. Finally, multicast-to-unicast splitting lets you distribute your streams to clients in areas of your network that do not support multicast.

These last two methods are known as *push-splitting* because the origin server blindly "pushes" the stream to the network for the WAE to receive by joining the multicast group.

When deploying streaming media, using suite of performance-acceleration tools helps the quality of audio-video broadcasts remain consistent and reliable. The Cisco ANS architecture includes Cisco ACNS software for automating content replication to WAEs in multiple sites where the content is locally cached and client requests are served. Further conserving bandwidth, live stream splitting allows the origin server to avoid replicating the stream for each receiving client. P

Further Reading

- Cisco Application Networking Services/ Application Acceleration
  cisco.com/packet/182_5a1
- Content Networking Fundamentals, by Silvano Da Ros (Cisco Press, ISBN: 1587052407)
  cisco.com/packet/182_5a2

PACKET ADVERTISER INDEX

cachefile

Cyber Quote "To succeed as a team is to hold all of the members accountable for their expertise." • Mitchell Caplan, CEO, E*Trade

Using Fingerprints to Secure Networks

University of Buffalo, New York, research findings could help eliminate the need to remember a dizzying array of passwords, or aid forensics specialists, according to Venu Govindaraju, professor of computer science and engineering, and director of the school's Center for Unified Biometrics and Sensors. Govindaraju and his team have determined the minimum surface area required for fingerprint scanning to achieve a level of security that is roughly comparable to the security achieved with a six-letter password. The algorithm takes into account the fact that even a legitimate fingerprint doesn't always look the same due to the way a person presses on a pad or because of moisture or other factors. [networkworld.com]

Net Lingo

Identity chaos—When a user has multiple identities and passwords across a variety of networks, applications, computers and/or computing devices. [whatis.com]

NEXT-GENERATION VOIP GEAR HITS NEW HIGH IN 2005

The next-generation voice-over-IP equipment market hit a new high in 2005 with US$2.5 billion in revenue, according to Infonetics Research. This was double the next-generation voice equipment revenue from the prior year. Sales are projected to increase by 145 percent over the next four years to reach US$6.2 billion by 2009. [networkingpipeline.com]

PERSONAL RELATIONSHIPS EXPAND ON THE WEB

The Internet allows people to build social networks that support personal decisions, according to findings in a report published by Pew Internet & American Life Project. Refuting previous assumptions that the Web limits social interaction, the report states that technologies such as the Internet, e-mail, instant messaging, and cell phones allow people to develop both "core ties" (social contacts involving a very close relationship with a person) and "significant ties" (people somewhat closely connected to an individual) at a more global level. Study respondents reported an average 23 core ties and 27 significant ties. The report is based on the findings of two daily tracking surveys on Americans' use of the Internet, and involved interviews with 2,200 adults age 18 and older. [**clickz.com**]

the 5th wave

"Think of our relationship as a version of Red Hat Linux—I will not share a directory on the love-branch of your life."

©The 5th Wave
www.the5thwave.com

POP QUIZ

Level: CCNA

IP Routes

Answers

QUESTIONS ON PAGE 16.

1. Enter the ip route destination-network [mask] {next-hop-address | outbound-interface} [distance] [permanent] global command. Example: RouterB (config) #ip route 172.17.0.0 255.255.0.0 172.16.0.1
2. The rule of split horizon is that it is never useful to send information about a route back in the direction from which the original update came.
3. AD is an integer from 0 to 255 that rates the trustworthiness of routing information received from a neighboring router. The AD is used as the tie-breaker when a router has multiple paths from different routing protocols to the same destination. The lower the path's AD, the more likely it is to be used.
4. By periodically broadcasting the entire routing table out all active interfaces. This method is often called "routing by rumor."
5. Convergence is when all routers have consistent knowledge and correct routing tables.

Source: CCNA Flash Cards and Exam Practice Pack