

PACKET

CISCO SYSTEMS USERS MAGAZINE

THIRD QUARTER 2004

ROUTING INNOVATION

Rising Expectations
in IP Networking 34

Cisco CRS-1:
Reinventing the Router 41

Deploying Video Telephony 23

Detecting Network Threats 13

SPECIAL REPORT:
Intelligent Networking 53



CISCO.COM/PACKET

PACKET

CISCO SYSTEMS USERS MAGAZINE

THIRD QUARTER 2004
VOLUME 16, NO. 3



34

ON THE COVER

Turning the Corner on Innovation

34

Market demands and sophisticated new applications are accelerating architectural innovation in IP routing. Cisco turns the corner with the new CRS-1 Carrier Routing System and enhancements to Cisco IOS® Software.

Reinventing the Router

41

With unparalleled capacity and raw horsepower, the Cisco CRS-1 provides the fault-tolerant, multiple-service networking service providers require to sustain anticipated growth in IP services over the next decade.

IOS: Routing's Crown Jewel

47

From its public debut in 1987 to the recent delivery of Cisco IOS XR for fault-tolerant routing at 92 Terabit-per-second speeds, Cisco IOS Software continues to evolve with the times.



53

SPECIAL REPORT

Intelligent Networking

53

An intelligent, systems-based approach to networking can substantially reduce complexity while increasing functionality. Learn more about Cisco's vision of the smarter network.



57



71



81

IN EVERY ISSUE

Mail	3
Calendar	5
Acquisitions	7
Networkers	6
Tech Tips	21
Advertiser Index	88
Cache File	90
The 5th Wave	90

TECHNOLOGY

VIDEO TELEPHONY: Deploying Video Telephony	23
---	-----------

Cisco CallManager 4.0 extends voice features to video over a common, user-friendly infrastructure that can be deployed to the desktop.

SECURITY: Deflector Shield	28
-----------------------------------	-----------

Fruits of Cisco Riverhead Networks acquisition help to mitigate distributed denial-of-service attacks.

ENTERPRISE SOLUTIONS

Turbo-Charged TAC	57
--------------------------	-----------

A virtual customer interaction network for Mercedes-Benz USA accelerates auto diagnosis and puts the brakes on telephony costs.

Routed Radio	61
---------------------	-----------

New Cisco Catalyst® 6500 Series Wireless LAN Services Module blends wired and wireless networks.

Radio Meets Multicast	63
------------------------------	-----------

Radio broadcaster GWR Group lowers costs by replacing satellite, data, and voice networks with multicast VPN.

Virtual Firewall Management	67
------------------------------------	-----------

Network administrators can manage multiple security contexts using Cisco PIX® Device Manager Version 4.0.

SERVICE PROVIDER SOLUTIONS

Wholesale BLISS	71
------------------------	-----------

Z-Tel Communications taps Cisco BLISS solution for unique wholesaler/retailer opportunity.

Taking to the ROADM	75
----------------------------	-----------

Reconfigurable optical add/drop multiplexer (ROADM) technology poised to spur metro dense wavelength-division market.

Calculating New Routes Faster	78
--------------------------------------	-----------

Cisco IOS® Software enhancements speed IS-IS network convergence.

SMALL AND MIDSIZED BUSINESSES

IP VPNs Gain Momentum	81
------------------------------	-----------

Small and midsized companies can save time and money by out-tasking their IP VPNs to a managed services provider.

DEPARTMENTS

From the Editor	1	Technically Speaking	84
Innovation and Standardization		IP Security or Secure Sockets Layer?	
User Connection	5	Cisco's Pete Davis discusses why you don't have to choose one over the other.	
CIPTUG IP Telephony Feature Request System • Cisco Career Certifications Updates		New Product Dispatches	85
Tech Tips & Training	9	What's new from Cisco over the past quarter	
Is Your Network Ready for Voice? • Threat Detection • Insider's Tips on Earning Your CCIE in Security • IP Multicast at a Glance • Reader Tips		NetPro Expert	89
		Expert advice on outdoor wireless LAN infrastructure	

PACKET MAGAZINE

David Ball
Editor-in-Chief

Jere King
Publisher

Jennifer Redovian
Managing Editor

Susan Borton
Senior Editor

Joanie Wexler
Contributing Editor

Robert J. Smith
Sunset Custom Publishing
Production Manager

Michelle Gervais, Nicole Mazzei,
Mark Ryan, Norma Tennis
Sunset Custom Publishing
Production

Jeff Brand, Bob Jones
Art Direction and Packet Redesign

Emily Burch
Designer

Ellen Sokoloff
Diagram Illustrator

Bill Littell
Print Production Manager

Cecelia Glover Taylor
Circulation Director

Valerie Marliac
Promotions Manager

Scott Griggs, Jordan Reeder
Cover Photograph

Special Thanks to the Following Contributors:

Leonard Bonsall, Jeff Brand, Karen Dalal,
Bob Jones, Janice King, Valerie Marliac,
Sam Masud

Advertising Information:

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

View Packet magazine at cisco.com/packet.

Publisher Information:

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

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

Packet Magazine

PO Box 2080
Skokie, Illinois 60076-9324

USA

Phone: 847-647-2293

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

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

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

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

This magazine is printed on recycled paper.



10%

TOTAL RECOVERED FIBER

CISCO SYSTEMS

Reprinted with permission from *Packet*® magazine (Volume 16, No. 3), copyright © 2004 by Cisco Systems, Inc. All rights reserved.

FROM THE EDITOR

Innovation and Standardization

If you're a regular reader of *Packet*®, you've no doubt noticed our new look. *Packet* has been redesigned to match a new look and feel that has been incorporated throughout all of Cisco's communications vehicles. From the commercials you see on TV, to the boxes that deliver your latest networking components, the company is adhering to a cohesive design philosophy that is collectively referred to in marketing circles as a *corporate identity system*. The theory is, if you're spending money on individual communications, each with its own audience, objectives, and agenda, you also want them to work together for a higher purpose—in this case, to build brand awareness in the marketplace. A corporate identity system makes individual components (whether a white paper, data sheet, or a magazine) work together for a greater good.

As I sat down to write this letter, I thought, *how can I tie Packet's redesign into this issue's theme of routing innovation?* Then it occurred to me: what we are experiencing at *Packet* is the same inevitable evolution that occurs in the world of networking—innovation to standardization—the standardization of the most practical and useful innovations to serve a greater good, that of widespread adoption and integration.

To advance the state of the art in any given field, there must be innovation. Throughout its 20-year history, Cisco has pioneered many innovations that continue to profoundly affect not only networking, but, to quote Cisco Chief Executive Officer John Chambers, the very way the world "works, lives, plays, and learns." However, as important as innovation is, working with the standards bodies ensures that the advancements achieved can be used by everybody. Few companies have invested as much effort in standards development as Cisco. A few examples of the company's contributions to industry standards include Border Gateway Protocol (BGP), Dynamic Packet Transport/Resilient Packet Ring (DPT/RPR), Multiprotocol Label Switching (MPLS), and Layer 2 Tunneling Protocol (L2TP). For more Cisco innovations, see "Turning the Corner on Innovation," page 34.

Companies reap huge benefits from standards-based networking technologies. While it might seem that conformance to industry standards would stifle creativity, the opposite is true. When all products and technologies adhere to industry standards, vendors must differentiate their products by other means. This competition between network equipment suppliers brings out the best in each vendor and continually pushes technology forward.

Over the years, *Packet* has won its share of awards for innovative design, photography, and illustrations. So, while we may have a smaller design palette with which to stretch our creative muscle, we will continue to work hard to differentiate ourselves with innovative editorial. To that end, a new column, "NetPro Expert" (see page 89), has been added to help satiate your appetite for technical tips and advice. Each quarter, this column will provide excerpts from a particularly interesting Q&A session held with one of Cisco's technical experts on the popular *Cisco Networking Professionals Connection* online community (cisco.com/go/netpro).

Look for more integration with NetPro forums on our newly designed Packet Online Website, coming soon. And let us know what you think of our new look by writing to us at packet-editor@cisco.com.

David A. Ball

David Ball
Editor-in-Chief
daball@cisco.com



THIRD QUARTER 2004 PACKET 1

MAIL

A Question of Timing

In reference to Yang Difei's Reader Tip [Second Quarter 2004], I'm surprised that an editor's note wasn't included. I like the functionality of the **reload** command and use it frequently when performing remote administration, but **reload in 60** gives you one heck of a waiting period for the router to revert to its prior configuration. I prefer to make changes to my equipment in small increments and use an appropriate reload in time of between 2 and 5 minutes. If you misconfigure a WAN interface and lose your connection, you've probably also lost the connectivity for several users.

—Gerri Costa, Promasa, New Orleans, Louisiana, USA



Diary Inspires Interest

After reading the second installment of Jimmy Kyriannis's "Deployment Diary" [First Quarter 2004], I went back and read the first part of the series [Second Quarter 2003]. On page 47, Kyriannis says he tested the new core while a "leaf" off the current production network with 2 million independent connections. He also stated that later they would test with 5 million connections. How can anyone possibly test this many connections? I think it's questionable that anywhere close to 2 million connections or "flows" would exist at any one time on a large campus network given the brief, transitory nature of many types of connections between routers.

—Mike Granger, EDS Corp., Louisville, Colorado, USA

The following is a response from author Jimmy Kyriannis.—Editors

The manner in which I conducted the test is fairly straightforward. To validate the Cisco Express Forwarding-based load-sharing algorithm, I didn't actually have to establish a complete connection with any end systems, but I did need to show that the traffic successfully traversed the Tetrahedron Core as described in the load-sharing algorithm documentation. Here's a brief outline of my test method.

Correction

The article "Branching Out" [Second Quarter 2004, page 80] contained factual errors regarding First Albany Capital's network deployment. A corrected version of the article is available at cisco.com/packet/163_2a1. We apologize for the errors.—Editors

1. I placed a UNIX system on a network that was attached to an access router connected to the Tetrahedron Core. That network was a /24 subnet, meaning that it could support a maximum 256 IP addresses.

2. I configured the UNIX system to use 250 IP addresses on its single Gigabit Ethernet interface.

3. I wrote an execution script to do the following:

- Randomly select a source IP address from one of the above 250 (in some of the tests, I used just a single source IP address)
- Randomly select any global destination IP addresses, up to a total of 5 million
- Execute a traceroute from that selected source IP address to that destination IP address using a max ttl that would ensure that the traffic would get past the far-end access router attached to the Tetrahedron Core and not actually reach its destination out on the Internet. (I think I would get more than a few complaints if I actually did contact 5 million systems!)
- Collect the output of all of the traceroutes

4. I then wrote an analyzer script that took the output of the traceroutes and reported on the statistical distribution of paths through the Tetrahedron Core that each src-dst-ip flow selected.

It was interesting to discover that the Cisco Express Forwarding load-balancing algorithm did not yield fairly distributed usage across all links until 16,384 destinations were selected. My impression is that this is a mathematical artifact of the bucket algorithm developed by Cisco engineers; this didn't bother me, because

on a large-scale campus network such as ours we see far more than 16,384 flows running through the core at any particular time.

Case of Mistaken Identity

I am anxiously waiting, no doubt along with many other *Packet* readers, to hear the explanation as to why Cisco's "Security Advocate," Mr. Aceves, is wearing Alison's badge in the photo on page 37 [First Quarter 2004]. In most companies I am sure there are policies which greatly frown upon such activities.

—Colin A. Kopp, Province of British Columbia, Victoria, B.C., Canada

We received a record-breaking number of letters regarding the photo in the article "Security Advocates," in which Richard Aceves is shown wearing someone else's employee identification badge. Borrowing badges is not a security best practice, and is certainly not a policy that Packet or Cisco condones. When our photographer suggested the shoot take place in the lab, Richard discovered that his access to the lab had expired—Cisco requires periodic electrostatic discharge concepts exams for continued access to the labs. The lab manager was aware of the situation, and Richard was allowed to borrow a badge from one of his employees to proceed with the photo shoot. Unfortunately, we did not spot the errant badge in the photo until the article had already gone to print, but it is gratifying to see how many of our readers are paying such close attention.—Editor

Send your comments to Packet

We welcome your comments and questions. Reach us through e-mail at packet-editor@cisco.com. Be sure to include your name, company affiliation, and e-mail address. Letters may be edited for clarity and length.

Note: The *Packet* editorial staff cannot provide help-desk services.

User Group Influences New Cisco IP Telephony Features

What started with a long list of features, a request for help in prioritizing them, and a point system using so-called “Cisco bucks” back in 2001 has evolved into a valuable program for learning which Cisco IP telephony product features users really want.

Over the past few years, Cisco and CIPTUG—the official users group for companies that operate Cisco IP telephony products—have honed a process for gathering the most desired hardware and software feature ideas from CIPTUG members and prioritizing them for Cisco product managers.

“This process is a great mechanism to receive customer input for our product development,” says Marc Ayres, product manager in the Voice Technology Group at Cisco. “It’s an excellent tool, it’s been formalized, and we take the results seriously. We listen to all customer feedback, from the product enhancement requests we get from our sales force to the one-on-one customer meetings and EBCs [Executive Briefing Centers].”

CIPTUG leaders say the ability to work collectively to communicate with Cisco is central to the program’s influence. “All alone, you are one of thousands of companies out there pitching your ideas and needs to Cisco,” says Mark Melvin, Feature Advocacy Committee chairperson for CIPTUG and IP telephony network engineer for Cisco Gold Partner APPTIS, Inc. “You’re much more likely to get an important feature—get it sooner—by participating in this process.”

Customers Have Their Say

The results speak for themselves. In October 2003, more than 50 IP telephony feature requests—or one-third of the total ideas at the time—were ranked as priorities by voting CIPTUG members and shared with Cisco. Of that list, Cisco committed to developing 22, and all 22 have already been released or are on the roadmap for an upcoming release.

In the most recent voting period, during May of this year, 51 of 144 features spanning six product categories received enough points to make the priority list that Cisco product managers are reviewing now. “It helps to know that many companies from different industries would use a particular feature,” Ayres says. “We’re listening but can’t guarantee we’ll be able to fulfill every request because so many variables go into selecting a feature for a product.”

One such variable is the fact that, because Cisco adheres to industry standards and incorporates open application-programming interfaces in its product design, many companies are creating features and applications that work with Cisco IP telephony products. A new enhancement to the CIPTUG feature request system will give Cisco the ability to flag feature requests that would be better addressed by third-party ecosystem partners. Melvin explains, “This gives the membership one more avenue for sharing their needs and increases the likelihood the feature will be implemented.”

The Process in Action

CIPTUG members can submit feature ideas to the group’s Website (ciptug.org) at any time. Cisco and CIPTUG are working with six product categories: Cisco CallManager, Cisco Unity™ unified messaging software, voice gateways, IP phones, wireless IP phones, and management tools such as CiscoWorks IP Telephony Environment Monitor (ITEM).

In addition to allocating 200 points across the suggested features, each company can add comments about how that feature would be used or what it might look like displayed on a phone or device. Demographic data on the voting companies—information such as the industry and how many phones are installed—also tells Cisco how broad the use of a feature could be.

Cisco product managers and CIPTUG members meet frequently to discuss new feature requests and to improve the feature request system.

The more than 200 members of CIPTUG comprise companies in all industries. “We have a diverse set of users, from finance to healthcare to education to retail,” Melvin says, “With input from call-center operators, insurance companies, universities, and many cities and school systems—the diversity makes our input even more valuable.”

CIPTUG Member Benefits

In addition to the feature request program, CIPTUG offers Web-based presentations, discounts on training and books, collaborative opportunities through its dedicated Website, and an annual users event. The 2004 meeting will feature product roadmap presentations, panel discussions, a partner exhibit area, and opportunities to speak one on one with Cisco technology experts. The event takes place September 27–29 in Orlando, Florida. For more information, visit ciptug.org. ■

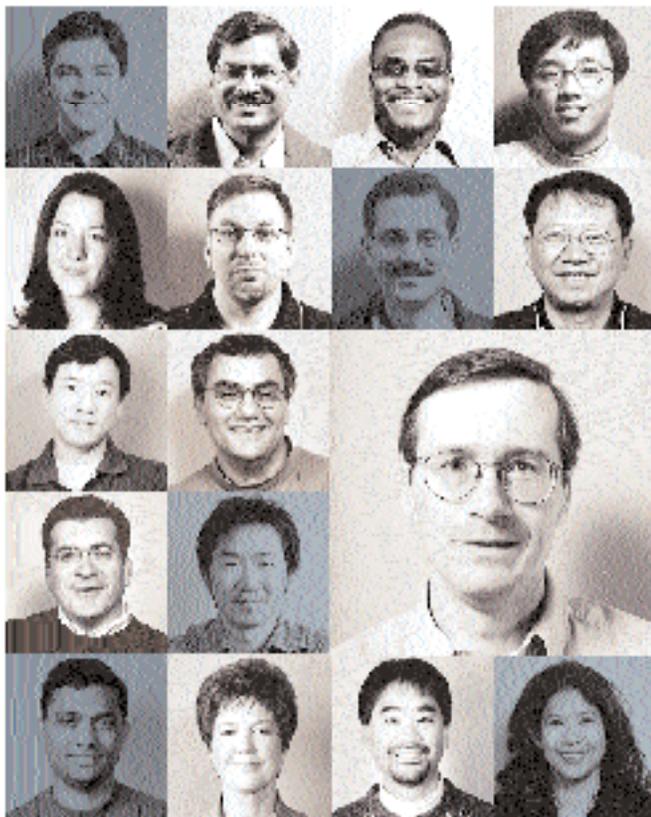
CISCO WORLDWIDE EVENTS

September 5–10	Cisco Powered Network Operations Symposium, Paris, France
September 28–30	Networkers Japan, Tokyo, Japan
November 4–6	Networkers China, Beijing, China
November 16–19	Networkers Mexico, Mexico City, Mexico
December 13–16	Networkers EMEA, Cannes, France
March 8–10, 2005	Networkers Korea, Seoul, Korea

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

Recently Announced Cisco Acquisitions

Acquired	Key Technology	Employees	Location
Actona Technologies	Developer of wide-area file-services software that helps companies store and manage data across geographically distributed offices. Actona technology will help Cisco expand the functionality of its branch-office access routers with intelligent network services that allow users at remote sites to access and transfer files as quickly and easily as users at headquarters sites. The acquired technology also allows enterprises to centralize file servers and storage and better protect and cost-effectively manage their remote office data. Actona's 48 employees based in the US and in Haifa, Israel, will join the Routing Technology Group at Cisco. Actona was founded in 2000.	48	Los Gatos, California, USA
Parc Technologies	Develops traffic engineering solutions and software for routing optimization. Parc's route server algorithms, which break up network routing problems involving complex quality-of-service constraints, can help service providers deliver high-quality services while improving network utilization and reducing capital expenditures. Cisco will incorporate the technology into its Multiprotocol Label Switching Management product line as part of the Cisco IP Solution Center. Parc's employees will join Cisco's Network Management Technology Group.	20	London, United Kingdom
Procket Networks	High-end routing company that develops concurrent services routers and has expertise in silicon and software development. The Procket engineering team and intellectual property are expected to make valuable contributions to the evolution of service provider and enterprise networks, as well as Cisco's next-generation routing technologies. About 120 employees from the company, which was founded in 1999 to build customized semiconductors for routers, will join Cisco's Routing Technology Group.	120	Milpitas, California, USA



HELLO WORLD
:)

THE EVOLUTION CONTINUES

NETWORKERS

TOKYO, JAPAN
September 28 - 30, 2004

SEOUL, KOREA
March 8 - 10, 2005

BEIJING, CHINA
November 4 - 6, 2004

GOLD COAST, AUSTRALIA
September 19 - 22, 2005

CANNES, FRANCE
December 13 - 16, 2004

www.cisco.com/networkers



Cisco Career Certifications Latest Offerings

A new storage networking specialization is the latest offering of the Cisco Career Certifications program.

“Engineers with routing and switching expertise who are called upon to support storage-area networks that are built with Cisco equipment need to know how to operate that equipment,” says Cindy Hoffmann, a program manager in the Internet Learning Solutions Group at Cisco. “The Cisco specialization trains candidates to plan, design, implement, troubleshoot, and operate Cisco MDS 9000 Series storage networking products.”

Like most Certifications courseware, content for the storage track is developed by Cisco experts but delivered by Cisco Learning Partners or training companies authorized by Cisco.

The Cisco Qualified Specialist program, which allows professionals to specialize in a particular technology such as IP telephony, network security, or wireless, is built upon the core, associate-level CCNA® and CCDA® certifications. The optical track is one exception—it does not require CCNA or CCDA status because general knowledge of networking is not necessary for managing an optical network.

Cisco also offers a storage specialization for its resellers through the Cisco Channel Partner Program.

For more information, visit cisco.com/packet/163_3e1.

Get Your Certificate by E-Mail

For certified professionals who prefer to receive an electronic certificate or want to receive their certificate more quickly, Cisco has an answer.

Candidates who complete the CCNA, Cisco Qualified Specialist, or any career certification other than CCIE® (CCIE recipients receive a plaque) can now receive the certificate electronically so it can be printed or shared with others through e-mail.

In May of this year, Cisco began offering candidates who complete their certifications a choice of a paper certificate or electronic delivery of a PDF file that cannot be modified. Either option generates the certificate, a wallet card, and a letter signed by Cisco CEO John Chambers.

Candidates who receive their first certification are notified by Cisco through e-mail and can select either a paper or electronic certificate free of charge at that time. Opting for both is US\$15. Already-certified individuals who want to order an additional paper or



FRAME IT The certificate that proves an individual has completed a Cisco Career Certification has a new look and is also available for electronic delivery.

electronic certificate can do so for \$15 per order. Additional or new orders can be made on the Cisco Certifications Community Website (cisco.com/go/cert-community) or the Cisco Career Certifications Tracking System (cisco.com/go/certifications/login). Electronic delivery takes a few days, while the paper certificate typically reaches recipients in 6 to 8 weeks.

“Some people want a printed certificate provided by Cisco that they can frame and an electronic copy they can send to prospective employers or friends and family—or even print out themselves,” says Abby Douglas, a program manager in the Internet Learning Solutions Group at Cisco.

As part of the new electronic service, Cisco updated the certificate and built a new process for verifying certificate authenticity. “It matters to those who have earned a Cisco certification that others can’t misrepresent themselves,” says Don Field, senior manager of certifications in the Internet Learning Solutions Group at Cisco.

Each certificate has a 16-digit number so that anyone examining the certificate, whether electronic or paper, can validate its authenticity on Cisco.com. In addition, certified individuals can use a Web-based tool to give others the ability to verify their certifications. “Because Cisco cannot by law verify a certification unless it has permission or a request from the certified professional, we’ve given them control of that process,” Douglas explains. ■

Is Your Network Ready for Voice?

Measuring Delay, Jitter, and Packet Loss for Voice-Enabled Data Networks

With the emergence of new applications such as voice and video on data networks, it is becoming increasingly important for network managers to accurately predict the impact of these new applications on the network. Not long ago, you could allocate bandwidth to applications and allow them to adapt to the bursty nature of traffic flows. Unfortunately, that's no longer true because today applications such as voice and video are more susceptible to changes in the transmission characteristics of data networks. Therefore, network managers must be completely aware of network characteristics such as delay, jitter, and packet loss, and how these characteristics affect applications.

Why You Need to Measure Delay, Jitter and Packet Loss

To meet today's business priorities and ensure user satisfaction and usage, IT groups and service providers are moving toward availability and performance commitments by IP application service levels or IP service-level agreements (SLAs).

Prior to deploying an IP service, network managers must first determine how well the network is working, second, deploy the service, such as voice over IP (VoIP), and finally, verify that the service levels are working correctly—which is required to optimize the service deployment. IP SLAs can help meet life-cycle requirements for managing IP services.

To ensure the successful implementation of VoIP applications, you first need to understand current traffic characteristics of the network. Measuring jitter, delay, and packet loss and verifying classes of service (CoS) before deployment of new applications can aid in the correct redesign and configuration of traffic prioritization and buffering parameters in data network equipment.

This article discusses methods for measuring delay, jitter, and packet loss on data networks using features in the Cisco IOS® Software and Cisco routers.

Delay is the time it takes voice to travel from one point to another in the network. You can measure delay in one direction or round trip. One-way delay calculations require added infrastructure such as Network Time Protocol (NTP) and clock synchronization and reference clocks.

NTP is deployed to synchronize router clocks and also when global positioning system (GPS) or another trusted reference time is needed in the network.

Accuracy of clocks and clock drift affect the accuracy of one-way delay measurements. VoIP can typically tolerate delays of up to approximately 150 ms one way before the quality of a call is unacceptable to most users.

Jitter is the variation in delay over time from point to point. If the delay of transmissions varies too widely in a VoIP call, the call quality is greatly degraded. The amount of jitter that is tolerable on the network is affected by the depth of jitter buffer on the network equipment in the voice path. When more jitter buffer is available, the network is more able to reduce the effects of the jitter for the benefit of users, but a buffer that is too big increases the overall gap between two packets. One-way jitter measurement is possible and does not require clock synchronization between the measurement routers.

Packet loss severely degrades voice applications and occurs when packets along the data path are lost.

Your success or failure in deploying new voice technologies will depend greatly on your ability to understand the traffic characteristics of the network and then applying your knowledge to engineer the appropriate network configurations to control those characteristics.

Measuring Network Performance

Key capabilities in the Cisco IOS Software can help you determine baseline values for VoIP application performance on the data network. The ability to gather data in real time and on demand makes it feasible for IT groups and service providers to create or verify SLAs for IP applications; baseline values can then be used to substantiate an IP SLA for VoIP. Cisco IOS Service Assurance Agent (SAA) technology is a component of an IP SLA solution and the Round Trip Time Monitor (RTTMON) MIB, which enable the testing and collection of delay, jitter, and packet loss measurement statistics. Active monitoring with traffic generation is used for edge-to-edge measurements in the network to monitor the network performance.

You can use the CiscoWorks Internetwork Performance Monitor (IPM) network management

application or the IOS command-line interface (CLI) to configure and retrieve data from the RTTMON MIB, or choose from a wide selection of Cisco ecosystem partners and public domain software to configure and retrieve the data. In addition, the CiscoWorks IPM features are now also available in the WAN Performance Utility (WPU) module of CiscoWorks IP Telephony Environment Monitor (ITEM) network management software.

Deploying Delay/Jitter Agent Routers

You can measure delay, jitter, and packet loss by deploying almost any Cisco IOS device, from a Cisco 800 Series Router on up.

Two deployment scenarios are possible: You can either purchase dedicated routers for SLA measurements or use current routers within the network. Place the routers in a campus network along with hosts to provide statistics for end-to-end connections. It is not practical to measure every possible voice path in the network, so place the dedicated routers in typical host locations to provide a statistical sampling of typical voice paths.

In the case of VoIP deployments using traditional phones connected to Cisco routers using FXS station ports, the router to which the phones are connected

also serves as the delay/jitter measurement device. Once deployed, the operation collects statistics and populates Simple Network Management Protocol (SNMP) MIB tables in the probe router. You can then access the data either through the CiscoWorks IPM, or through simple SNMP polling tools and other third-party applications.

Additionally, after baseline values have been established, you can configure operations to send alerts to a network management system (NMS) station if thresholds for delay, jitter, and packet loss are exceeded.

Simulating a Voice Call

One of the strengths of using Cisco IOS SAA as the testing mechanism is that you can simulate a voice call. In Cisco IOS Software Release 12.3(4)T and later, you can configure the VoIP codec directly in the CLI and simulate a voice call. This release also includes voice quality estimates, Mean Opinion Scores (MOS), and Planning Impairment Factor (PIF) scores.

Earlier versions of the Cisco IOS Software enable you to estimate a VoIP codec using the correct packet size, spacing, and interval for the measurement data and enter the appropriate parameters. The CoS can be set on data or VoIP tests, which allows you to verify how well QoS is working in the

network. Examples of how to simulate a voice call are shown below.

With Cisco IOS Software Release 12.3(4)T or later, you can use the VoIP jitter operation to simulate a test call:

```
rtr 1
type jitter dest-ipaddr 10.1.1.2 dest-port 14384
codec g711alaw
rtr schedule 1 start-time now
```

With earlier IOS releases before 12.3(4)T you can use the rtp/udp even port numbers in the range of 16384 to 32766. The user then approximates 64 kbit/s, and the packet size is 200 bytes ((160 bytes of payload + 40 bytes for IP/UDP/RTP (uncompressed)). You can simulate that type of traffic by setting up the jitter operation as shown below.

The jitter operation accomplishes the following:

- Send the request to rtp/udp port number 14384
- Send 172 byte packets (160 payload + 12 byte RTP header size) + 28 bytes (IP + UDP)
- Send 3000 packets for each frequency cycle
- Send every packet 20 milliseconds apart for a duration of 60 seconds and sleep 10 seconds before starting the next frequency cycle

The parameters in the example above give you 64 kbit/s for the 60-second test period.

$((3000 \text{ datagrams} * 160 \text{ bytes per datagram}) / 60 \text{ seconds}) * 8 \text{ bits per byte} = 64 \text{ kbit/s}$

The configuration on the router would look like this:

```
rtr 1
type jitter dest-ipaddr 10.1.1.2 dest-port 14384 num-
packets 3000
request-data-size 172**
frequency 70
rtr schedule 1 start-time now
```

Note that IP+UDP is not considered in the request-data-size, because the router internally adds them to the size automatically.

Delay/Jitter Probe Deployment Example

The two routers below would simulate voice calls of 64 kbit/s every 60 seconds and record delay, jitter, and packet loss in both directions. Note that the delay calculations are round-trip times and must be divided by two to arrive at the amount of one-way delay unless NTP is implemented for one-way delay measurements.

```
router1#
rtr responder
rtr 1
type jitter dest-ipaddr 10.1.2.1 dest-port 14384
```

```
codec g711alaw
tos 160
frequency 60
```

```
rtr schedule 1 start-time now
```

```
router2#
rtr responder
rtr 1
type jitter dest-ipaddr 10.1.1.1 dest-port 14385
codec g711alaw
tos 160
frequency 60
```

```
rtr schedule 1 start-time now
```

Command-Line Data Examples

To view the results you can use the IOS `show` command at the command line for the jitter operation. Additionally, you can use the command-line data for real-time monitoring and troubleshooting of delay, jitter, and packet loss. For an example of the CLI output, refer to cisco.com/packet/163_4b1.

Monitoring Thresholds

You can use the CLI, CiscoWorks IPM, or the WPU in CiscoWorks ITEM to configure features and monitor data. You can use this data to manage IP SLAs that have been created for VoIP. After you have determined baseline values, you can reconfigure the jitter operations to monitor the network. When predetermined delay and jitter service-level thresholds are reached or exceeded, NMS stations will be alerted.

After you have established baseline values through the initial data collection, you can monitor the delay, jitter, and packet loss levels in the network with the embedded alarm features of Cisco IOS SAA.

The Cisco IOS SAA `threshold` command sets the rising threshold (hysteresis) that generates a reaction event and stores history information for the operation. Cisco IOS SAA can measure and create thresholds for round-trip time delay, average jitter, connectivity loss, one-way packet loss, jitter, and delay.

Sample Service Assurance Threshold Configuration

```
router1#
rtr 100
rtr reaction-configuration 100 threshold-falling 50
threshold-type immediate action trapOnly
```

Understanding the traffic characteristics of the network before you deploy new advanced applications is the key to successful implementations. Delay, jitter, and packet loss greatly affect VoIP applications. Your success or failure in deploying new voice technologies will depend greatly on your ability to understand the traffic characteristics of the network and then applying your knowledge to engineer the

appropriate network configurations to control those characteristics.

◆ ◆ ◆

This article was developed by the Cisco Advanced Services Network Reliability Improvement team, which specializes in network high availability and operational best practices. In addition to using the techniques discussed in this article, you should have good operational practices in place to achieve higher levels of availability such as 99.999 (“five nines”) percent. ■

FURTHER READING

- Cisco IOS SAA technology
cisco.com/go/saa
- Cisco IOS SAA for VoIP
cisco.com/packet/163_4b2
- CiscoWorks Internetwork Performance Monitor (IPM)
cisco.com/packet/163_4b3
- CiscoWorks ITEM
cisco.com/packet/163_4b4
- White papers on operational best practices for network availability
cisco.com/packet/163_4b5
- Cisco Network Availability Improvement Services program
cisco.com/packet/163_4b6

Threat Detection

Identifying and Classifying Network Threats with Cisco IOS Software

By Ramya Venkatraman

Networks are continually becoming more intelligent and complex. Because the network plays an increasingly critical role in the daily functioning of most business environments, it is also rapidly evolving as the choice target of threats and attacks. The ever-increasing complexity of networks and intelligent services is often dwarfed by the increased sophistication of emerging network threats and attacks.

Three key areas of security that must be addressed early on are *threat detection and identification*, *attack containment*, and *mitigation*. This article provides insight into the first of these important security areas—threat detection and identification—and focuses on some key Cisco IOS® Software features that enable you to inspect traffic and identify potential threats.

First, Assess the Risk

Threats can be classified by source, internal or external; or by type, spoofing, spam, denial of service (DoS), or worms. Basic categories of attacks that threaten a network device or the network infrastructure can be broadly classified as follows:

Spoofing and impersonation—A hacker gains access by making the network think that he is a “trusted” sender. This can be due to weak or compromised user accounts and passwords or by spoofing IP addresses. Probes and scans such as port scanning, icmp unreachable messages, network commands such as whois, finger, ping, and the like, help in mining information about the network topology. In addition, protocol analysis on captured data that contains sensitive information also helps forge identity and spoof IP addresses.

DoS/distributed DoS (DDoS)—These attacks are caused by flooding the network with requests that can fill circuits with attack traffic, overwhelm network devices, slow down critical network services, and ultimately impact the network’s ability to support services. The main characteristic of any DoS/DDoS attack is hijacking a system by bombarding it with a spate of spurious traffic to process in a short span of time. Examples of such attacks include TCP SYN flooding, ICMP echo requests, TTL expiration, and UDP (fraggle) and fragmentation attacks.

Malicious code—Examples of malicious code include viruses and various worms such as Nimda, Code Red, and Slammer. Once launched, worms are self-

replicating programs and can rapidly propagate without any manual intervention. Viruses are self-replicating programs that usually require some form of human intervention to infect other systems. Malicious worms can propagate Internet-wide in a matter of a few minutes, leading to serious denial of service, downtime, and data loss in the infected hosts.

Spam—Although an indirect threat, spam is rapidly gaining ground as one of today’s main security concerns. Consulting firm Ferris Research estimates that spam now represents more than half of Internet e-mail traffic volume, and the cost of spam to enterprises in the US has more than doubled in the past year. To propagate spam, senders are increasingly relying on various tactics such as unauthorized Border Gateway Protocol (BGP) route injection, AS route hijacking, and asymmetrical routing with spoofed IP addresses.

How to Identify and Classify Threats

The first step in attack detection is gathering relevant information about its characteristics and devising a relevant threat classification strategy. This discussion focuses on identifying and classifying threats based on attack types.

Develop a network baseline. A vast majority of DoS attacks are designed to overload network devices. These attacks are usually characterized by anomalies such as an overwhelmingly large number of input buffer drops, significantly higher than usual CPU utilization levels, or link saturation. To identify such deviations from expected behavior, we first need to determine the normal behavior under a no-threat condition. This is typically accomplished by a process called *network baselining*, which helps security managers to define network performance and network resource usage for different time periods, under typical operating conditions. Investigating current link usage levels, CPU usage, memory usage, syslog entries, and other overall performance parameters are an important part of baseline profiling. Any deviations or policy violations from the network baseline should be investigated carefully, as they are potential indicators of an attack or anomaly. Examples of such behavior include:

RAMYA VENKATRAMAN is a technical marketing engineer in Cisco’s Internet Technologies Division. For the past four years, she has worked in numerous QoS and security projects at Cisco, and has been a regular speaker at Networkers and a periodic contributor to *Packet*®. She can be reached at ramyav@cisco.com.

Discover more about defending your network against threats at the Cisco Networking Professionals Connection “Security” forum: cisco.com/discuss/security.

- Large number of input buffer drops and malloc failures; could be indicators of an attack induced to exhaust resources or cause excessive memory fragmentation
- Unexplained spikes in CPU usage; could be caused by hacker-initiated scans and probes that usually consume a lot of processing power
- A sudden increase in link utilization levels; could be the result of DoS attacks or worm activity that generates inordinately large volumes of traffic
- Any other abnormal behavior such as inexplicable syslog entries, large number of threshold breaches, RMON alerts, and so on

Cisco IOS for Threat Detection and Classification

Given its ubiquitous presence across communication networks, Cisco IOS Software is the ideal platform to launch security policies to thwart attacks and help defend networks. Following are some ways to proactively identify and classify various network attacks using tools already built into Cisco IOS Software.

NetFlow with Anomaly Detection

Cisco NetFlow is the primary and most widely deployed DoS identification and network traffic flow analysis technology for IP networks in the industry

today. It is supported in most Cisco platforms via ASICs or Cisco IOS and Cisco Catalyst® Operating System (CatOS) software, and provides valuable information about traffic characteristics, link usage, and traffic profiling on the network.

NetFlow classifies packets by way of flows. Each flow is defined by its unique seven-key characteristics: the ingress interface, IP protocol type, type-of-service (ToS) byte, source and destination IP addresses, and source and destination port numbers. This level of flow granularity allows NetFlow to easily handle large-scale traffic monitoring. The NetFlow seven-tuple provides enough data for baseline profiling and determining the “who, what, when, where, and how” of network traffic.

A network traffic anomaly is an event or condition in the network characterized by a statistical abnormality compared to typical traffic patterns gleaned from previously collected profiles and baselines. NetFlow allows users to identify anomalies by producing detailed accounting of traffic flows. Deviations from the typical traffic patterns are indicative of changing traffic patterns, an early sign of potential attacks. NetFlow is usually deployed across the edge of a service provider’s network to monitor edge and peer interfaces, as these are the “typical” ingress points for most attacks. The router maintains a live Cisco IOS NetFlow cache to track the current flows.

The **show ip cache flow** command can be used to view a snapshot of the high-volume flows stored in the router cache (see figure).

IP flow information can be exported from the NetFlow cache to an external collector for further analysis. Flow data from multiple collectors can be mapped to identify the network nodes under attack and also to determine the attack characteristics. Analysis of this exported data is helpful in determining the necessary threat classification criteria enforced by IOS features such as ingress access control lists (ACLs), Network-Based Application Recognition (NBAR), and Unicast Reverse Path Forwarding (uRPF).

There are several freeware tools that can analyze NetFlow data, including cflowd, flow-tools, and autofocus. Vendors such as Arbor, Mazu, and Adlex provide GUI-based collector application tools for large-scale data collection from multiple collectors, analysis for DoS/DDoS attack detection, and centralized reporting. For example, security engineers can detect and prevent DoS attacks by using Cisco NetFlow to collect attack information such as source and destination IP, port number, packet size, and protocol type, and then send the information to a threat detection correlation tool, such as Panoptis, for anomaly detection.

Access Control Lists with IP Options

Cisco IOS access lists are the most commonly adopted technique to classify and deny access to a router at the network edge. An ACL with a series of permit statements is used to filter and characterize traffic flows of interest and trace “spoofed” packet flows back to their point of origin. Increasing numbers of DoS attacks are associated with various options being set in the IP header. Cisco IOS ACLs also have the capability of filtering packets based on various IP options in the packet header. ACL counters are used to determine which flows and protocols are potential threats due to their unexpectedly high volume. After the suspect flows are identified, permit ACLs with logging option can be used to capture additional packet characteristics.

Consider the following example:

```
access-list 101 permit icmp any any echo-reply
access-list 101 permit icmp any any echo
access-list 101 permit udp any any eq echo
access-list 101 permit udp any any eq echo any
access-list 101 permit tcp any any established
access-list 101 permit tcp any any
access-list 101 permit ip any any
```

```
interface serial 0/0
ip access-group 101 in
```

Access-list 101 permits all packets, but the individual access list entries (ACEs) can be used to categorize

the most common attack vectors, namely ICMP flooding, UDP echo attacks, and TCP SYN floods. Now the user can issue the **show access-list** command to display the access-list packet match statistics and diagnose for any potential threats.

```
Router# show access-list 101
Extended IP access list 101
  permit icmp any any echo-reply (2354 matches)
  permit icmp any any echo (1368 matches)
  permit udp any any eq echo (18 matches)
  permit udp any any eq echo any (7 matches)
  permit tcp any any established (100 matches)
  permit tcp any any (25 matches)
  permit ip any any (1015 matches)
```

The output indicates a large number of incoming ICMP echo request and reply packets—an indication of a potential ICMP flood attack or smurf attack. The log-input keyword is enabled to collect further information on the suspect packet stream such as the input interface or source IP address.

```
access-list 101 permit icmp any any echo-reply
log-input
access-list 101 permit icmp any any echo log-input
```

IP Source Tracker

To effectively block or limit an attack directed toward a host, we must first trace the origin of the threat. *Source tracking* is the process of tracing the source of the attack through the network from the victim back

show ip cache flow

```
router_A#sh ip cache flow
IP packet size distribution (85435 total packets):
 1-32  64  96  128  160  192  224  256  288  320  352  384  416  448  480
 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000

 512  544  576  1024  1536  2048  2560  3072  3584  4096  4608
 .000 .000 .000 .000 1.00 .000 .000 .000 .000 .000 .000
```

IP Flow Switching Cache, 278544 bytes
 2728 active, 1368 inactive, 85310 added
 463824 age polls, 0 flow alloc failures
 Active flows timeout in 30 minutes
 Inactive flows timeout in 15 seconds
 last clearing of statistics never

Protocol	Total Flows	Flows /Sec	Packets /Flow	Bytes /Pkt	Packets /Sec	Active (Sec) /Flow	Idle (Sec) /Flow
TCP-X	2	0.0	1	1440	0.0	0.0	9.5
TCP-other	82580	11.2	1	1440	11.2	0.0	12.0
Total:	82582				11.2	0.0	12.0

SrcIF	SrcIPAddress	DstIf	DstIPAddress	Pr	SrcP	DstP	Pkts
Eto/0	132.122.25.60	Se0/0	192.168.1.1	06	9 AEE	0007	1
Eto/0	139.57.220.28	Se0/0	192.168.1.1	06	7 08D	0007	1
Eto/0	165.172.153.65	Se0/0	192.168.1.1	06	C B46	0007	1

Source Interface (points to SrcIF column)

Flow info Summary (points to the summary table)

Flow Details (points to the flow details table)

SHOW THE FLOW The **show ip cache flow** command enables a snapshot of high-volume flows stored in the router cache.

to the attacker. Though ACLs can be leveraged to traceback attacks, there is a potential performance impact when excessive packet filters are inserted into an actual production network environment. The Cisco *IP Source Tracker* feature generates all the essential information to trace the ingress point of attack into the network all the way to the network edge, with minimal impact on performance.

After a host is diagnosed to be under attack via Net-Flow, users can enable simultaneous tracking of multiple destination IP addresses on the entire router by globally enabling the `ip source-track` command. Each line card CPU collects data about the traffic flow to individual destination IP addresses in an easy-to-use format and periodically exports this data to the router. The `show ip source-track` command can be used to display complete flow information for each inbound interface on the router including detailed statistics of the traffic destined to each IP address. This statistical granularity allows users to determine which upstream router to analyze next. By determining the source port of attack on each device, a hop-by-hop traceback to the attacker is possible. This step is repeated on each upstream router until the entry point of attack on a border router is identified.

Following is a sample configuration for IP source tracking on all port adapters in a router to collect traffic flow statistics to host address 172.10.1.1 for 3 minutes, create an internal system log entry, and export packet and flow information for viewing to the route processor every 60 seconds.

```
Router(config)# ip source-track 172.10.1.1
Router(config)# ip source-track syslog-interval 3
Router(config)# ip source-track export-interval 60
```

To display detailed information of the flow, enter the `show ip source-track <ip-address>` command

```
Router# show ip source-track 172.10.1.1
```

Address	SrcIF	Bytes	Pkts	Bytes/s	Pkts/s
172.10.1.1	P01/2	131M	511M	1538	6
172.10.1.1	P02/0	144G	3134M	6619923	143909

The output indicates interface POS 2/0 as the potential upstream attack path. You can now disable ip source-track on the current router and enable it on the upstream router to track the next preceding hop.

Unicast Reverse Path Forwarding

A large number of DoS and DDoS attackers employ spurious or rapidly altering source IP addresses to navigate around threat detection and filtering mechanisms. The uRPF feature helps mitigate attacks caused by the introduction of spoofed IP addresses into a network by discarding IP packets that lack a verifiable IP source address; uRPF

forwards only packets that have legitimate source addresses that are consistent with the IP routing table. If the source IP address is known to be valid and reachable through the interface on which the packet was received, the packet is forwarded or else dropped. Unicast reverse path checks should be deployed at the network edge or the customer edge of an ISP and should not be used in conjunction with asymmetric routing.

The uRPF feature with ACL logging adds an additional diagnostic capability by enabling reverse path forwarding check on an interface in a “pass-through” mode. In this mode, all RPF violations are logged using the ACL log-input feature. If a packet fails a unicast RPF check, the ACL is checked to determine if the packet should be dropped (using a deny ACL) or forwarded (using a permit ACL). This feature can be selectively applied to an interface to detect network threats that use spoofed IP addresses. The ACL logging counter and match counter statistics are incremented to reflect statistics for packets with spurious IP addresses. The network operator can scan the ACL log output and the counters to detect and gather more information on any potential DoS attacks.

Consider the following example:

```
int serial0/0
 ip address 172.168.100.1 255.255.255.0
 ip verify unicast reverse-path 101
!
access-list 101 deny ip 172.168.101.0 0.0.0.127
any log-input
access-list 101 permit ip 172.168.101.128
0.0.0.127 any log-input
```

Frames sourced from 172.168.101.75 arriving at serial0/0 and failing the uRPF check are logged by the ACL log statement and dropped by the ACL deny

Continued on page 88

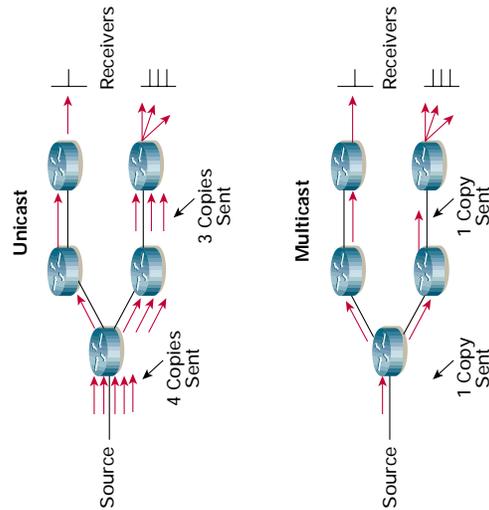
FURTHER READING

- Cisco Feature Navigator, for Cisco platform and IOS release support
cisco.com/go/fn
- Cisco NetFlow
cisco.com/packet/163_4c2
- IP access lists
cisco.com/packet/163_4c3
- IP access lists with IP options selective drop
cisco.com/packet/163_4c4
- IP Source Tracker
cisco.com/packet/163_4c5
- IP unicast Reverse Path Forwarding
cisco.com/packet/163_4c6
- RAW IP Traffic Export
cisco.com/packet/163_4c7

Courtesy of Cisco Enterprise Marketing

Why Should I Care About IP Multicast?

Many applications used in modern networks require information (voice, video, or data) to be sent to multiple end stations. When only a few end stations are targeted, sending multiple copies of the same information through the network (unicast) causes no ill effects. However, as the number of targeted end stations increases, the negative effects of duplicate packets can rise sharply. Deploying applications such as streaming video, financial market data, and IP telephony-based Music on Hold without enabling network devices for multicast support can cause severe degradation to network performance.



What Problems Need to Be Solved?

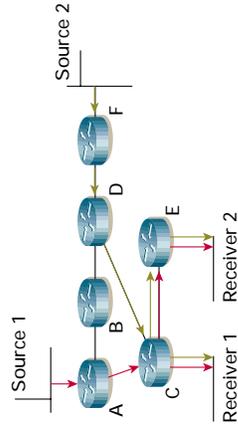
Multicasting requires methods to efficiently deploy and scale distributed group applications across the network. This is accomplished by using protocols that reduce the network load associated with sending the same data to multiple receivers and alleviate the high host/router processing requirements for serving individual connections.

Internet Group Membership Protocol

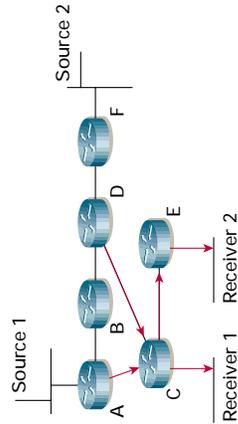
Internet Group Membership Protocol (IGMP) allows end stations to join a multicast group. Joining a multicast group can be likened to subscribing to a session or service where multicast is used. IGMP relies on Class D IP addresses for creating multicast groups. When a multicast session begins, the host sends an IGMP message throughout the network to discover which end stations have joined the group. The host then sends traffic to all members of that multicast group. Routers "listen" to IGMP traffic and periodically send out queries to discover which groups are active or inactive on particular LANs. Routers communicate with each other using one or more protocols to build multicast routes for each group.

Multicast Distribution Trees

Multicast-capable routers create distribution trees that control the path that IP Multicast traffic takes through the network to deliver traffic to all receivers. The two basic types of multicast distribution trees are source trees and shared trees.



With source trees (or "shortest path trees"), each source sends its data to each receiver using the most efficient path. Source trees are optimized for latency but have higher memory requirements, as routers must keep track of all sources.



With shared trees, multicast data is sent to a common point in the network (known as the rendezvous point or RP) prior to being sent to each receiver. Shared trees require less memory in routers than source trees, but might not always use the optimal path, which can result in packet delivery latency.

A Layer 2 switch will forward all multicast traffic, which reduces network efficiency. Two methods—Cisco Group Management Protocol (CGMP) and IGMP Snooping—were developed to mitigate this inefficient switch behavior.

CGMP allows Cisco Catalyst® switches to make Layer 2 forwarding decisions based on IGMP information. When configured on switches and routers, CGMP ensures that IP Multicast traffic is delivered only to ports that are attached to interested receivers or multicast routers. With CGMP running, any

router receiving a multicast join message via a switch will reply back to the switch with a CGMP join message. This message allows Layer 2 forwarding decisions to be made.

IGMP Snooping improves efficiency by enabling a Layer 2 switch to look at Layer 3 information (IGMP join/leave messages) sent between hosts and routers. When an IGMP host report is sent through a switch, the switch adds the port number of the host to the associated multicast table entry. When the switch hears the IGMP leave group message from a host, the switch removes the host's table entry. IGMP requires a switch to examine all multicast packets and, therefore, should only be implemented on high-end switches.

Multicast Forwarding

In unicast routing, traffic is routed from the source to the destination host. In multicast forwarding, the source is sending traffic to several hosts, represented by a multicast group address. The multicast router must determine which direction is upstream (toward the source) and which is downstream (toward the hosts). When there is more than one downstream path, the best downstream paths (toward the group address) are chosen. These paths may or may not be the same path that would be chosen for a unicast packet, a process called Reverse Path Forwarding (RPF). RPF is used to create distribution trees that loop free.

Protocol Independent Multicast

Protocol Independent Multicast (PIM) can work with whichever unicast routing protocols are used to populate the unicast routing table. PIM uses the unicast routing information to perform the multicast forwarding function, and it uses the unicast routing table to perform the RPF check instead of building up a completely independent multicast routing table. It includes two different modes of behavior for dense and sparse traffic environments.

In PIM Dense Mode, the multicast router floods traffic messages out all ports (a "push" model). If a router has no hosts or downstream neighbors that are members of the group, a prune message is sent out telling the router not to flood messages on a particular interface. Dense mode uses only source trees. Because of the flood and prune behavior, dense mode is not recommended.

PIM Sparse Mode uses an "explicit join" model, where traffic is sent only to hosts that explicitly ask to receive it. This is accomplished by sending a join message to the RP. Anycast RP provides load balancing, redundancy, and fault tolerance by assigning the same IP address to multiple RPs within a PIM Sparse Mode network multicast domain.

Cracking the Code



Insider's Tips on Earning Your CCIE in Security

By Yusuf Bhajji

Introduced in 2001, the CCIE® Security certification has evolved into one of the networking industry's most respected high-level security certifications. To become a CCIE Security expert you must pass both the written qualification exam and hands-on lab exam security. This article provides tips on resources and materials available to help you prepare for the exams.

Exam Changes

The Cisco Certifications program announced changes to the CCIE Security track this year, including significant changes to the written and lab exams. Blueprints available on the CCIE Website (cisco.com/go/ccie) outline the topics covered on the exams, so study these carefully.

Version 2.0 of the CCIE Security written exam strengthens coverage of technologies that are critical to highly secure enterprise networks. New topics such as wireless security, the Cisco Catalyst® 6500 Series security modules, and security applications such as VPN Management Solution (VMS) test candidates on security technologies and best practices. The complete blueprint for the security written exam is available online at cisco.com/packet/163_4d1. Recent changes are indicated on the blueprint in bold type.

The new revised CCIE Security lab exam preconfigures much of the core routing and switching on the devices, allowing more exam time for security-specific technologies. Topics covered more extensively on the new exam include:

- Firewalls (hardware and software)
- Virtual private networks (VPNs)
- Intrusion protection
- Identity authentication
- Advanced security technologies
- Mitigation techniques to respond to network attacks

The new content goes into effect at all exam locations beginning October 1, 2004. The preconfiguration of basic routing and switching does not make the exam easier; candidates must still configure advanced routing and switching elements and must be able to troubleshoot problems that result from the security configurations. The complete blueprint for the Security lab exam is available at cisco.com/packet/163_4d2.

Planning and Resources

An abundance of material is available to help you prepare for CCIE certification. However, be selective

and choose materials that are approved or provided by Cisco and its Authorized Learning Partners.

Books: Many Cisco Press and other vendor books are available to assist in preparing for CCIE exams. Check the current list on the CCIE Website at cisco.com/packet/163_4d3. No single resource contains all the information you need so plan to add multiple books to your collection.

Trainings: Although training is not a prerequisite for CCIE certification, the CCIE Website lists courses that might be helpful to you in studying subject matter you have less direct experience with. For a list of recommended training courses, visit cisco.com/packet/163_4d4.

Bootcamps: Many candidates ask me to recommend a security bootcamp. In my opinion, bootcamps are intended to give an overview of the lab, offer tips and tricks for exam taking, and provide mock scenarios that help you gauge your readiness. To gain the most benefit, study the technologies involved before attending a bootcamp.

Cisco.com Website: Many candidates overlook one of the best resources for useful material and technical information: Cisco.com. A plethora of sample scenarios are available on the tech support pages for each Cisco product and technology. These articles reflect current trends and demands and include sample diagrams, configurations, and invaluable IOS® **show** and **debug** command outputs.

Online Forums: Forums can be invaluable for preparation. Qualified CCIE experts and other security engineers are available around the clock to answer your queries and work through your technical problems. Some Cisco forums include:

- Cisco Networking Professionals Connection: cisco.com/go/netpro
- Cisco Certifications Community: cisco.com/go/certcommunity
Online resource for those who hold at least one Cisco certification.
- Cisco Certifications Online Support: cisco.com/go/certsupport
Q&A on certification-related topics.

Cisco Documentation CD: Make sure you can navigate the Cisco documentation CD with confidence because this is the only resource you will be allowed to refer to during the exam. Make the CD part of your regular study; if you are familiar with it, you can save time during the exam.

Practice Labs: When studying technologies such as IPSec, AAA (accounting, authentication, and authorization), firewalls, and others, you might find you can easily gain proficiency using them as standalone technologies, but integrating multiple technologies is more difficult. Find practice labs with real-world scenarios that require you to integrate multiple technologies. Practicing complex lab exercises will develop your exam strategy and help you refocus and revise your study plan.

In addition to technical skill, good time management and a solid exam-taking strategy is also important to your success. Practice labs also help you improve your time management and test-taking approach.

Equipment (home lab versus rental racks): Although acquiring a personal home lab is ideal, it can be costly to gather all the equipment to build a security rack. You can start with just a few devices—for example, three to four routers, a switch, and a Cisco PIX® Firewall. For the hardware devices that are costly to obtain, such as the IDS Sensor or VPN 3000 Concentrator, consider renting the equipment online from one of the many vendors that provide such services. Type “CCIE rack rental” in your favorite online search engine.

A current list of equipment covered on the CCIE lab exam is available at cisco.com/packet/163_4d5.

Recipe for Success

Here are some important tips and strategies from my own experience proctoring the lab exam and watching others take it.

Read the entire exam first. Read the entire test book before you begin your lab exam. Do not skip any details or sections.

Redraw your topology. Before you start the lab exam, I strongly recommend that you redraw your entire topology with all the details available. This will help you visualize your network and map the entire topology as packet flows. This map serves as a snapshot of your entire network.

Practice good time management. Make a good strategic plan to complete all the sections in the time provided. Divide the exam into categories such as Layer 2, Layer 3, backup scenarios, VPN, attacks, etc., and then work out how much time you will spend on each question, keeping in mind the point

value of each question. Allow enough time near the end of the exam to verify your solutions.

Clarify the exam questions. You must clearly understand the requirements of each question on the exam. Making assumptions can get you into trouble. During the lab, if you are in doubt, approach the proctor and verify your understanding of the requirements. Clarifying a question can make the difference between passing and failing your exam.

Keep a list. During your exam, make notes on configurations and settings as you work. For example, when configuring your device for a firewall, add access control lists (ACLs), configure filters, tunnel endpoints, and tweak routing. Keep a separate list for the items that you have not been able to address or where you have not achieved the required result and need to revisit an item.

Expect the unexpected. You might be caught off guard by an unfamiliar exam topic or question. Don't stress too much over this. Work on the things you are more comfortable with first and go back to the more difficult ones.

Practice troubleshooting. You must know how to troubleshoot problems with your configurations by using the available tools. However, although troubleshooting is important, make sure you don't lose too much time troubleshooting a 2- or 3-point question. Try to move on and return again later.

Test your work. Never rely on a configuration you did in the early hours of the exam. An item that you configured a few sections earlier could become broken and nonfunctional. Always validate your solutions toward the end of the exam. Keep in mind that points are awarded for working configurations only.

Do not memorize. Your goal should be to master the technology and the architecture.

A Final Word

I hope that the preceding tips and information will encourage you to pursue CCIE certification. Achieving your CCIE can be a great source of satisfaction and can boost your career to the next level. The secret to success on CCIE, as with most endeavors, is motivation, dedication and consistency. In the long run, being an expert in the field of security networking is not just a destination, but an ongoing journey.

For more information, visit the CCIE Website at cisco.com/go/ccie. ■

FAHIM HUSSAIN YUSUF BHAIJI, CCIE No. 9305, is the content lead for Cisco CCIE security certification and exam proctor in Sydney, Australia. Bhaiji recently published a book on preparing for CCIE Security, *CCIE Security Practice Labs* (Cisco Press 2004). He can be reached at yusuff@cisco.com.

Reader Tips

Packet® thanks all of the readers who submitted technical tips this quarter. While every effort has been made to verify the following reader tips, *Packet* magazine and Cisco Systems cannot guarantee their accuracy or completeness, or be held responsible for their use.

Configuration

TIP Using X.25 to Configure Integrated Systems

I use the X.25 Protocol to integrate Call Data Records (CDR) data for billing systems (mediation). These are primarily mobile switches using X.25 protocols to integrate the CDR, remote terminal (OMT or CTL) and OMCS. I use X.25 over TCP/IP (XOT) to integrate all of these functions using reliable IP media. Traditionally, X.25 provides 64k bandwidth, but by changing the clock parameters you can also achieve more than 64k. The following configuration is useful for anyone working with Global System for Mobile Communications (GSM) operators or for PSTN network providers.

Router # x25 routing xot-use-interface-defaults

```
interface Serialx/x
  description XXXXXXXX
  no ip address
  encapsulation x25 dce ietf
  x25 address XXXXXXXX
x25 htc 32
  x25 win 7
  x25 wout 7
  x25 ips 256
  x25 ops 256
  x25 subscribe flow-control always (this is the most
  important command)
  clockrate 64000
  lapb T1 2000
  lapb T2 800
  lapb N2 7
  lapb k 2
```

Route:

```
Router # x25 route < x.25 address > xot < remote IP
address >
```

—*Muhammad Ali, Mobilink-GSM, Islamabad, Pakistan*

TIP Avoiding Cisco CallManager Application Server Reconfiguration When Using DID Numbers

Because enterprise-level IP telephony networks are so dependent on system features, when integrating these networks with application servers such as Cisco IPCC Express, Cisco Personal Assistant, Cisco Unity™, and Cisco Meetingplace, I create private internal directory

numbers when I configure the computer telephony interface (CTI) route points for these services. Many customers require that the application servers must accommodate PSTN-based calls through the use of Direct Inward Dial (DID) access numbers. To do this, create a CallManager Translation Pattern that uses a DID number which then redirects calls to the private directory number of the specific application CTI route point. When a customer wants to add, delete, or change DID numbers, this method is much easier to manage instead of doing an elaborate reconfiguration of CTI route points and application server configurations.

—*Michael Cotrone, CCIE® No. 8411, Datatnet Services, Inc., Greensboro, North Carolina, USA*

Troubleshooting

TIP Recovering Lost Passwords on Remote Devices

Configuring a Simple Network Management Protocol (SNMP) read-write (RW) community ahead of time enables me to modify the configuration of a device if I need to recover a lost password from a remote router or switch. I use these steps:

1. Set the copy mode (1.- TFTP; 3.-RCP): `snmpset ipAddress RW-Community .1.3.6.1.4.1.9.9.96.1.1.1.1.2.83119 i 1`
2. Set the source configuration type to copy (1.- Network; 3.-Startup-config; 4.-Running-Config): `snmpset ipAddress RW-Community .1.3.6.1.4.1.9.9.96.1.1.1.1.3.83119 i 4`
3. Set the destination configuration type to copy (1.- Network; 3.-Startup-config; 4.-Running-Config): `snmpset ipAddress RW-Community .1.3.6.1.4.1.9.9.96.1.1.1.1.4.83119 i 1`
4. Set the TFTP server IP address: `snmpset ipAddress RW-Community .1.3.6.1.4.1.9.9.96.1.1.1.1.5.83119 a TFTP-SRV-ipAddress`
5. Set the name of the file that contains *my device* configuration: `snmpset ipAddress RW-Community .1.3.6.1.4.1.9.9.96.1.1.1.1.6.83119 s "Mydevice Config.txt"`
6. Set the **create and go** command: `snmpset ipAddress RW-Community .1.3.6.1.4.1.9.9.96.1.1.1.1.14.83119 i 1`

Then I modify the password in a file named *My-deviceConfig.txt* and run the command again, modifying the following lines:

1. Set source configuration type to copy (1.-Network; 3.-Startup-config; 4.-Running-Config):

```
snmpset ipAddress RW-Community
.1.3.6.1.4.1.9.9.96.1.1.1.1.3.83119 i 1
```
2. Set destination configuration type to copy (1.-Network; 3.-Startup-config; 4.-Running-Config):

```
snmpset ipAddress RW-Community .1.3.6.1.4.1.9.9.96.1.1.1.1.4.83119 i 4
```

Be careful when you modify and upload the configuration to the device, and remember that the destination is Running-Config, so you must ingress to the device to change the password again and then write this to the startup configuration.

For more information about copying configurations using SNMP, see cisco.com/packet/163_4f1.

—Rodrigo Barroso, *Petrobras Energia S.A., Buenos Aires, Argentina*

TIP Troubleshooting DoS Attacks

Multiple large-sized packets injected into your network from any source, including a host PC, can bring your network to a dead crawl. In the worst case, they can even shut down operations. To determine which host or node is sending or receiving suspiciously large and multiple “packets” (no pun intended),

enable **ip accounting output-packets** in the interface that you suspect they pass through. Then use the command **sh ip accounting output-packets** to view the output in real time. Even packet and byte sizes are displayed, which can help you identify what kind of traffic is present in your link. For example:

```
Router(config)# interface FastEthernet 0/1
Router(config-if)# ip accounting output-packets
Router# sh ip accounting output-packets
```

—Alfred Romero Jr., *WeCare Technology Services Corp., Makati City, Philippines*

Editor's note: The preferred, more scalable, method is to use NetFlow on ingress interfaces to try to find the type of traffic (see cisco.com/packet/163_4f2). Because NetFlow keeps statistics on flows, you can more easily isolate the protocols involved. To enable NetFlow on interfaces, use the interface configuration command **ip route-cache flow**. Support for NetFlow can vary depending on your platform and code version. For older platforms that do not support NetFlow, IP accounting can be useful, although it tends to negatively affect performance. ■

SUBMIT A TIP

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

Tech Tips

Learn about wireless security capabilities in Cisco wireless products. New centrally managed, dynamic per-user, per-session Wired Equivalent Protocol (WEP) capabilities in Cisco Aironet® Software Release 11.0 and Cisco Access Control Server (ACS) 2.6 address wireless security issues. cisco.com/packet/163_4g1

Troubleshoot wireless network connectivity. This document helps you identify and troubleshoot common wireless network connectivity problems including configuration, interference, and cable issues. cisco.com/packet/163_4g3

Learn about DiffServ tunneling modes for MPLS networks. This document describes the Differentiated Services (Diff-Serv) Tunneling Modes available for implementation in Multiprotocol Label Switching (MPLS)-based network environments. cisco.com/packet/163_4g4

Troubleshoot Cisco IP Phone connection issues. This document describes how to solve connectivity problems with the Cisco VT Advantage video telephony solution. cisco.com/packet/162_4g5

Read about best practices for NTP network management. This white paper describes a hypothetical process definition for conducting network management functions for the Network Time Protocol (NTP), which organizations can customize in order to meet internal objectives. Includes process and task definitions, as well as configuration and report format examples. cisco.com/packet/162_4g6

Learn about security and VPN resources. View the free, on-demand Cisco technical support seminar, “Using the Cisco Technical Support Website for Security and Virtual Private Network Issues.” cisco.com/techsupport/seminars

Deploying Video Telephony

Cisco CallManager 4.0 extends voice features to video over a common, user-friendly infrastructure that can be deployed to the desktop.

By Tom Schepers

NETWORKERS 2004

This article is based on a session presented at the Cisco Networkers 2004 users conference. To learn more about Networkers, visit cisco.com/networkers.

Video telephony leverages the intelligence of IP telephony to provide advanced features that are not available in traditional IP videoconferencing deployments: call forwarding, call hold, call park, class of service restrictions, ad-hoc conferencing, bandwidth controls, enhanced digit manipulation, and call rerouting, to name a few. The result? Enterprises can retain their existing H.320 and H.323 investments while benefiting from a user-friendly, more feature-rich environment for large-scale video deployments.

Video communication capabilities have been integrated into Cisco CallManager 4.0—extending several voice features to video that benefit end users, network administrators, and enterprises as a whole (for a comprehensive list of Cisco CallManager video telephony features, visit cisco.com/packet/163_5a1). Among the benefits, users enjoy a simple interface, leveraging the same dial plan structure as their IP phone deployment in a familiar user environment. With the ability to create multipoint conferencing, users can also manage more effective meetings and schedules. For administrators, video telephony provides a single infrastructure that leverages a common graphical interface and common features for all voice and video communications. A common IP infrastructure for all communications not only provides an enterprise with reduced cost of ownership and faster return on investment (ROI), but also provides greater reliability and ease of maintenance because video calls do not have to be done over separate ISDN lines. This allows users to more readily and easily adapt to a system that can now be deployed to the desktop.

Video Call Control and Resilience

Video call control within Cisco CallManager 4.0 functions essentially the same as it does for audio. Call setup signaling is handled by CallManager, resolving dialed numbers based on the dial plan deployed within the CallManager clusters. The Cisco IOS® Gatekeeper provides a logical trunk to the CallManager cluster, which allows existing H.323 and H.320 devices to be integrated into CallManager (see figure, page 24). Video calls typically include Real-Time Transport Protocol (RTP) streams, in each direction, for audio, video, and far-end camera control (FECC), and a sequence of call control signaling messages. This bearer traffic is not handled by CallManager but is routed directly between endpoints.

Because Cisco CallManager routes all H.323 call signaling (for example, H.225/H.245), the enhanced functionality, such as call forwarding, call park, and shared lines, can be transparently provided for H.323 devices. In addition, digit manipulation is not reflected back to the calling endpoint, so there are no special requirements for the endpoints to support having their calls rerouted or manipulated.

For video calls, Cisco CallManager 4.0 includes the additional logic to handle negotiation of the video codec (H.261, H.263), resolution, frame rate, and H.323 annexes. The region and location settings for admission control have also been enhanced to provide for accounting of video bandwidth on a per-call and aggregate basis. For video calls, the negotiated bandwidth for an H.323 device typically includes both audio and video; for example, a 384-kbit/s video call is comprised of 64-kbit/s audio and 320-kbit/s video channels. Video capabilities are provided for calls between devices within a cluster and between clusters (for example, via inter-cluster trunks).

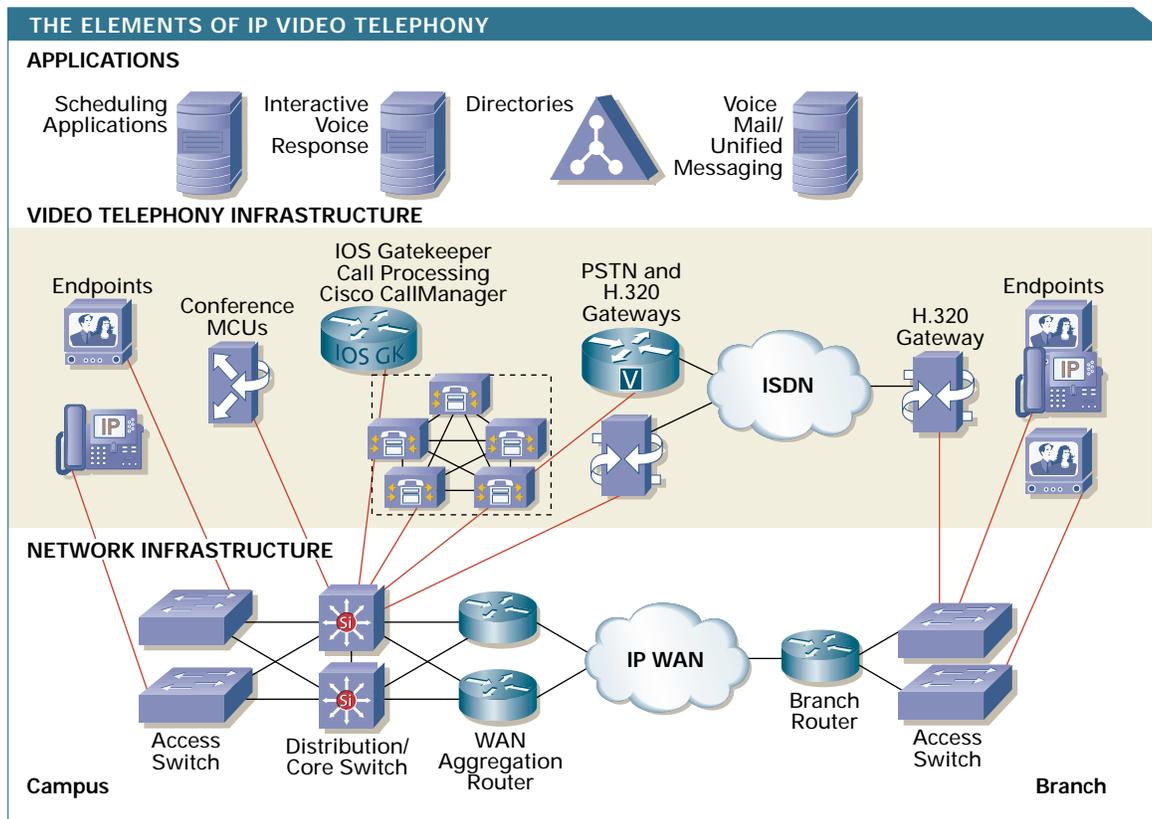
Cisco CallManager clustering, as well as Cisco IOS Gatekeeper clustering using the Alternate Gatekeeper (Alt-GK) feature, provide for a resilient environment to protect video telephony from component failures. While CallManager and many H.323 devices support Alt-GK, not all H.323 devices do, in which case Hot Standby Router Protocol (HSRP) can be used to provide resilience of the gatekeeper elements. Alt-GK is a more robust implementation than using HSRP because Alt-GK provides for load balancing and the ability to locate gatekeepers in diverse network locations (HSRP requires that the gatekeepers be on the same IP subnet).

Skinny Client Control Protocol (SCCP) video endpoints—whether a Cisco VT Advantage USB camera used in conjunction with a Cisco IP Phone, or a Tandberg video endpoint that uses SCCP—register directly to the Cisco CallManager. For calls to video-capable endpoints, CallManager opens the logical channels for video automatically if the originating endpoint also has video capabilities as defined in the endpoint setup in CallManager. SCCP endpoints will also provide a richer set of messaging to end users (for example, indicating the reason for a failed call, such as unavailable bandwidth). Endpoint configuration, listed under the “Phones” menu on CallManager, allows users to define the necessary adjunct definitions for the endpoint, such as region, location, call forwarding on busy or no answer, Automated Alternate Routing (AAR) groups, digit manipulation or translations, calling search space, partition, Media Resource Group List (MRGL), and directory number(s).

In addition, SCCP video endpoints behave like an IP phone. For example, when users take the device off hook to make a new call, a

RINGING UP VIDEO

Video call control within Cisco CallManager 4.0 functions essentially the same as it does for audio. Call setup signaling is handled by CallManager, resolving dialed numbers based on the dial plan deployed within the CallManager clusters.



dial tone is played; users can press the phone's softkey buttons to invoke features and supplementary services.

Alternate Routing Using the PSTN

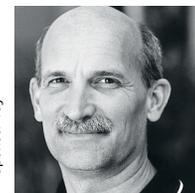
H.323 gateways can be used for alternate routing of video calls over the public ISDN network. SCCP, Media Gateway Control Protocol (MGCP), and IOS H.323 gateways can also be used for alternate routing of video calls as audio-only using the PSTN. Cisco CallManager retries a video call as audio-only under certain conditions: upon failure of region and locations admission control, when using H.323 video gateways to provide routing over the PSTN in the event of admission control or possible network failure, or when the gateways are audio-only devices. Unlike with traditional H.323 deployments, the user does not have to redial to get the alternate route. CallManager will manipulate the dialed digits as necessary, adding a PSTN access code (9, for example), along with the long-distance access code and area code, to create a fully qualified number for routing via the public network. An SCCP endpoint will provide indications that alternate routing is in effect. AAR is available for calls between locations managed by the same CallManager cluster, and for calls between CallManager clusters.

Multipoint Conferencing

Cisco CallManager supports several methods for users to participate in multipoint video calls, including ad hoc, scheduled, and reservationless. Each method requires a Cisco IP/VC 3500 Series Multipoint

Conference Unit (MCU), which supports both SCCP and H.323 protocols. SCCP is used for ad-hoc conferences, and H.323 is used for scheduled and reservationless conferences. With the phone or SCCP video endpoint interface, a user can establish an ad-hoc videoconference by pressing the "Conf" softkey and then dialing additional participants into the call. The participants can be on any other SCCP endpoint or audio-only endpoints, as well as H.323 or H.320 video endpoints.

H.323 devices typically register to an H.323 gatekeeper and are defined within CallManager as "H.323 Clients." The administrator can apply settings to each endpoint, such as directory number, region, location, MRGL, and so on. H.323 MCUs and H.323/H.320 gateways, such as the Cisco IP/VC 3500 Series videoconferencing products, also register to the gatekeeper and are defined in CallManager as "H.323 Gateways." The administrator can then apply settings to the device, but instead of defining a directory number, route patterns are used to reach these devices. A route pattern can point either directly to the device in



Spencer Toy

TOM SCHEPERS, consulting systems engineer at Cisco, is the presenter of "Designing and Deploying IP Video Telephony Networks" at the Networkers 2004 Cisco users conference. He can be reached at tscheper@cisco.com.

CallManager or to a route list containing one or more route groups to provide alternate routing in the event that one of the MCUs or gateways is unavailable.

Alternatively, the route pattern could point to an H.225 gatekeeper-controlled trunk. For calls to an H.323 MCU conference, the route pattern would be constructed to match the service prefix defined in the MCU for the type of conference you want to join. For example, a service for continuous presence, H.263, 384-kbit/s, 30-fps conferences may be defined as 82* (where the * can be any digit(s) 0 through 9 and any number of digits). The CallManager will be configured with a route pattern that states all calls beginning with 82 (such as 82XXX) are to be routed to the MCU, either directly by defining the MCU as an H.323 gateway in CallManager or via the H.225 trunk; in the latter case, the gatekeeper receives the call setup and forwards the call to the MCU registered with that service prefix.

Likewise, for calls to an H.320 gateway, the route pattern would also be constructed to match the service prefix configured in the gateway. But in this case, the service prefix simply defines how many ISDN channels the call should use. For example, a 384-kbit/s service may be defined as service prefix 9#*. The CallManager would be configured with a route pattern that states all calls beginning with 9 (such as 9.@, where @ represents all PSTN patterns supported by the North American Numbering Plan, or NANP) are to be routed to the gateway, either directly by defining the gateway as an H.323 gateway in CallManager or to a pool of gateways contained in a route list/route group(s), or via the H.225 trunk. In the latter case, the gatekeeper receives the call setup and forwards the call to the gateway(s) registered with that service prefix.

With digit manipulation, users do not have to dial the # character. A user simply dials "9+1+area code+number," for example, and CallManager can prepend the # before routing the call to the gateway.

When using the gatekeeper to reach the gateway(s), the gateways use Resource Availability Indications/Resource Availability Confirmation (RAI/RAC) messaging to tell the gatekeeper whether or not there are enough open ISDN B-channels available to support another call. If there are not, the gateway sends an RAI message indicating that it should be taken out of the gatekeeper's list of available gateways. It will send another RAI message when enough channels are open so that the next call request can be successfully serviced.

Call Accounting and Performance Monitoring

Call accounting, using the Cisco CallManager CDR Analysis and Reporting (CAR) tool, provides additional information for video calls, including but not limited to:

- IP addresses and port numbers
- Codec (H.261, H.263)
- Bandwidth (in each direction)
- Resolution (CIF/QCIF, for example)
- Calling name/number
- Called name/number

Reports can be generated using the CAR tool to monitor the amount of bandwidth being used for video, the number of calls made by a specific endpoint, and usage statistics for MCUs and gateways. Performance monitoring can be used to track the number of active calls; calls completed; calls rejected due to lack of resources; locations bandwidth available and the number of times bandwidth at a location has been exceeded; and much more. This is done using the Real-Time Monitoring Tool (RTMT) in Cisco CallManager Serviceability.

See the sidebar, "Cisco CallManager Video Telephony Configuration," on page 26 for a summary of configuration steps.

H.323 Integration

In recent years, enterprises have increasingly been investing in H.323 videoconferencing solutions. As such, the evolution to video telephony must provide for the integration of existing H.323 equipment, including endpoints, gateways, MCUs, and scheduling systems. Cisco CallManager provides this integration by using the Cisco IOS Gatekeeper. All H.323 devices continue to register to the gatekeeper, but all H.225 and H.245 call signaling is routed to CallManager for dial plan resolution, call accounting, and supplementary services. The Cisco IOS Gatekeeper uses a default routing mechanism that results in all call setup signaling initiated by H.323 devices to be forwarded to CallManager for resolution. CallManager then takes control of the call and performs all digit analysis, digit manipulation, bandwidth controls, and class of service restrictions. Conversely, when CallManager signals a call setup to an H.323 device that is defined within CallManager (not one that is accessed via a route pattern and H.225 trunk), the gatekeeper does not need to be present. Because CallManager already knows the IP address of the H.323 device, CallManager initiates call setup directly to the device.

H.323 endpoints offer varying degrees of integration. Although they cannot initiate the supplementary services available for SCCP endpoints, H.323 endpoints can take advantage of the unified dial plan, AAR, shared lines, hunt groups, call accounting, and other features that provide intrinsic value to the H.323 deployment.

While conforming to the standard, not all H.323 endpoints will support the same services, particularly supplementary services. With Empty Capabilities Set (ECS), an endpoint can be the target of any supplementary services (such as call hold, park, conference,

Cisco CallManager Video Telephony Configuration

Step 1: Define CAC parameters for video, both regions and locations.

Step 2: Define any SCCP bridges.

Step 3: Add H.323 MCUs, either via a route pattern to the H.225 trunk to the gatekeeper, or define the MCUs within CallManager directly as "H.323 Gateways." Define route patterns for each MCU service prefix.

Step 4: Define the MRGLs required to ensure that the appropriate resources are allocated, depending on the conference initiator.

Step 5: Add H.323 gateways, either via a route pattern to the H.225 trunk to the gatekeeper, or define the gateways within CallManager directly. If you choose the latter, also define the AAR configuration and the route list/route group this gateway should be a member of. Digit manipulation for prefixing required digits to access the PSTN should be part of this configuration.

Step 6: Define the H.323 gatekeeper(s).

Step 7: Define the H.225 trunk(s) to the gatekeeper(s).

Step 8: Define endpoints, along with the required attributes such as directory numbers, AAR groups, and MRGL.

Step 9: Configure the "Retry Video Call as Audio" setting on each type of video-capable device according to whether you want CallManager to perform this behavior or reroute the call via AAR instead. If you choose the latter, configuration of AAR groups and External Phone Number Mask on each endpoint is also required.

For all of the device configuration steps, you will also need to define the advanced settings such as partition, calling search space, and MRGL. Finally, maintain the system by using all available monitoring and troubleshooting tools, such as RTMT, CAR, the embedded call trace facilities, and alarms/traps in CallManager Serviceability.

or transfer) but cannot initiate these functions. Without ECS support, an H.323 endpoint will drop the call if these services are invoked to it.

Deployment Scenarios

The deployment models available for video telephony are essentially the same as for IP telephony, including single site; multisite centralized call

processing; multisite distributed call processing; Voice- and Video-Enabled VPN (V³PN) and telecommuter VPN environments; service provider managed and hosted multitenant environments; and so on. The video devices deployed can consist of SCCP only, H.323 only, or a combination of both. MCUs, gateways, and gatekeepers fit into each of these scenarios as well.

Deploying SCCP devices is straightforward, because they register directly to the CallManager, download their configuration from a central TFTP server, and are under the complete control of CallManager. Deployments that include H.323 devices require the addition of an H.323 gatekeeper. The gatekeeper and CallManager are linked via an H.225 trunk. Depending on the deployment model, the gatekeeper serves as either an endpoint gatekeeper (the gatekeeper that all the H.323 endpoints register; it is configured to route all calls to CallManager) or an inter-cluster trunk gatekeeper (the gatekeeper that provides dial plan resolution and CAC between different CallManager clusters in a distributed call processing model). In both cases, the gatekeeper requires the definition of one or more local zones, zone prefixes, and technology prefixes.

For centralized deployments, all call processing is handled by a cluster of CallManagers located at the central site. Branch offices in this environment contain no local call processing but are controlled by the central CallManager cluster. One or more endpoint gatekeepers would also reside at the central site, adjacent to the CallManager cluster, providing the integration between H.323 devices and the Cisco CallManager 4.0 deployment. It is recommended that the endpoint gatekeeper have different zones defined for each type of endpoint: one for endpoints, one for the CallManager servers, one for MCUs, and one for gateways. Zone prefixes are used to route all calls to the CallManager zone, and technology prefixes are used to route the call to the correct CallManager server. Following is an example endpoint gatekeeper configuration:

```
gatekeeper
zone local endpoints xyz.com
zone local callmanagers xyz.com
zone local gateways xyz.com
zone local mcus xyz.com
zone prefix callmanagers 0*
zone prefix callmanagers 1*
zone prefix callmanagers 2*
zone prefix callmanagers 3*
zone prefix callmanagers 4*
zone prefix callmanagers 5*
zone prefix callmanagers 6*
zone prefix callmanagers 7*
zone prefix callmanagers 8*
zone prefix callmanagers 9*
zone subnet callmanagers 10.1.1.10/32 enable
no zone subnet callmanager default enable
zone subnet gateways 10.1.1.11/32 enable
```

```

no zone subnet gateways default enable
zone subnet mcus 10.1.1.12/32 enable
no zone subnet mcus default enable
no zone subnet endpoints 10.1.1.10/32 enable
no zone subnet endpoints 10.1.1.11/32 enable
no zone subnet endpoints 10.1.1.12/32 enable
no zone subnet endpoints 10.1.1.13/32 enable
no zone subnet endpoints default enable
gw-type-prefix 1# default-technology
no use-proxy endpoints inbound-to-terminal
no use-proxy endpoints outbound-from-terminal
endpoint ttl 60
no shutdown

```

The H.225 trunk is defined in CallManager Administration to register in the “callmanagers” zone with the technology prefix 1#. The zone prefixes applied to the callmanagers zone force all calls to be routed to it, and then the default technology prefix is used to route the call to the CallManager H.225 trunk. This procedure ensures that endpoints are not allowed to call the MCUs or gateways directly, so CallManager remains in control of all call routing and is able to generate call detail records (CDRs) for every call.

The MCUs and gateways can either be located centrally or placed in each branch office to provide localized services specific to a branch, such as local gateway resources to access the public ISDN/PSTN network. Device pools and MRGLs control which MCU is used by each branch, and calling search spaces and route lists/route groups control which gateways are used. CallManager controls all bandwidth and CAC functions, and AAR is available if the WAN is oversubscribed.

The endpoint gatekeeper can be deployed in a redundant fashion by using HSRP or Gatekeeper Clustering. Gatekeeper Clustering is a newer, more efficient mechanism available in Cisco IOS Software Release 12.2(2)T or higher. It has many benefits over HSRP including the ability for the gatekeepers to be geographically dispersed to provide even greater fault tolerance and special redundancy; every gatekeeper in the cluster keeps active state of which endpoints are registered and which calls are active. However, it requires the endpoints to support the H.323v3 Alt-GK field passed back to the endpoint during registration.

Many H.323 video endpoints on the market do not yet support the Alt-GK feature, and so HSRP can be used instead. HSRP is transparent to the endpoints; however, with HSRP the gatekeepers share a logical (virtual) IP address and, thus, must be physically located in the same IP subnet. In addition, only the active gatekeeper maintains state; the others are essentially asleep until they sense that the active router is down, at which point the secondary gatekeeper will come on line without any knowledge of pre-existing calls.

For distributed deployments, each CallManager cluster handles local call processing for the devices and branches that it controls, as described above, and an inter-cluster trunk gatekeeper may be deployed to provide dial plan resolution and CAC between the different CallManager clusters. It is recommended that this gatekeeper be configured with a zone for each CallManager cluster. Zone prefixes are then applied to route calls between the different zones, based on the directory numbers that each CallManager cluster services; the default technology prefix 1# is used to route the call to the inter-cluster trunk registered within that zone. Gatekeeper bandwidth commands are applied to limit the amount of bandwidth allowed between each zone. Bandwidth commands can also be used to limit the amount of bandwidth allowed per call. Following is a sample configuration for two clusters located in different sites, St. Louis and Chicago:

```

gatekeeper
zone local stlouis xyz.com 10.2.1.1
zone local chicago xyz.com
zone subnet stlouis 10.2.1.0/24 enable
no zone subnet stlouis default enable
zone subnet chicago 10.2.3.0/24 enable
no zone subnet chicago default enable
zone prefix stlouis 1636*
zone prefix chicago 1773*
gw-type-prefix 1# default-technology
bandwidth interzone stlouis 1408
bandwidth session stlouis 768
bandwidth interzone chicago 1408
bandwidth session Chicago 768
endpoint ttl 60
no shutdown

```

The bandwidth interzone command regulates the aggregate amount of bandwidth allowed to and from that zone, and the bandwidth session command regulates the amount of bandwidth allowed per call. The H.323 specification dictates that the bandwidth values be entered as 2 x the call bit rate. For example, a 384-kbit/s video call would be entered as 768 in the gatekeeper. A G.711 audio-only call would be entered as 128 in the gatekeeper. The interzone command is the sum of all audio and video calls that you want to allow to and from that zone. For example, the 1408 number used in the configuration above would allow for 5 G.711 audio calls and one 384-kbit/s video call ($128 * 5 + 768 = 1408$).

◆ ◆ ◆

To learn more about deploying video telephony using Cisco CallManager, see the *Cisco CallManager System Guide, Release 4.0* at cisco.com/packet/163_5a2. For more on Cisco IP/VC 3500 Series MCUs, gateways, and enhanced media processors, see the corresponding administration guides at cisco.com/packet/163_5a3. ■

Deflector Shield

Cisco acquires Riverhead Networks for mitigating distributed denial-of-service attacks.

By Gail Meredith Otteson

Distributed denial-of-service (DDoS) attacks are “weapons of mass disruption.” Unlike attacks that compromise data or steal information, DDoS attacks can shut down business for days, or even weeks. Until recently, enterprises and service providers frequently had to resort to defensive tactics which, because of their lack of granularity, often had the effect of “completing the DoS” for the attacker. “Either way, the attacker wins, because business stops,” says Steve Woo, director of marketing in Cisco’s Internet Switching Business Unit, Layer 4–7 Services. “People know when they are being attacked—that’s the easy part. The issue is to stop the attack without stopping business.”

Business continuance in the face of disruption is essential to business survival. Losses due to DDoS interruptions can be devastating, affecting revenues and productivity. Attacks can increase IT expenses and expose organizations to litigation. Customer confidence is damaged—sometimes permanently. The Yankee Group reports that a series of DDoS attacks in February 2000 against Amazon, eBay, Yahoo, and other major Websites caused an estimated cumulative loss of US\$1.2 billion. Today, potential losses are even higher.

DDoS attacks have grown in scale and stealth, making them harder to detect and difficult to mitigate. A typical DDoS attack recruits hundreds, or thousands, of “zombie” hosts to launch an attack against a single target. Zombies are drafted from the millions of unprotected computers that are connected to the Internet through high-bandwidth, “always-on” connections. Attackers implant malicious software onto these machines and then launch attacks with a single command. Owners are unaware that their PCs are sending undetectable volumes of DDoS traffic. Multiplied over thousands of zombies, the cumulative amount of traffic thrown at a target overwhelms its resources, making it unavailable to legitimate users.

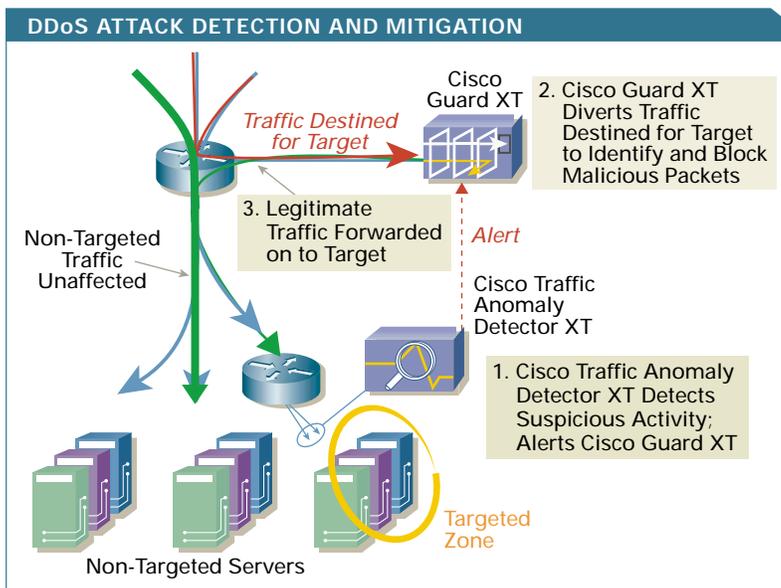
Attack targets can include a provider network infrastructure or any data center resource. Targets might be e-commerce, database, and application servers; network services such as Web, Domain Name System (DNS), and e-mail systems; network routers; security devices such as firewalls and intrusion detection systems (IDS); and access links.

There are many ways to detect attacks in progress, and Cisco has developed many tools and techniques to block them. However, fine-grained, application-specific mitigation has been a challenge—that is until 2002, when Riverhead Networks introduced a solution that blocks malicious traffic and allows legitimate transactions to continue, resulting in business as usual.

The Self-Defending Network

Cisco completed its acquisition of Riverhead in March 2004, incorporating the company’s unique DDoS detection and mitigation technology into the Cisco security portfolio. “There’s nothing else like it on the planet,” says Roland Dobbins, network engineer in the IT Internet Services Group at Cisco. “It’s an important tool in the Cisco security toolkit.”

The DDoS solution includes the Cisco Guard XT and Cisco Traffic Anomaly Detector XT, adding critical functionality to the Cisco Self-Defending Network strategy (see “The Self-Defending Network,” *Packet*® First Quarter 2004, cisco.com/packet/163_5b5), which can automatically identify threats, react in a situationally appropriate manner, and ensure service continuity during an attack. The Guard and Detector are vital components in a multilayer defense strategy for public-facing data centers and Web services in



BUSINESS AS USUAL Many devices can detect a DDoS attack and alert the Cisco Guard. The Guard tells the router to divert all traffic destined for the targeted device to itself. It analyzes and “scrubs” traffic, dropping malicious packets, then forwards legitimate traffic to the target, maintaining business continuity during an attack.

large enterprises and government agencies, and for service providers that offer managed hosting and Web connectivity services. Cisco is already adapting the Detector and Guard into integrated modules for its Cisco 7600 Series Router and Cisco Catalyst® 6500 Series Switch platforms.

The solution protects against two basic types of attacks: flooding and application.

- A **flooding attack** overwhelms network links and equipment with a high volume of TCP, UDP, or Internet Control Message Protocol (ICMP) packets, rendering network resources unavailable for valid traffic and causing inline security devices to fail under the load.
- An **application attack** uses the expected behavior of protocols such as TCP and HTTP to the attacker's advantage by tying up computing resources and preventing them from processing legitimate transactions and requests. Examples include HTTP half-open and HTTP error attacks.

Other Security Tools

Cisco has devised many tools and techniques to detect and mitigate DoS attacks using existing technologies. These devices serve critical security roles in a defense-in-depth architecture and are also important tools in the Cisco security toolkit. They include the following:

- **Firewalls** are primarily used to enforce static security policies.
- **Intrusion detection systems**, while useful for detection of attacks for which signatures are available, alone do not provide scalable, granular mitigation of DoS attacks.
- **Routers** play an important role in the Cisco Guard mitigation process. Access control lists (ACLs) and Remotely Triggered Blackholes (RTBH) are extremely useful but do not typically include a behavior-based feedback mechanism to assist in limiting “collateral damage.”
- **Load balancers** are not designed to combat DDoS application attacks but can be used to help spread heavy loads.

Effective Mitigation Strategy

Dedicated DDoS protection must accomplish four things. First, it must mitigate attacks, not just detect them. Next, it must accurately distinguish between legitimate and malicious traffic, enabling service continuity. Third, it must be deployed in a topologically appropriate manner that allows maximum protection for high-value assets (including other security devices such as firewalls and IDS). Last, it must scale in a predictable, cost-effective manner.

The Cisco Guard XT and Cisco Traffic Anomaly Detector XT interact with Cisco routers to create an effective solution that meets all four requirements. The four-step solution includes *detection*, *diversion*, *analysis and filtering*, and *forwarding* (see figure on page 28).

Detection

The Cisco Guard watches and learns normal traffic patterns, then dynamically creates policies and thresholds based on the observed behavior. The Detector watches for DDoS activity using anomaly-based algorithms so that it can identify new types of attacks on day zero. When that activity varies, the Detector alerts the Guard with detailed information about the atypical traffic and its target.

Many Cisco customers already use devices that can detect DDoS attacks, and these devices can also be configured to alert a Cisco Guard. The devices include Cisco IDS appliances, the Cisco Catalyst 6500 IDS Module (IDSM-2), and the Arbor Networks Peakflow service provider anomaly-detection system, which is based on Cisco NetFlow technology. All of these detection systems can be configured to trigger diversion through the Cisco Guard during an attack; network operations personnel can also elect to trigger the Guard manually if needed.

Diversion

Once the Guard has been alerted to a potential attack, it begins the *diversion phase*.

The Guard begins diversion with a Border Gateway Protocol (BGP) announcement to the nearest upstream router. The router sends all traffic destined for the DDoS target to the Guard. Traffic to other destinations continues to nontargeted zones through the network topology and is unaffected by the diversion of traffic destined for the target.

Analysis and Filtering

The Guard analyzes and filters diverted traffic, dropping malicious packets and forwarding legitimate ones. To accomplish this, the Guard uses a unique, patent-pending technology called the Multi-Verification Process (MVP) to “scrub” flows. This purification process has five modules:

- **Packet filtering**—both static and dynamic DDoS filters block nonessential traffic from reaching the victim. Static filters, which are user-configurable, ship with preset default values. Dynamic filters are inserted by other modules based on observed behavior and detailed flow analysis, delivering real-time updates that either increase verification levels or block identified malicious sources and flows.

- **Active verification**—verifies the legitimacy of packets entering the system and eliminates the risk of discarding valid packets. However, advanced DDoS

Continued on page 31

Protecting the Little Guys: Long-Diversion Method

While it might be appropriate for service providers to deploy a single Guard for each large enterprise customer with multigigabit access links, it is not cost-effective to deploy many Guards near low-speed links to smaller volume customers. Service providers can efficiently protect these customers using the Long-Diversion method, in which a single Guard is deployed at a central network location, with Detectors near edge links into customer premises. Attack traffic identified by the Detectors at the edge is “long-diverted” from multiple BGP peering routers to the central Guard, where it is scrubbed and forwarded to its original destination, often through Generic Routing Encapsulation (GRE) tunnels or other topologically appropriate reinjection methods.

Some service providers already use the Cisco Guard and Detector to offer managed DDoS protection services. Equipped with a Guard, these providers no longer need to shut down service to one targeted customer to protect everyone else on the network; instead, they can preserve service-level agreements (SLAs) to protect both their own and their customers’ revenues and business continuity.

Rackspace Managed Hosting is a managed hosting provider headquartered in San Antonio, Texas. With a commitment to “fanatical support,” it did not want to tell customers suffering from DDoS attacks that it could not help them. As a beta tester and Cisco reference customer for the Guard, Rackspace was among the first to offer managed DDoS services. Through its PreventTier offering, Rackspace provides dedicated, subscription, and ad-hoc DDoS mitigation services to meet the various requirements of its 5600 customers. The Guard automatically mitigates about 80 percent of its daily DDoS attacks, and with Rackspace expert management, it easily conquers the other, more creative assaults.

“The nice thing about the Guard is that it doesn’t sit in the critical path,” says Paul Froutan, vice president of engineering at Rackspace. “It doesn’t add a point of failure to our system, and that’s very important to us.”

attacks use legitimate IP source addresses, so this step merely blocks clumsier attacks, then whitelists flows from legitimate addresses and passes them to the anomaly recognition module for further analysis.

- **Anomaly recognition**—monitors traffic not stopped by the static filters or active verification modules and compares it to baseline behavior patterns, looking for deviations from patterns of legitimate sources seen during normal operation. Attack sources and types are identified at this stage, providing guidelines for the *Packet Filtering* module to install dynamic filters to block malicious traffic or performing more detailed analysis.
- **Protocol analysis**—processes suspicious flows identified by the anomaly recognition module, looking for application-specific attacks. It detects misbehaving protocol transactions, including incomplete transactions or errors.
- **Rate limiting**—an optional feature, rate limiting performs per-flow traffic shaping to prevent misbehaving flows from overwhelming the target while more detailed monitoring takes place.

Forwarding

Once the Guard has verified legitimate flows, it forwards them to the target, maintaining service continuity during attacks. This final step differentiates the Cisco Guard from any other DDoS mitigation technology or product.

Scalability and Clustering

The nature of DDoS attacks requires a highly scalable solution that can successfully process massive packet volume. The Cisco Traffic Anomaly Detector XT has two Gigabit Ethernet interfaces, for 2-Gbit/s monitoring at 3 million pps of up to 90 zones simultaneously. The Cisco Guard XT also has two Gigabit Ethernet interfaces, allowing 1-Gbit/s mitigation up to 1 million pps. The Guard can process up to 1.5 million concurrent connections, protecting an average of 15 concurrently attacked zones, depending upon server type and zone size. It can defend against up to 100,000 zombies and deliver legitimate traffic to the target with less than 1-msec latency.

Both devices can be deployed within a day and are manageable through a command-line interface or a Web-based user console.

A pair of Guards is usually sufficient to protect a midsized service provider network or a large enterprise demilitarized zone (DMZ) network (a DMZ allows external Internet users to access public servers, including Web and FTP servers, while maintaining security for the company’s private LAN). Where more capacity is required, organizations can cluster up to eight Guards behind a single Cisco Catalyst 6500 Series Switch, enabling multigigabit protection in very high volume or multiple-target attacks. ■

For a white paper on Defeating DDoS attacks, visit cisco.com/packet/163_5b1.

INNOVATION





ON INNOVATION

A new era dawns in IP networking.

TURNING THE CORNER

The requirements of the IP routing market are rapidly maturing beyond best-effort data networking.

In the many years since the Internet boom began, routers have been hard at work in service provider backbones and enterprise networks, successfully delivering packets to their destinations. Most of the Web-based data applications in common use—e-mail and file sharing, for example—have tolerated moderate levels of packet loss, latency, and jitter with minimal impact on end users.

Over time, routers have advanced incrementally to support far greater levels of network availability and quality of service (QoS).

Great Expectations

As in any industry, however, expectations only continue to rise. In addition, new applications for IP networks keep emerging—and some of these applications are far more finicky about network performance than e-mail. Consider, for example, the strict latency and jitter sensitivities inherent in real-time IP voice subscriber services and wholesale voice backhaul applications. Then there are forthcoming IP virtual private networks (VPNs) with requirements for end-to-end “committed information rates” and the tricky multicast and QoS requirements of video-on-demand service delivery.

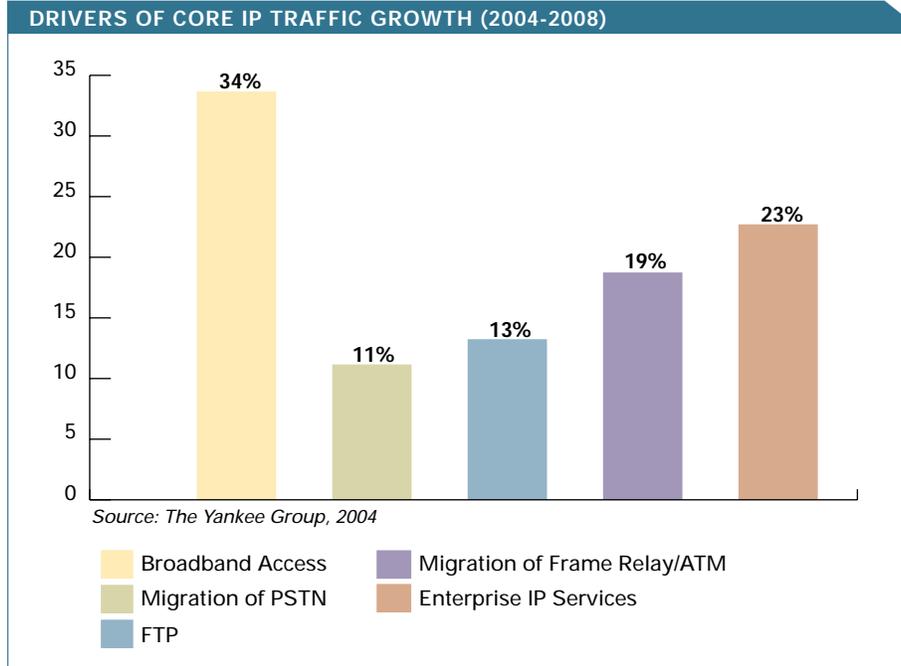
These services represent only a tip of the IP iceberg. The demands of service providers, enterprises, and consumers—and the sophistication of new applications—have reached a point where it has become necessary for the IP routing industry to begin turning a corner on architectural innovation.

To meet scalability and performance expectations in the coming years, owners of IP routing infrastructures will soon need a more available, scalable, and flexible services environment that can deliver on the true vision of network convergence. This vision—one built on converged IP Multiprotocol Label Switching (IP/MPLS) packet infrastructures and able to consolidate the many communications services that today still require separate networks—will be constructed using routing systems with fundamentally different architectures than those that have served the industry well in the past. These new routing systems will be capable of delivering multi-decade scalability, continuous system availability, and unprecedented service flexibility. They will help to alleviate much of the management complexity and costs associated with growing service provider points of presence (POPs) to add capacity for new services and subscribers.

Winds of Change

Why is the industry ripe for change now?

First, service providers would like their IP networks to begin yielding higher revenues. One way to achieve this goal is to deploy new services for which they can charge premium prices. At this juncture, the fees that carriers are able to charge for best-effort data networking services are declining rapidly in a commodity market. Being able to combine traditional best-effort services and “premium” services (those with strict guarantees for



IP TRAFFIC EVERYWHERE Core network traffic is set to explode, driven largely by network consolidation, broadband services, and enterprise IP services.

bandwidth, latency, jitter, and packet loss) onto one network requires router architectures that can deliver 99.999 percent (“five nines”) availability or better, scale without disruption, and deliver extensive traffic classification and queuing capabilities using sophisticated high-speed packet processors.

The new Cisco CRS-1 Carrier Routing System provides all of these capabilities with a massively distributed, “service-aware” architecture that enables nondisruptive scaling of interfaces, processors, and capacity. It supports complete partitioning of resources and provides packet forwarding mechanisms that can perform deep-packet inspection at wire speed. This allows it to service traffic with potentially thousands of queues per interface (see article, “Reinventing the Router,” page 41).

“This is a significant differentiator for Cisco,” says Mark Bieberich, program manager in the Communications Network Infrastructure group at the Yankee Group, a Boston, Massachusetts-based networking researching firm.

“The CRS-1 can apply QoS and traffic management for specific services or network functions using its partitioning capabilities,” he observes. “Service providers have begun migrating mission-critical Frame Relay, ATM, and private-line traffic to an IP/MPLS network. As this migration effort progresses, the IP/MPLS network must match service-level agreements [SLAs] for those types of services,” says Bieberich.

	1985	1986	1987	1988
	MEIS Subsystem is first Cisco product to ship	Cisco AGS (Advanced Gateway Server) is first commercial product shipped	Interior Gateway Routing Protocol (IGRP) is developed, the first protocol to permit the building of large internets	Multipoint Communications Interface ships, the industry's highest-speed network interface

New Age of IP Networking

Meanwhile, consumers increasingly presume that they can do nearly everything related to communications using the Web, their computing devices, and personal communicators. These tasks have evolved beyond basic text e-mail to bundle voice, still video (camera), video messaging, live chat, online gaming, and any number of other services. The delivery of these services requires new levels of performance—not just pure speed, but also tight control over latency, jitter, and network availability.

Given the explosion in intranet- and Internet-based Web activity, combined with the influx of traffic created by the consolidation of ATM, Frame Relay, private-line, and voice networks, it is easy to conceive how the sheer volume of traffic joining IP/MPLS backbones is skyrocketing (see figure). All this communication is driving the need for routers to gain pure horsepower for scalability and performance. In fact, based on primary research conducted in 2004 with worldwide Tier 1 service providers, the Yankee Group predicts a healthy annual growth rate in IP/MPLS core traffic of 117 percent through 2006.

Eighty-five percent of the world's top 20 revenue-generating service providers already have network-consolidation projects underway, according to Bieberich. "These projects validate that carriers are gaining confidence that router architectures will make networks scalable and flexible enough to meet their multiple-service delivery needs," he says.

What have been missing, according to David Willis, vice president of technology research services at META Group, a networking research firm in Stamford, Connecticut, are the "very high levels of hardware scalability and redundancy that ensure very low failure rates."

What are the innovative developments allowing the industry to forge ahead into this new era of IP networking? They include the following:

- Architectures in devices such as the new Cisco CRS-1 that have been designed to deliver the levels of scalability, availability, and service flexibility required for service providers to build converged packet infrastructures and less complex POP architectures
- Performance in carrier and enterprise router architectures alike designed to scale and to suffer no degradation as additional services are turned on
- Maturing standards for the MPLS suite of control-plane protocols
- QoS advances in router hardware to better enforce prioritization and resource reservation markings signaled by router control planes

Router Reinforcements

Router hardware and software designs are beginning to borrow massively parallel processing and modular process-isolation concepts from the computing and telephony industries. One goal is to enable a given router to deliver the five-nines availability that is expected from public switched telephone network (PSTN) switches.

Historically, it has been possible to design routed networks that can deliver five-nines availability by deploying redundant routers in multiple, complex routing tiers, but such uptime was not consistently available from individual routers, points out Brian Daugherty, product marketing manager for Core and Edge Routing at Cisco. But that is changing with the Cisco CRS-1, he says, because of its "always-on," highly distributed hardware and software architecture, which distributes packet forwarding and control-plane processing in a way that greatly minimizes the effects any hardware or software failure can have on overall system availability.

Cisco IOS® XR—the latest member of the Cisco IOS Software family—has been developed specifically to address the scalability, availability, and flexibility requirements of converged packet infrastructures. Its highly modular nature allows for extremely granular process isolation and distribution, so that critical system processes can be started, stopped, or upgraded individually and even moved automatically to take advantage of processor resources anywhere in a multishelf system. Additionally, notes Daugherty, complex state information used by many system processes can even be maintained across process restarts to allow for hitless upgrades and fault recovery.

States Robert Whiteley, an associate analyst at Forrester Research in Cambridge, Massachusetts: "Cisco has leapfrogged the industry with the CRS-1 to build a product on par with the PSTN."

Whiteley, for example, says he is most impressed with the CRS-1's switch fabric. The router, unlike older architectures in the industry, has a three-stage switch fabric that is upgradable in-service, dynamically self-routed, and well architected for delivering multicast traffic. For example, the router can natively replicate multicast traffic directly within the fabric for up to

1989	1990	1992	
Border Gateway Protocol (BGP) is developed and implemented on Cisco routers	Development of cBus and cBus controller and deployment of FDDI, the first high-speed technology interface; additional Ethernet interfaces with up to six Ethernet ports on a cBus card are developed, enabling high-speed switching	Cisco IGS is the first remote access router introduced AGS+ modular router chassis and the <i>ciscoBus</i> five-slot high-speed backplane are introduced NetCentral network management software introduced	Cisco's first patent, No. 5,088,032, is received for IGRP (Feb.) Cisco Communication Server Family introduced (May)

1 million multicast groups, offloading the need for multicast packet replication from the packet processors.

“By the time a packet reaches the output interface, all the work is done. In the old days, a packet wouldn’t be replicated in the actual switch fabric. Instead, it would reach a line card, then go to the switch fabric, then back again, and so forth. It was inefficient,” Whiteley says.

According to Whiteley, it is difficult to retrofit core router switching fabrics and line cards to handle multicast, which he predicts is going to be very important going forward for applications like video on demand. “Now, the multicast replication process is graceful, and it takes place at wire speed,” he says.

Inklings of Innovation

Among the characteristics of the router architectures that will usher in a new generation of IP networking are the following:

- Massively parallel processing
- Checkpointing of state information
- Process distribution
- Service partitioning
- Fault isolation
- CPU and memory separation among applications
- Multiple logical internal routers within a multichassis device
- Deep-packet inspection of multiple services across thousands of queues at wire speed for QoS

These developments exemplify the innovation that will usher the industry into a new era of communications.

Minimizing Disruptions

Cisco’s Daugherty points out that enabling network operators to scale their POP architectures nondisruptively and to extend the lifespan of equipment in a given POP are also a sign of the times. As traffic volumes explode and the traffic from multiple networks consolidates within a given POP, the past approaches cannot scale—from a cost, reliability, or manageability standpoint.

“Historically, the approach has been to add more routers,” says John Doyle, director of marketing for Core and Edge Routing at Cisco. “But with the performance-sensitive services merging into a given POP, not only do network operators need to be able to scale their networks without service disruption, they also need to alleviate the extra administrative burden that comes with adding more hardware, redundancy, and interconnections.”

This consolidation spills over to enterprise networks as well, both in large sites and small. In branch offices, for example, with limited technical staff, simple high-performance integrated systems will emerge for the same reasons that service provider POPs require simplification (see sidebar, “Enterprise Requirements”).

MPLS Matures

Given that IP was created as a simple and connectionless protocol, MPLS was able to bring some semblance of deterministic performance and behavior to IP by predetermining paths and marking MPLS labels for priority QoS. MPLS Traffic Engineering—preselecting paths through a network based on performance or other administrative criteria—is yet another application of MPLS.

History has demonstrated that vision can sometimes lag implementation, given the realities of the standards process and interoperability testing. So while the industry has been making strides with MPLS for many years, the key standards needed to kick MPLS into full action have recently solidified, rendering the control-plane protocol suite finally ready for prime time on a large scale.

Some of these include Internet Engineering Task Force (IETF) standards for Layer 2 tunneling and interworking through MPLS. This means that the legacy Layer 2 subscriber services that have for so long generated handsome revenues for carriers—namely, Frame Relay and ATM—can now all be harmoniously converged alongside newer IP services in an IP/MPLS backbone. The standards for these capabilities—including tunneling between either like or dissimilar endpoints (for example, Frame Relay to Frame Relay or Frame Relay to Ethernet through an IP/MPLS backbone) are now in place.

To further ease service provisioning and management in converged IP/MPLS networks, operations, administration, and maintenance (OAM) features have finally become available for MPLS-based IP networks. MPLS management tools help service

1992

Cisco 3000 Series low-end router platform launches (Aug.)

CiscoWorks router management software introduced (Sept.)



Cisco 4000 Series modular routers for regional and branch offices unveiled (Sept.)

Three-phase program for ATM interfaces is mapped out (Oct.)

1993

Cisco 7000 Series high-end, multiprotocol router platform redefines high-performance routing (Jan.)

Cisco 2000 Series remote access router platform extends the enterprise network to remote sites (June)

Patent No. 5,274,631 for Computer Network Switching System (Dec.)

Cisco is first multiprotocol router vendor to support national ISDN-1 standard (Dec.)

First ATM interface for a router is developed and implemented on the Cisco 7000 Series

providers guarantee service levels for MPLS-based IP VPN services, for example, independent of subscriber interface, while also fulfilling SLAs for traditional Layer 2 services tunneled through MPLS in a converged-network environment.

Software Toughens Up

META Group's Willis considers the management capabilities inherent in the Cisco CRS-1 IOS XR an industry innovation. He observes that Extensible Markup Language (XML) support in the software enables the CRS-1 to work directly with any existing operations support system (OSS) and to take "more of a systems view than an individual-box view in terms of management."

Overall, "IOS XR turns away from being all things to all people to a purpose-built operating system directly tailored to the needs of carriers," Willis says.

Forrester's Whiteley agrees. "Other router vendors have modularized their software, though not to the same extent," Whiteley says. "Cisco took things a step further, by virtualizing the processes and distributing them to any processing resource across multiple chassis. If you separate BGP and OSPF [routing protocols] within the management plane that connects the two functions, you can much more easily troubleshoot a problem."

He says such a setup is a boon to real-time services, such as voice over IP (VoIP). "Now, carriers have the correct foundation for the reliability they need to offer the real-time and converged services we hear so much about," Whiteley says. "They also have the ability to deeply inspect packets at 40-Gbit/s speeds [the speed of the CRS-1 line cards] for QoS, so they can lay the entire proper framework."

Enterprise Requirements

Router innovation is not reserved solely for the service provider backbone. While Tier 1 carrier core networks have the largest requirements from a pure scalability perspective, real-time application traffic generated by even the smallest networks will commingle with packets in the heart of the largest service provider backbones.

The concepts of being able to turn on additional services without performance degradation or service disruption, the need for five-nines availability, and the goals of minimizing administrative complexity and improving price-performance apply to network operators of all sizes.

With such goals in mind, Cisco data center and branch office routers continue to integrate services, such as many aspects of security technology, voice, and video. Most recently, Cisco enterprise routers gained capabilities to optimize edge routing in sites that are dual-homed, based on best-path performance characteristics at the time of transmission and least-cost routing.

For more on the latest developments in the enterprise routing space based on enhancements to Cisco IOS Software, see "IOS: Routing's Crown Jewel," page 47.

Moving On

The networking industry is making its way from running a circuit-switched telephony network for voice, a Frame Relay/ATM network for business data, and a best-effort IP network for consumers (at a minimum) to one next-generation network that supports all requirements. Convergence of this nature has always been a goal, but getting there has been more of a technical challenge than the industry might have envisioned when the commercial Internet took off, and both service providers and router vendors were challenged to simply "keep up" with demand.

The world's network operators are poised to move off their service-specific infrastructures to converged packet infrastructures based on IP/MPLS to handle the next era of networking. At the end of the day, the sheer volume of traffic and the stringent performance requirements of the applications to be supported by tomorrow's networks no longer allow network operators to continue purchasing isolated hardware devices to scale their networks. Rather, large, very fast routers designed to deliver unprecedented levels of scalability, availability, flexibility, and management ease—while vastly simplifying network architectures—will serve network operators well for at least the next decade. ■

1994

Cisco 2500 Series for small and branch offices introduced (Jan.)

Patent No. 5,280,500, method and apparatus for multilevel encoding for LANs (Jan.)

CiscoFusion internetworking architecture is unveiled (Feb.)

Cisco Catalyst® Switch, the first intelligent switch for client/server workgroups, is introduced (Feb.)

First Cisco ATM switch is shipped (Sept.)

Cisco 7000 Router Family is enhanced with a Silicon Switch Processor that nearly triples the routers' throughput (Sept.)

IP Multicast routing technologies introduced that enable massively scalable distribution of data, voice, and video streams efficiently to millions of users

New interface for Cisco 7000 Series—the fruit of an OEM agreement between IBM and Cisco—represents the first time a multiprotocol router can connect directly to a mainframe ESCON channel

Hot Standby Router Protocol (HSRP) introduced; HSRP overcomes previous limitations that host-based network software imposed on "network convergence"—the ability of the host to adapt to changes in network topology

REINVENTING THE ROUTER

A Peek Under the Hood of the Cisco CRS-1



By Gail Meredith Otteson

Router speeds and feeds will always be critical factors in overall network performance. But to meet the IP industry's next-generation availability and scalability expectations, advances in pure capacity must join innovative architectural designs that address other business and operational issues, as well.

The Cisco CRS-1 core router—the first member of the Cisco Carrier Routing System (CRS) family—is indeed unparalleled in terms of capacity and raw horsepower, able to service millions of customers simultaneously. But at least as important, it raises the industrywide routing bar architecturally by enabling the continuous operation of IP networks.

The smart, innovative engineering behind the Cisco CRS-1 moves the IP services community from best-effort data networking to the fault-tolerant, multiple-service networking service providers have long envisioned, with the feature flexibility and capacity they need to sustain the anticipated growth in IP services over the next decade.

1995

Cisco Catalyst 5000 Series is the first multilayer modular switch to combine switching, routing, and VLAN capabilities (March)

Cisco 7500 Series is first router to have a packet-over-SONET interface (Aug.)

Cisco 750 Series ISDN router is introduced (Nov.)

Patent No. 5,473,599 for a system and protocol for routing data packets from a LAN host through a virtual address that belongs to a group of routers (Dec.)

Fast Ethernet Interface Processor for Cisco 7000 and 7500 series routers is the first Fast Ethernet interface in any IP router

1996

AS5200 is first universal access server family introduced (Jan.)

Patent No. 5,519,704 for Reliable Transport Protocol for internetwork routing (May)

Cisco 7200 Series Router extends high-end capabilities to wider range of network environments (June)

Tag Switching technology, the precursor to Multiprotocol Label Switching (MPLS), is introduced (Sept.)

The new class of router supports an aggregate throughput of 92 Tbit/s on a multi-shelf system, divided into 1152 40-Gbit/s slots, offering a variety of interfaces. The Cisco CRS-1 offers the world's first OC-768c/STM-256c interface on a router.

The Cisco CRS-1 achievement represents a significant advance in routing technology, with more than 50 patents on both hardware and software components. Cisco has invested half a billion dollars in its development, drawing upon its 20 years of routing expertise, lessons learned with the large-scale deployment of routers in service provider and enterprise networks, and close collaboration with its leading service provider customers over the past four years.

The Cisco CRS-1 allows service providers to phase out multiple single-service networks in favor of a single, converged network.

"Service providers cannot continue to operate single-service networks and remain profitable," says Tony Bates, vice president and general manager of engineering at Cisco. "Virtually no one is investing in next-generation circuit switches going forward. Those product lifecycles are ending."

The Cisco vision of a truly converged, high-speed packet infrastructure is one that supports today's data, voice, and video services while also accommodating future growth in capacity and capabilities. Networks built with the Cisco CRS-1 system will offer the flexibility and control that enable future consumer-scale, high-value services such as video on demand and video telephony. Both these services require inexpensive bandwidth to gain traction with consumers; therefore, the next-generation

IP infrastructure must significantly reduce cost per unit of bandwidth through network convergence.

Service providers must also protect their profits through reduced capital and operational expenditures. The capacity of the Cisco CRS-1 system allows service providers to reduce the average number of point-of-presence (POP) elements from hundreds to dozens. Existing Cisco 12000 Series routers can be redeployed from the core to the edge for robust, converged edge services.

"Reducing the number of elements and interconnects in the POP represents substantial cost savings," says Mike Volpi, senior vice president and general manager of the Routing Technology Group at Cisco. "At the same time, with the Cisco CRS-1, we're asking service providers to consolidate many eggs into one basket. So it is critical that Cisco delivers a system that is highly available—not just big and fast."

Hardware Architecture

Developing the Cisco CRS-1 "was the Cisco equivalent of NASA's [US National Aeronautics and Space Administration's] race to the moon in terms of the level of drive, investment, and invention required," says David Tsiang, Distinguished Systems Engineer in the Carrier Core Multiservice Business Unit at Cisco. "We've created a radically different architecture. Pieces of the new technologies will trickle into other Cisco products over time, and eventually every customer will benefit from these innovations."

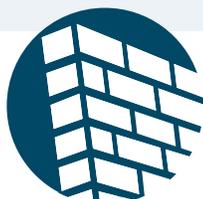
The Cisco CRS-1 architecture draws upon concepts from the computing world, the telephony industry, and lessons learned from previous Cisco product architectures, such as delivering no single point of failure and in-service upgrades. The Cisco 7500 Series Router, for example, proved the concept of distributed processing, which became an inherent design feature of both the Cisco 12000 Series routers and the Cisco CRS-1 platforms that succeeded it.

The single-stage, crossbar switching fabric of the Cisco 12000 Series scales to about 1.28 Tbit/s. Pushing scalability to the next level, Tsiang and his team developed a three-stage, eight-plane switching fabric for the Cisco CRS-1 based on the Benes architecture, a mathematical algorithm originally developed for telephone networks (see Figure 1).

"It's a deterministically nonblocking architecture with connectionless data flows," explains Tsiang. "We achieve the equivalent performance of connection-oriented traffic by randomizing the data paths through the switch fabric. It balances traffic evenly across all data paths."

Like many core routers, the Cisco CRS-1 converts packets into cells for travel across the switching fabric, because packet sizes vary widely according to their application. A TCP ACK is 40 bytes in length, while a data packet may be 1500 bytes or larger. The Cisco CRS-1 uses a cell size of 136 bytes with the ability to pack two packets or portions of a packet in a cell for efficient utilization and performance.

The three-stage switch fabric design can guarantee nonblocking behavior even at the sub-port level. In addition, where some core routers replicate packets at ingress, the Cisco



1997

Patent No. 5,617,417 for ATM communication in inverse multiplexing over multiple communication links (April)

First voice over IP (VoIP) and fax over IP products introduced (Oct.)

Cisco 12000 Series Router for service providers and carriers is introduced, the first completely distributed, modular router with the ability to scale more than 100 times the original capacity (Dec.)

Cable data product line launches (Dec.)

1998

Cisco Catalyst 8500 Series modular campus switch routers announced (April)

Patent No. 5,793,763 for security system for Network Address Translation systems (Aug.)

First industry cable modem for SOHO and telecommuters based on the DOCSIS ITU J.112 standard is introduced (Sept.)

Gigabit Ethernet and Layer 3 routing in switches is introduced (Oct.)

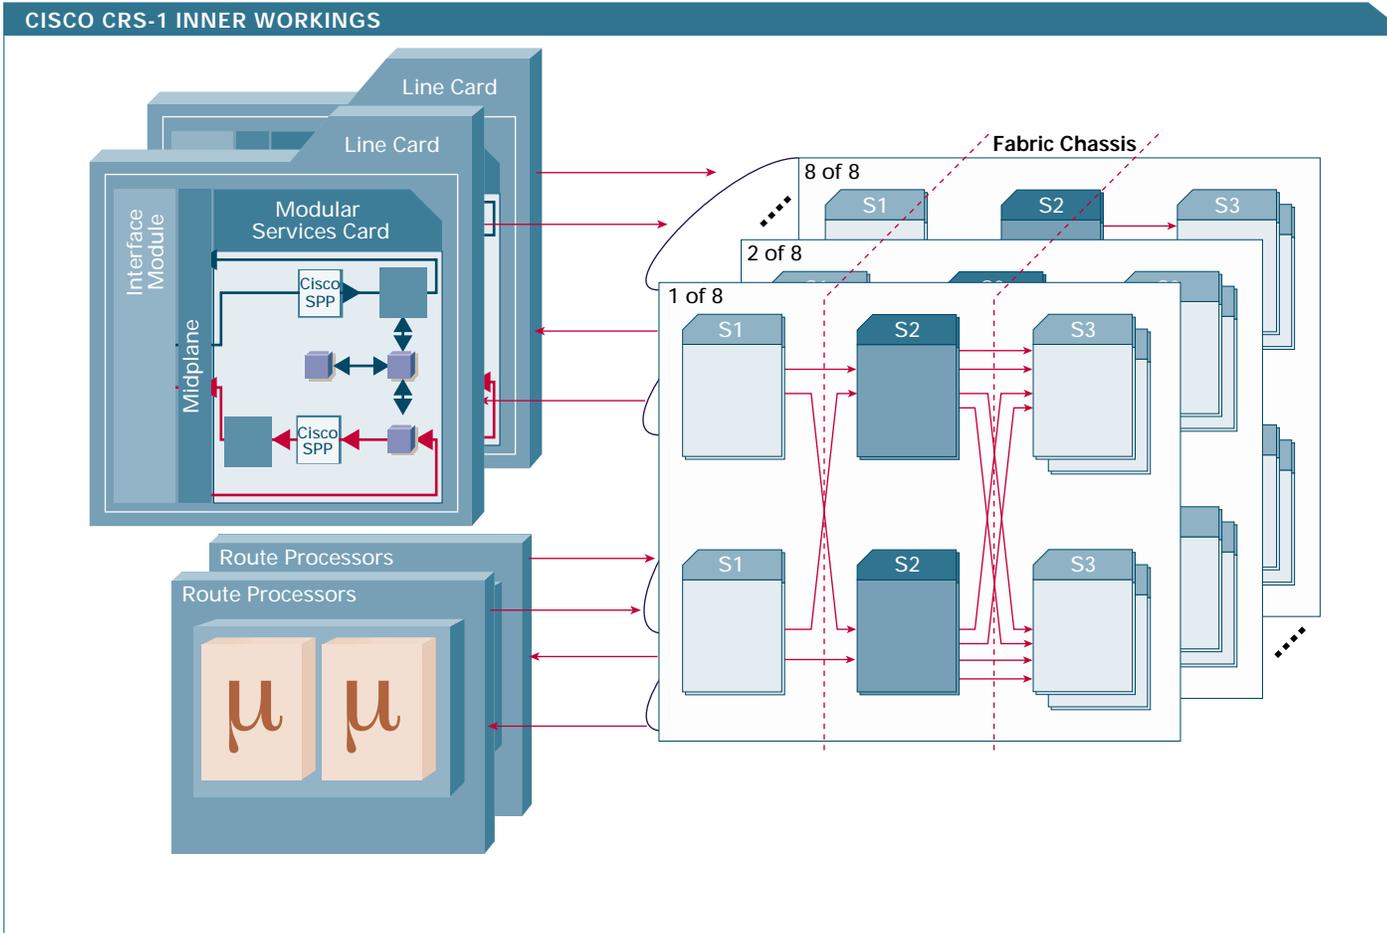


FIGURE 1 The Cisco CRS-1 hardware architecture delivers 40 Gbit/s per line-card slot or 1.28 Tbit/s per single-shelf system and 92 Tbit/s per multishelf system.

CRS-1 system replicates packets in multistage egress. The first stage of the switching fabric directs packets to second stages, where packets are replicated to multiple third stages as required. Packets are replicated on the third stage and forwarded to egress line cards, where they are replicated again before forwarding to egress ports.

A Cisco CRS-1 chassis has 16 slots for line cards or additional route processing cards. A line card has two components: the Modular Services Card (MSC), which performs packet processing, and an Interface Module (see sidebar, "Cisco CRS-1 Interface Modules").

The MSC has a two-stage forwarding architecture with two processors, one dedicated to ingress and the other dedicated to egress. The patented Cisco Silicon Packet Processor (SPP) on the MSC is a 100 percent-programmable ASIC composed of 188 32-bit RISC

1998	1999		
Cisco 800 Series routers for small offices and corporate telecommuters are introduced (Nov.)	Cisco Catalyst 4000 and 6000 series modular gigabit chassis switches are introduced (Jan.)	SONET/SDH, but is optimized to carry IP traffic and applications (Feb.); DPT is now used across Cisco routing platforms	Patent No. 5,883,893 for VoIP technology innovation with a transport layer protocol for compressed voice, fax, and modem data (March)
	New Dynamic Packet Transport (DPT) technology offers the reliability and restorability associated with traditional transport technologies, such as	First vendor to ship a Resilient Packet Ring (RPR) solution using DPT	Cisco 7100 Series of integrated VPN routers launched (May)



processors that operate like the massively parallel processors in supercomputers. Each processor on an SPP operates independently, processing packets completely before forwarding them. Unlike sequential processing architectures, where multiple ASICs partially process packets, this massively parallel architecture is easily programmable and scalable.

Redundant route processors execute routing protocols, system management, accounting, and shelf controller functions with up to 4 GB of DRAM and a 40-GB hard drive for storing logging information and dumps. Service providers can increase system performance with the addition of Distributed Route Processor (DRP) cards that insert into a slot on the chassis. Each DRP card uses dual PowerPC Symmetrical Multiprocessing CPU clusters, double the power of a single route processor.

A standalone configuration supports a single line-card chassis without the need for a fabric chassis. A complete multishelf configuration has up to 72 Cisco CRS-1 line-card chassis and eight Cisco CRS-1 fabric-card chassis.

Sprint Drives the Internet at 40 Gbit/s

Sprint, a global communications provider with more than 26 million customers in over 100 countries, collaborated with Cisco engineers on the design and development of the Cisco CRS-1, including beta testing. In June 2004, Sprint tested the platform with a successful 40-Gbit/s transmission over the live Sprint Internet between the cities of San Jose and Stockton, California, a busy data route.

“We ran the test during ‘rush hour,’” says Oliver Valente, vice president of technology development and chief technology officer at Sprint.

Valente anticipates that a converged, multiservice network will provide greater scalability and more functionality at a lower cost over multiple networks.

“Sprint wants to collapse its many single-function networks into one network that supports multiple services. We believe the Cisco CRS-1 platform will allow us to realize that backbone within two years with fewer moving parts,” he says. “Where we have 100 routers, we can reduce that to 10,” Valente continues. “When we get the code for multichassis [deployment], we expect nothing else will come close in terms of scalability, or probably ever will, since no one else can afford the research and development.”

Valente says he also believes that the platform can support ATM-grade service-level agreements (SLAs).

Cisco IOS XR Software

The Cisco CRS-1 hardware architecture provides a highly scalable, reliable framework, yet the heart of the system is the microkernel-based Cisco IOS® XR Software, which is fully interoperable with Cisco IOS Software on existing platforms or any other standards-based networking platforms. From the ground up, the software architecture was designed to ensure continuous system operation. It also addresses the mathematical complexities of routing through a massive system with memory-protected process operation and exceptional service flexibility.

Cisco CRS-1 Interface Modules

The Cisco CRS-1 offers the following interface modules, delivering 40 Gbit/s to a single line card:

- 1-port OC-768c/STM-256c packet over SONET (POS)
- 4-port OC-192c/STM-64c POS
- 16-port OC-48c/STM-16c POS (due late 2004)
- 4-port 10 Gigabit Ethernet (due late 2004)
- 8-port 10 Gigabit Ethernet (due late 2004)

1999

Next-generation stacking with Cisco Catalyst 3500 Series XL is introduced (May)

Patent No. 5,937,057 for call-center VoIP technology (Aug.)

Cisco 1600 Series becomes the fastest selling router in company history

Cisco AVVID (Architecture for Voice, Video and Integrated Data) for enterprise networks is introduced (Sept.)

Patent No. 5,959,968 for Port Aggregation Protocol (Sept.)

Cisco teams with 10 leading companies to create standards for wireless Internet technology (Oct.)

2000

Parallel Express Forwarding (PXF) Network Processor is introduced

Patent No. 6,101,599 for contextual switching in a parallel processing pipeline array

Ternary Cams (TCAMs), used to support wire-speed, “high touch” packet processing, are introduced; Cisco is the first company to deploy TCAMs in Layer 3 products and has filed more than a dozen patents on the use of TCAMs in packet classification and forwarding

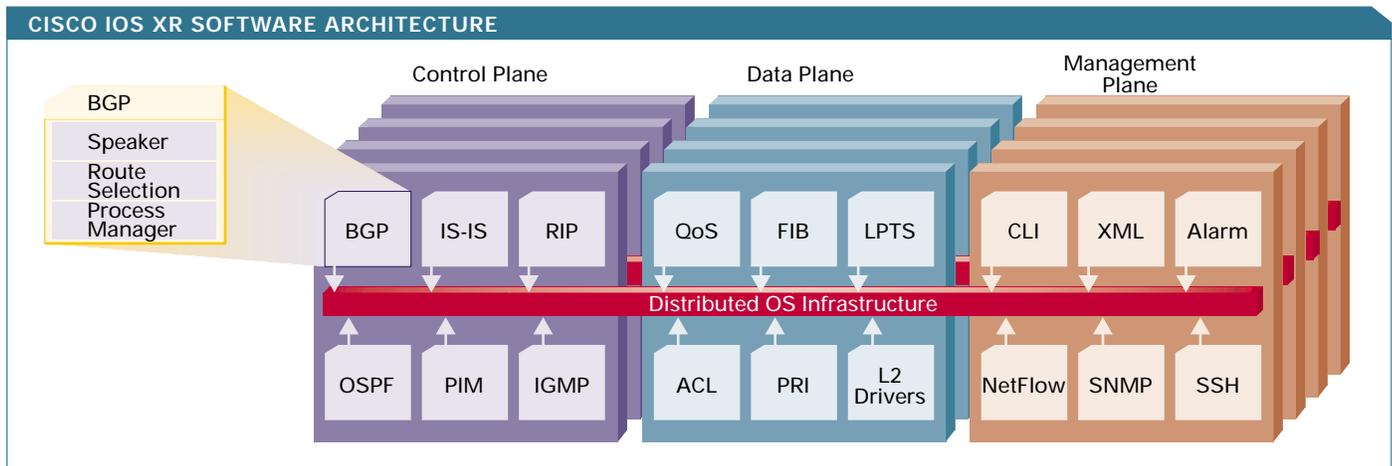


FIGURE 2 Cisco IOS XR Software is a modular, distributed operating system built with a microkernel-based, memory-protected architecture that supports hitless process restarts and in-service software upgrades.

“The asynchronous distributed system was built upon the ideas from GRID computing, cluster computing, parallel processing, and supercomputing,” says David Ward, Distinguished Systems Engineer in the Carrier Core Multiservice Business Unit at Cisco. “Since none of the models completely fit the need of a distributed networking device, all the models were used as the various different applications demanded.”

Cisco IOS XR is modular, adding an entirely new level of reliability to Cisco routing by isolating faults and processes. It has a memory-protected, microkernel architecture and complete separation of control, data, and management planes. Within each plane, operations are organized into smaller objects or threads based on function (see Figure 2). For example, Multiprotocol Label Switching (MPLS) is deployed as a set of modules.

Each thread or module can be distributed to different processing resources—such as quality of service (QoS) into an egress Cisco SPP on the line card and routing protocols on the central route processors.

“There are many CPUs available, each with a 4-GB memory pool,” explains Ward. “This allows us to distribute applications running in the system to each CPU and memory pool, to optimize for scaling and performance. Also, each application is memory-protected for fault tolerance and restartable for high availability.”

Cisco IOS XR also provides a level of physical protection between processes by distributing them inside the system, Ward explains. “You can separate and load balance large-memory applications such as the routing information base [RIB] from smaller-memory applications such as memory agents and other routing and signaling applications.”

For resilience, the microkernel performs only essential processing elements such as message passing, memory protection, and process or thread scheduling. Outside the kernel

2000				
<p>Cisco Catalyst 4006 and inline power are introduced to the midmarket (Jan.)</p> <p>Patent No. 6,018,650 in wireless technology (Jan.)</p> <p>Patent No. 6,044,081 in VoIP technology (March)</p>	<p>Cisco Aironet® 340 Series wireless LAN solution for small and midsized businesses and corporate enterprises is introduced (March)</p> <p>Carrier-class Cisco 10000 Series Edge Services Router is introduced (April)</p>	<p>Patent No. 6,049,533 in wireless/mobility technology (April)</p> <p>Method for integrating hardware encryption technology into Cisco 1700 Series is developed (April); shrinks technology to fit into the size of a PCMCIA card</p>	<p>Cisco Metro 1500 Series MAN DWDM platform is introduced (May)</p> <p>First-ever Internet-transported, digitally screened movie makes motion picture history (June)</p> <p>Patent No. 6,097,718 in IP routing technology (Aug.)</p>	

are elements such as file system, network drivers, and process management. This granular modularity allows service providers to upgrade or restart certain processes without disrupting the entire system. It also allows for a new level of fault tolerance so that drivers and other pieces of the operating system once considered “low level” are now isolated.

The Cisco CRS-1 system can dynamically and automatically restart any software module without disturbing other modules. The system supports three restart modes: cold, warm, and hot.

“Any process can restart in a modular system,” says Ward. “What’s interesting about our resynchronization method is how the system communicates between dependent processes using checkpoint shared-memory databases and hot-standby infrastructure. It’s much faster than relearning through full information exchange between dependent applications. We can achieve very rapid failover between route processors because routing and signaling protocol adjacencies recover very quickly.”

The modular software architecture facilitates in-service, hitless software upgrades. Cisco will release Cisco IOS XR upgrades as packaged sets of modules, with maintenance releases along functional boundaries such as routing, security, and MPLS. Upgrading modules instead of the entire operating system adds stability and is easier to “undo” should unforeseen irregularities occur. It also gives service providers the new ability to run many versions of software in the system, enabling flexibility such as tailoring particular line cards or routing applications for specific applications and services.



FIGURE 3 The Cisco CRS-1 multishef system raises the bar throughout the IP networking industry toward converged networks with continuous operation.

Another service-flexibility feature is the ability to create logical routers within a multishef Cisco CRS-1 system. Logical routing allows service providers to consolidate multiple smaller routers into a single system without losing separation by customer or service, avoiding the need to redesign a POP. Where virtual routers share infrastructure components such as databases, each logical router is built around physically separate route processor modules, with full replication of all applications and databases, along with complete segmentation of the switching fabric. Any route processor can control any subset of line cards across all chassis.

For scalability and performance, routing tables and internal databases are distributed. Each line card only knows its own routes, and each stage of the switching fabric only knows its own forwarding tables. Distributing routing information keeps database size manageable—and scalable—to enable efficiency in a massive system.

Like other high-end Cisco routing and switching systems, the CRS-1 supports both IP Version 4 and IP Version 6 within the same chassis. This meets the emerging needs of early IPv6 adopters, while paving the way toward a global transition over the next decade.

The Cisco CRS-1 has many manageability features such as dashboard alarm management, a high-availability fault manager, a programmable Extensible Markup Language (XML) Web interface, and the ability to create users and groups in an authentication, authorization, and accounting (AAA) system. It ships with an element management system (EMS) based on the successful Cisco Transport Manager, originally designed for managing Cisco optical networking elements. It offers a flexible framework with well-known APIs for integration with existing operation support systems (OSSs).

A New Networking Era

The Cisco CRS-1 is ushering the IP industry into a new era of networking. When the Internet first commercially caught on, the routers in place supported data networking applications that could tolerate small variances in network availability, latency, packet loss, and jitter. But as businesses and consumers place greater reliance on the Internet and other IP networks for all their communications requirements, the new public network must evolve to support those services and the availability and scalability challenges they pose.

The innovation inherent in the Cisco CRS-1 hardware and software builds upon proven routing concepts and borrows from other industries to leapfrog current-generation infrastructures. As a result, the new platform is poised to escort the IP industry into a whole new age of reliable data, voice, and video networking. ■

FURTHER READING

- Cisco CRS-1 Home Page cisco.com/go/crs
- Cisco IOS Software Family and IOS XR cisco.com/packet/163_6b1

2000	2001	2001	2001	2001
Patent No. 6,115,468 for power feed for Ethernet telephones via Ethernet link (Sept.)	Cisco Catalyst 6000 Series Intrusion Detection System integrates IDS functions directly into the switch (Sept.)	Very Short Reach Optics is introduced, enabling cost-effective scaling of IP networks and a significantly lower-cost solution for 10-Gbit/s intraPOP connections	Patent No. 6,173,386 for parallel processor with debug capability (Jan.)	Next-generation Cisco Aironet 350 Series wireless LAN pioneers enterprise deployment of wireless networks (Jan.)
	Patent No. 6,147,996 for IP switching technology innovation (Nov.)		Cisco 12400 Series Router, delivering 10 Gigabit IP/MPLS services, ships (Jan.)	Cisco Smart IAD2400, a family of smart integrated access devices, launches (Jan.)

IOS: ROUTING'S CROWN JEWEL

Industry's reigning router software extends security, performance features for enterprises.

Cisco IOS® Software has long been at the heart of routing innovation in the networking industry. Routing advances embedded in software may not seem as tangible as a shiny new chassis or line card. But many of the rich network services that have distinguished Cisco routers for nearly two decades can be credited directly to the intelligent features built into the routers' operating system.

From the introduction in 1987 of the Cisco Interior Gateway Routing Protocol (IGRP)—the industry innovation that first enabled enterprises to build routed, multiprotocol internetworks—to the recent delivery of Cisco IOS XR for fault-tolerant routing at 92 Tbit/s speeds, Cisco IOS Software continues to evolve with the times.

Alongside Cisco IOS XR—purpose-built for next-generation service provider requirements (see “Reinventing the Router,” page 41)—existing enterprise versions of the software continue to gain rich new features. Many of the latest enterprise enhancements focus on security, an ongoing challenge for network and IT managers as new threats emerge.

For example, Cisco Network Admission Control (NAC), Cisco Public Key Infrastructure (PKI) enhancements, inline intrusion prevention, and embedded Layer 2 and IPv6 firewall filtering features keep viruses, worms, and other intruders at bay in enterprise environments.

2001

Patent No. 6,188,760 for a signal state management system that avoids the overhead of maintaining call state and complex signaling in a packet network gateway (Feb.)

iSCSI protocol for providing ubiquitous access to storage devices over IP networks is proposed by Cisco and IBM (April)

First vendor to announce IPv6 deployment capabilities (May)

Cisco ONS 15540 Extended Services Platform, a high-end metro DWDM optical networking system for enterprise and service provider networks, is introduced (May)

Industry's first 10-Gbit/s Multi-service Transport Platform is launched (June)

Boomerang algorithms and the SODA algorithm for enterprise content delivery networks are developed (Nov.)

First 10 Gigabit Ethernet modules for the Cisco Catalyst 6500 Series are introduced (Nov.)



Other advances improve performance, cost savings, and uptime at the enterprise WAN edge.

A new capability called Cisco Optimized Edge Routing (OER), for instance, automatically selects the best route across multi-homed Internet service provider (ISP) connections according to customer-specified policies. And for improved WAN availability in smaller branch sites without redundant systems, Cisco Warm Reload improves system reboot times, lowering the downtime impact of software faults.

Finally, Cisco has consolidated its routing software into three “trains,” or versions: Cisco IOS T Software for enterprise access and core devices; Cisco IOS S Software for large enterprise core, service provider edge, and service provider core systems; and, going forward, Cisco IOS XR Software for the new Carrier Routing System family of service provider core routers.

NAC Bolsters Endpoint Security

An overriding, continuing security challenge is that networks have increasing numbers of user endpoints (client devices), as the number of remote and traveling users proliferates. Given the dispersion of the workforce, it can be difficult to keep all endpoints up to date on their antivirus signatures and operating system patches.

Compounding the problem are the multiple types of client devices, operating systems, access networks, and services in use by teleworkers and traveling employees that could compromise the corporate network. When off-net, any of these devices might be susceptible to a new infection from the public Internet; then, when they reconnect to the corporate network, they could potentially contaminate it.

“Organizations today use security safeguards such as antivirus software, firewalls, and other endpoint technologies. But at the end of the day, many often still get smacked,” says Russell Rice, a product marketing manager at Cisco.

However, Cisco NAC, at a minimum, enforces one simple rule: If you’re not clean, you don’t get on the network.

A NAC-enabled Cisco router at the enterprise edge automatically checks that settings on client devices attempting to connect to the internal network are current before granting access. This safeguard supports initial blocking and also keeps ill-managed devices from harboring viruses already swept from the rest of the network.

Such tight control, Rice says, is crucial, because “the time it takes for a virus or worm to cause massive amounts of harm has gone from hours to seconds.”

When a noncompliant device is detected, NAC may pursue one of several courses, depending on the enterprise’s policy. It may simply deny access or restrict access to only a server that delivers the most current updates, for example. Network managers might want to set policies for what types of devices have access to which resources based on a number of dynamics, including device type, connection type, and operating system.

Cisco NAC is a Cisco-initiated, industry-wide effort, says Deepak Kini, a Cisco technical marketing manager, in that it supports endpoints running Microsoft Windows NT, XP, and 2000 operating systems, enabling it to check on OS patches. It also works with antivirus software from McAfee, Trend Micro, and Symantec. An open API enables other interested companies to add their systems to the group.

NAC, available in Cisco IOS Software Release 12.3(8)T, is available on all Cisco access routers from the Cisco 800 Series Router to the Cisco 7200 Series Router. The feature is also supported on Cisco network security management and access products such as the Cisco Secure Access Control Server (ACS), Cisco Security Agent, and CiscoWorks Security Information Management Solution (SIMS).

Keeping endpoints clean is a first step, but the dangers don’t stop there. Other threats include tampering, eavesdropping, and man-in-the-middle attacks. Intruders use man-in-the-middle attacks to gain information by intercepting a message, copying or manipulating it, and then retransmitting it to the originally intended receiver. In such cases, it appears that the two original parties are still communicating with each other directly.

Manageable PKI Architecture

Cisco IOS Software offers many features to protect sensitive communications. As networks grow, managing the data-security mechanisms can become cumbersome, so the software includes many innovative PKI features that simplify the provisioning and management of public and private key encryption.

For example, network managers might manually administer shared secret keys for an IPSec VPN with fewer than 100 endpoints. But manual administration and use of shared-secret

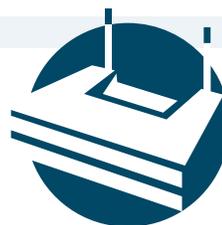
2001

Patent No. 6,314,110 for Method and Apparatus for Distributed Bandwidth Allocation for a Bidirectional Ring Media with Spatial and Local Reuse (Nov.); later known as DPT or RPR

Cisco Long-Reach Ethernet broadband networking solution, the industry’s first end-to-end product line for delivering 5-15 Mbit/s performance over existing Category 1/2/3 wiring, is introduced (Dec.)

Cisco ONS 15808 for delivering data in long-haul optical transport environments is introduced (Dec.)

Complete Optical Multiservice Edge and Transport (COMET) portfolio is introduced (Dec.)



IP Services Engine (ISE) is introduced; this Layer 3 forwarding engine for the Cisco 12000 Series forms the basis for a new line of unique programmable, high-performance, edge-optimized line cards

keys becomes difficult to manage when VPNs grow to hundreds or thousands of endpoints, especially if IPSec network security is applied in a fully meshed configuration.

Cisco's PKI implementation offers a scalable management system for public-private key cryptography. A user's public key is integrated into a digital certificate by the PKI Certification Authority (CA). The digital certificate can be used to vouch for a device's identity and permission to be on the network and provides a trustworthy mechanism for distributing the device's cryptography key material.

"A PKI can serve tens of thousands of users," says Brian Stiff, a technical marketing manager at Cisco. "It offers built-in safeguards such as automatic expiration of certificates at specified intervals and can specify what each certificate holder is entitled to access," he continues. "It's essential for any midsize to large-scale IPSec VPN."

Cisco IOS Software first offered an integrated CA—the foundation of a PKI—in Cisco IOS Software Release 12.3(4)T. One router in the network is established as the CA, and all devices in the network's security infrastructure enroll with it.

Recently, Cisco IOS Software Release 12.3(8)T made certificate deployment simpler with a GUI-based feature called Easy Secure Device Deployment (EzSDD). With EzSDD, remote-site routers are shipped directly to their destinations, where a user removes the device from the box, plugs it in, enters the introduction URL and a username and password. The router automatically registers itself with a security enrollment mechanism and receives a bootstrap configuration to set basic parameters or contact a provisioning system.

Without EzSDD, companies must ship all remote-site routers to a deployment center to have security devices provisioned before shipping the devices to their remote sites.

In addition, Cisco has recently added software features to deal with issues of digital certificate lifetime and revocation. Historically, one concern with PKI has been that the information in certificate revocation lists (CRLs) can become out of date for the period of time between when certificates have been revoked and when the databases storing information about their status have been refreshed, leaving a short-term security gap.

Two features have been added to Cisco IOS Software to close this potential security gap. First, the Cisco PKI-Authentication, Accounting, and Authorization (AAA) feature provides connectivity between certificate validation mechanisms on security devices and AAA servers. In addition, Cisco IOS Software also now supports the IETF-standard Online Certificate Status Protocol (OCSP), resolving the stale-CRL issue by offering a real-time mechanism for certificate status checking from distributed routers.

An OCSP server attaches directly to the CA server that issued the user certificate (and may have revoked it). The OCSP server has instant information as to whether the certificate is valid.

OCSP implemented on Cisco routers directly queries the OCSP server for the certificate's status. The OCSP server checks the certificate revocation database and immediately returns a message indicating the certificate's validity status.

SAMPLING OF NEW IOS FEATURES

Router Software Feature	Release
Network Admission Control (NAC)	12.3(8)T
Inline Intrusion Prevention System	12.3(8)T
Easy Secure Device Deployment (EzSDD)	12.3(8)T
Optimized Edge Routing (OER)	12.3(8)T
Transparent Firewall	12.3(7)T
IPv6 Firewall	12.3(7)T
PKI Certificate Authority (CA)	12.3(4)T
Online Certificate Status Protocol (OCSP)	12.3(2)T
Warm Reload	12.3(2)T
PKI-AAA	12.3(1)

FIGURE 1 Recent enterprise enhancements fall in the areas of security, performance, and high availability.

Other Security Features

Every week there seems to be another virus or worm to be concerned about. Cisco IOS Software is taking additional steps to increase the defense of the network by simplifying the ability to stop these attacks.

Cisco IOS Software's new inline Intrusion Prevention System, available in Cisco IOS Software Release 12.3(8)T, supports a Worm and Signature Attack File containing an initial basic set of 118 signatures. These signatures are from the signature database found on the Cisco.com Website. They have specifically been chosen because they are highly unlikely to generate false positives and because they have the highest likelihood of posing an actual network threat, according to Ruben Rios, a product manager at Cisco.

"Very soon, we will give customers the ability to customize the signatures they download, so that they can run only the signatures that make sense in their own environments. For example, pure Windows environments with no UNIX hosts need not look for UNIX vulnerability attack signatures and vice versa," Rios explains.

2002

Cisco Aironet 1200 Series, the first wireless LAN access point that supports both IEEE 802.11b and 5-GHz 802.11a radios, is introduced (April)

Globally Resilient IP, a set of Cisco IOS Software features that enable networkwide resilience, is unveiled (May)

Industry's first single-port, 10 Gigabit Ethernet router line card for scaling service provider Ethernet infrastructures is introduced, along with industry's first 10 Gigabit DPT/RPR line cards (June)

Cisco ONS 15600 Series Multi-service Switching Platform for

carriers is introduced (Sept.)
New Cisco 7200 Series Network Processing Engine (NPE-G1) more than doubles the processing power of Cisco 7200 Series routers (Sept.)

Cisco Catalyst 6500 Content Switching Module, which allows intelligent scaling of

Websites to very large numbers of servers, is introduced (July)
Cisco MDS 9000 Series multi-layer intelligent storage switches introduced (Aug.)

In-house designed ASICs and a Layer 3 switching module for greater control of data, voice, and video networks in Cisco

Catalyst 4000 Series switches are introduced (Sept.)

New Cisco 7200 Series Network Processing Engine (NPE-G1) more than doubles the processing power of Cisco 7200 Series routers (Sept.)

Eventually, access to all the same signatures available to customers of Cisco IDS Sensor appliances will be made available in Cisco IOS Software, Rios says. Users will be able to customize actions they take on a signature and even modify the signature itself in case a deviation of a known signature have might be coming, he explains.

The Cisco IOS Firewall has also gained two new capabilities. A feature called the Transparent Firewall, available in Cisco IOS Software Release 12.3(7)T, can be placed in an existing network without requiring that network to be re-subnetted to accommodate it, a task that is tedious and resource-intensive.

The feature is targeted at securing wireless LAN access nodes. Instead of reconfiguring remote sites to route between VLANs, the wiring closet switch to which wireless access points connect links to the router with the Transparent Firewall installed via a trunk port, which can be remotely configured. The firewall can be used to control user access permissions to corporate network resources—another line of defense in potentially keeping rogue users off the network.

The Cisco IOS Firewall is the only firewall able to concurrently conduct stateful filtering at both Layer 2 and Layer 3. The same release also makes available the IPv6 Firewall, which can statefully inspect IPv4 and IPv6 messages simultaneously.

Choosing the Best WAN Route

Improved security is one major vector of Cisco IOS Software innovation. Intelligent routing is another, with OER as an excellent example.

THREE TRAINS: THE CISCO IOS SOFTWARE FAMILY

	IOS T	IOS S	IOS XR
Target Networks	<ul style="list-style-type: none"> Access Enterprise Managed CPE/WAN edge 	<ul style="list-style-type: none"> Large enterprise core Service provider edge Service provider core (today) 	<ul style="list-style-type: none"> Service provider core (near term) Service provider edge (future)
Key Attributes	<ul style="list-style-type: none"> Broad platform support/small footprint Integrated security, voice, QoS Flexible feature-option packaging 	<ul style="list-style-type: none"> Enhanced scalability, availability, security Optimized for critical enterprise core and service provider edge networks Broad feature set for enabling flexible service-delivery 	<ul style="list-style-type: none"> Terabit-scale core IP/MPLS routing Unprecedented scalability and performance Continuous system operation Exceptional service flexibility
Target Applications	<ul style="list-style-type: none"> Firewall, intrusion detection IP telephony Wireless networking QoS IPv4 and IPv6 routing 	<ul style="list-style-type: none"> High-end platform support Core and edge IP/MPLS routing MPLS VPNs Any Transport over MPLS (ATOM) Enterprise core infrastructures 	<ul style="list-style-type: none"> Core IP/MPLS routing Large-scale peering POP consolidation Converged infrastructures Continuous system operation

FIGURE 2 The Cisco IOS Software versions, or “trains,” target specific platforms and markets and will continue to draw on proven Cisco IOS Software technologies.

“The route with the fewest hops isn’t necessarily the best-performing route,” says Kathleen Nguyen, a Cisco product manager.

A feature in Cisco IOS Software Release 12.3(8)T, OER is intended for deployment on enterprise border routers connected to multiple ISPs. OER can make real-time routing adjustments based on criteria other than fixed routing metrics. It uses data on traffic characteristics including latency, packet loss, link usage, reachability, and throughput gathered by other Cisco IOS Software capabilities such as NetFlow packet accounting and Service Assurance Agent (SAA) real-time performance monitoring.

Using the same Cisco IOS Software code base, an OER master controller runs on a Linux-based appliance to provide scalability, manageability, data history, enhanced GUI configuration, and reporting.

OER can perform route optimization for specific purposes such as minimizing costs or maintaining QoS for certain application traffic sensitive to metrics such as delay, points out Dan Gill, manager of the Internet Technologies Division technical marketing group at Cisco.

“Say you have a fixed-base and tiered-base cost structure on various links. Depending on performance requirements, OER can route your call to whichever has the lowest-cost link available at that moment,” he explains.

OER also generates reports that help enterprises manage service-level agreements (SLAs) with their ISPs, he adds.



2003

The Cisco SN 5420 Storage Router, the industry's first commercially available iSCSI platform, is launched (Oct.)

Cisco and SURFnet showcase the first transatlantic transmission of HDTV over IPv6 (Nov.)

Cisco StackWise technology in the Catalyst 3750 Series (April) and introduction of the Cisco Catalyst 3550-24 PWR Intelligent Ethernet Switch (Jan.)

The first Layer 2 Tunneling Protocol version 3 (L2TPv3) solution is released (Jan.)

Cisco 7301 Router, the highest performing single-rack-unit router for customer-edge application is introduced (April)

Cisco Wireless IP Phone 7920 based on Wi-Fi and IEEE 802.11b technology is unveiled (April)

Cisco Catalyst 2955 Ethernet Switch for rugged environments and the Catalyst 2940 for classrooms introduced (June)

Structured Wireless-Aware Network Initiative integrates wired and wireless LANs (June)

Higher Uptime at Single-Processor Sites

Some enterprise locations, such as certain branch sites, might not have redundant WAN access routers or routers with dual route processors for hot failover. However, under some circumstances, it might be beneficial or necessary to reload Cisco IOS Software onto these routers. To minimize downtime associated with a reload, there is a feature called Cisco Warm Reload.

The software can reload without reading and decompressing the software from Flash memory. Instead, the image restores the read-write data from a previously saved copy in RAM and restarts execution of Cisco IOS Software. This avoids the lengthy amount of time usually required for a Flash-to-RAM copy and image self-decompression.

The process works because when the router is first booted, Cisco IOS Software saves the initialized Cisco IOS Software data segment before it is changed. When a warm reload is requested, the saved data segment is restored and control is passed to the start of the Cisco IOS Software text segment. Usually, the time savings with Warm Reload, compared with a cold reboot, are between 80 and 90 percent—from minutes to seconds—depending on the router and configuration used.

Cisco IOS Warm Reload first became available in Cisco IOS Software Release 12.3(2)T.

Evolution of IOS

As mentioned, Cisco has simplified Cisco IOS Software by consolidating many versions into three primary ones. The aim is to hone Cisco IOS T, IOS S, and IOS XR software versions to work optimally on three groups of hardware platforms used by enterprise access, enterprise core, and service provider edge and core networks (see Figure 2, opposite page).

Each train contains option sets appropriate to the market it serves, making it less confusing for customers to figure out which version numbers and platforms they must mix and match in order for them to gain the specific capabilities they require.

However, each version “continues to build on the huge shared base of intellectual property and development expertise that has gone into IOS over the years,” says Holly Linden, product technology marketing manager at Cisco. “Our software strategy is to optimize the features and capabilities in each train for each customer group.”

Because of the common Cisco IOS Software base, each feature will operate the same way on any version that includes it, with the same command line interface.

Customers will now find three levels of customization in Cisco IOS Software:

- Infrastructure optimization—the choice of Cisco IOS T, S, or XR Software
- Product-specific optimization—releases specific to given hardware platforms to speed time to market for new hardware, cards, and price/performance enhancements

- Feature packages—streamlined to roughly eight for each hardware platform. “These are like option packages you get with a car; maybe you choose a package with air conditioning, a sunroof, antilock brakes, and automatic windows,” Linden points out.

And some of the Cisco IOS XR Software capabilities will migrate, when appropriate, to the S and T versions of the router software, says Linden.

For example, microkernel-based preemptive multitasking, a capability developed for Cisco IOS XR and the CRS-1 service provider platform, is on the roadmap to migrate to the Cisco IOS Software S train, she says. This capability enables a router running multiple parallel processes to shift capacity among them as needed or in the case of a failure of one process.

In this way and in others, Cisco IOS Software will continue to evolve, as it always has, to meet the continually changing requirements of the many markets it serves. ■

FURTHER READING

- Network Admission Control White Paper cisco.com/packet/163_6c1
- Public Key Infrastructure Data Sheet cisco.com/packet/163_6c2
- Transparent Firewall cisco.com/packet/163_6c3
- Optimized Edge Routing Data Sheet cisco.com/packet/163_6c4

2003

Cisco IP Phone 7970G with color display unveiled (Sept.)

Cisco teams with Network Associates, Symantec, and Trend Micro to address critical industry security issues; outlines Self-Defending Network Initiative to help companies identify, prevent, and adapt to security threats (Nov.)

Enhancements to the Cisco 12000 Series Router double the capacity of the world's largest core networks and provide additional edge service flexibility and extended investment protection (Dec.)

2004

Industry's first carrier routing system—the Cisco CRS-1—is introduced (May)

Guinness Book of World Records certifies the Cisco CRS-1 as the highest capacity router ever developed at 92 trillion bps of total throughput; the CRS-1 is the first networking technology to be recognized by Guinness World Records (July)

A Smarter Way to Network

How an Intelligent, Systems-Based Approach Reduces Complexity While Increasing Functionality

By David Ball



Organizations the world over are turning to technology to be more efficient and productive. Business-critical applications such as customer relationship management (CRM), supply chain management, and enterprise resource planning (ERP) are having a profound impact on the bottom line, allowing companies of all shapes and sizes to adapt in real time to changes in market conditions, and to be more responsive to the needs of customers, partners, and vendors. However, as these applications and services have proliferated, infrastructure complexity and costs have increased exponentially due to the vertical or “siloe” nature in which they typically have been deployed. Many organizations today have hundreds of separate applications and disparate databases, with minimal application integration. This complexity/cost scenario is even worse for service providers, because many have built entirely separate networks for each service they deployed.

Complexity is the IT professional’s constant companion. It appears in many forms and on all fronts. There is the complexity of security as organizations cope with continual and evolving threats from hackers, worms, and viruses. Scalability issues. The rising cost and complication of systems integration and management. Application interoperability issues. Performance and reliability concerns. The list goes on.

These burdens not only raise operating costs, they quickly sap IT’s ability to be a proactive partner of the business, being trapped instead in a discouraging cycle of having to respond to security issues, adds, moves, and changes, network performance degradations or application failures.

While they are challenged to reduce capital expenditures, IT and network professionals understand that business-critical applications and services are only as effective as the network they run on. So, they must continue to make certain that the network is reliable, secure, and accessible to those that need it, regardless of location. To do that, they are continually adding functionality, scale, and resilience to the network, which has the potential of adding further complexity. Worse still, all this added investment offers no guarantee that mission-critical applications and information are making the business more agile and responsive to change.

So, how do you reduce complexity and lower costs, while optimizing the delivery of applications and services that have the greatest impact on the organization? Cisco believes the answer lies in a new approach to the way networks are designed and built—a systems-based approach the company refers to as *intelligent networking*.

Intelligent Networking

To reverse the trend of ever higher customization and operational costs, there is a move in the industry from vertical application deployment to a horizontal approach where a more adaptable and feature-rich network acts as the foundation for a much higher degree of integration between all elements of the infrastructure.

“In the past, network managers have followed a strategy of keeping the network as simple as possible,” says Rob Redford, vice president of product and technology marketing at Cisco. “This strategy was effective when the primary challenges were bandwidth and scale. But now we face more complex challenges: increased application and service integration, better problem diagnosis and fault isolation, and service-level guarantees for the applications and services most critical to the business.”

To meet these challenges, Cisco is embedding intelligent features and capabilities into the network. But intelligent networking also means designing networks more intelligently—thinking first about the business challenges you're trying to solve, and then designing an integrated system that is adaptive and scalable enough to meet future needs. The three elements of intelligent networking are *active participation* by the network in the delivery of the application or service, a *systems* approach to networking, with the network and computing environment working together in an integrated fashion, and *policies* for linking business objectives and processes to network rules.

The Intelligent Approach

The network is the one element of the infrastructure that touches all others—from the middleware and applications, to servers and end users. It is, therefore, a logical place to implement changes that can cost-effectively scale to positively impact the entire organization.

When capabilities are administered at the endpoints (PCs and servers), changes must be made at every distributed node or server, causing management complexity and operational costs to rise exponentially. If, however, these capabilities can reside in the network, where it is easier to make centrally managed changes, they can scale more cost effectively and simplify operations. Through product development as well as partnerships with industry leaders, Cisco is working to identify and implement those capabilities that are best suited to the network arena, and how these capabilities will work in conjunction with other elements of the infrastructure on a systems basis.

"The network infrastructure should no longer be looked at as a passive means of connectivity," says Mario Mazzola, senior vice president and chief development officer at Cisco, "but as an active and integrated participant in the business process."

For networks to become more intelligent, they must be able to make better informed decisions regarding the handling of particular applications or streams of packets. "The network must look deeper into the payload," explains Redford. "It must look beyond the packet headers to understand what type of application it is, and what it is trying to do."

This is a logical evolution. Early routers looked only at the source and destination fields in the IP header to make routing decisions. As the need for more sophisticated route decisions emerged, Cisco added capabilities in software and hardware to look at extended

header fields. Determining quality of service (QoS) for prioritizing IP voice calls is a good example of how this sort of intelligence is used in networks today.

However, the problems of complexity, cost, and application optimization will not be solved by intelligence in the network alone. Different components of intelligence must be present at an integrated systems level (applications, services, middleware, and the network), and be controlled by business policies that set the agenda for the entire infrastructure. In fact, intelligence at the network level is not even possible without a systems-based approach to building the individual products that make up an integrated technology solution.

The Traditional Approach

Until recently, most infrastructure managers were primarily focused on saving money by lowering capital expenditures. They built custom infrastructures from best-of-breed products, trying to reduce expenses by integrating components themselves, and managing complexity to the best of their ability. Today, we're asking much more from our infrastructures. If we continue with the traditional approach, the time and money required to integrate today's larger, more complex infrastructures will likely eliminate any savings on equipment purchases. Standalone products don't necessarily integrate easily, if at all. Support can be an issue because multiple vendors need to be brought in to solve problems. Ultimately, the people and applications that drive the business suffer as issues cause delays. And management asks "where is the promised return on my investments?"

"The pieces are all there, but the current approach has run its course," says Redford. "More sophisticated tools and strategies are required to address the problems of exponentially increasing complexity."

The key is to look at your network and the functions it must perform from both a systems and an end-user perspective—what you need the network to do—rather than starting with boxes and linking them together. "If you forget about designing a network by connecting each component individually, but instead design a system in an integrated way," says Mazzola, "you'll be building a network that will look very different, and have much greater capabilities, with greater intelligence and power."

An intelligently integrated network represents a conceptual shift in terms of how a network can be more than a cost center; it can become an important resource in helping a company optimize its business processes and achieve its business objectives.

Building a Systems-Based Intelligent Infrastructure

To date, the networking industry has focused on solving customer problems one at a time, adding features, capabilities, and intelligence only at the individual product level. As a result, networking equipment manufacturers have developed components that tend to increase management complexities and fail to offer investment protection for existing infrastructure. With each individual product having its own management and feature design, the operation, management, and maintenance of networks have become more complex and expensive.

“The industry is at an inflection point,” explains Redford. “Vendors that continue down the path of focusing only on product-level innovation will increasingly fail to meet the customers’ evolving needs.” This is what inspires Cisco to continue its strategy of delivering innovative, technically advanced products in a systems-level design.

“Everything we are developing today fits into a systems framework,” says Mazzola. Cisco has defined a set of architectural baselines for common features deemed essential in a networked system: security, high availability, QoS, multicast, virtualization, and application optimization. These baselines are, in essence, a set of standards or specifications to ensure functionality and management consistency across a wide variety of products that together form technology solutions. “By increasing the use of common feature sets and hardware functions, we’re making

The Intelligent Approach

Active Participation—intelligent capabilities and features that allow the network to “understand” and actively participate in the delivery of applications and services

Integrated Systems—enables this intelligence across multiple products to create solutions that help to reduce the complexity and cost of your business processes

Policies—adapts this intelligent system to a customer’s specific business processes and objectives, and allows them to extend the benefits to partners and customers

products that are easier to deploy, integrate, use, and manage,” says Mazzola. “This approach utilizes existing open standards such as XML (Extensible Markup Language), as well as new standards we will develop with our partners and industry forums such as the IETF.” Going forward, products developed at Cisco will be evaluated as much for their ability to be deployed and managed as a system, as for their individual capabilities. “For instance, no one is allowed to move forward with a product unless cohesive security functionality is part of the development plan,” explains Mazzola.

Security—Moving Deep into the Network

It is now clear that for security to be effective, it must be pervasive throughout the infrastructure, and this is why Cisco places such emphasis on security as one of its major architectural baselines. As mentioned, the network touches every element of the infrastructure; therefore, it is in a unique position to not only monitor the transfer of information, but to enforce policies in a very coherent way.

For instance, it is nearly impossible to simultaneously control the condition of thousands, or hundreds of thousands, of endpoints. However, the network, designed with the appropriate architectural definitions and intelligence, can automatically check when a new client wants to be connected to the network to see if the new device adheres to the organization’s security policy. As the network increasingly plays the role of a catalyst, it can enable a higher degree of interoperability between PCs, servers, and policy appliances.

An example of this capability today is Cisco Network Admission Control (NAC). Cisco NAC utilizes intelligence in the network to enable an integrated system managed by customer-defined policies to bring the above security scenario to life.

Cisco NAC is part of the company’s security strategy for creating networks that are self-defending. It works with antivirus software on endpoints such as PCs and laptops to ensure that the device’s security status is compliant with local policy before network admission is granted. This is accomplished through agents on the device, in the policy server, and on the router or switch.

“Security vendors tell us that up to 30 percent of their support calls are for threats that have known solutions—problems that have already been solved,” explains Redford. “There’s a high re-infection rate of known worms and viruses because individual users are

not always up to date with the latest antivirus signatures and patches.” Cisco NAC prevents noncompliant devices from accessing the network. Depending upon the policy specified, if the device requesting access has not installed the latest signature patches, Cisco NAC could quarantine the device, routing it to a subnetwork from which it can only access the server to download the required updates. In this way, the reintroduction of known viruses into the network is contained.

Cisco created the NAC program in conjunction with leading antivirus software companies, including Network Associates, Symantec, and Trend Micro. This type of industry collaboration is essential to the creation of intelligent, integrated systems, and is a linchpin in Cisco’s vision of the *Intelligent Information Network*.

The Intelligent Information Network

While aspects of intelligent networking exist in many Cisco solutions today, the Intelligent Information Network is the company’s three-to-five-year vision for building networks that help organizations optimize business processes by enhancing the resilience and adaptability of the infrastructure, integrating new technology—such as IP telephony and wireless—without adding complexity, managing the escalating costs of systems integration, and increasing organizational agility. Cisco is driving this vision for a system that has greater capabilities, intelligence, and power by working with industry leaders, partners, and standards bodies to continually tighten the intelligence integration between the network and computing environments.

In the future, Mazzola sees network components designed with intelligence that interact with your business applications to enhance their performance. For instance, an ASIC on a router could look into packet payloads from an order entry system during an order crunch. In this scenario, the network would understand what’s happening in Layers 4 through 7 and be able to interpret packets carrying XML or Simple Object Access Protocol (SOAP), for example. “If it’s the end of the quarter and your transaction server has orders piling up, the network could inspect packets to look for big orders from important customers with short delivery times,” Mazzola notes.

How will the future be different from today? An Intelligent Information Network will look more deeply at network traffic—not only at packet headers, but deep into the payload to make better informed decisions regarding the handling of individual applications. The network will deliver resilience in terms of service-level performance, not just network per-

formance—for example, ensuring that information gets from source to destination, not just whether a connection from one point to another has a resilient path. Management will operate on a system-wide level, not just on a box or element basis. Assets will be dynamically allocated and managed, enabling resources to be applied to a particular request then freed up for other assignments—essentially converting fixed costs into variable costs.

Vision for the Future Starts Now

For information technology to meet the high-level objectives of the organization, the infrastructure must be designed to enable business decisions and optimize business processes in a way that reduces cost and complexity. Because the network is the foundation of that infrastructure, it must be resilient, integrated, and adaptive. This enables applications, processes, and services to be more effective, which in turn makes organizations and people more productive and profitable.

With intelligence enabled inside the network, organizations have a better understanding of how applications and services operate, allowing networks to make better decisions. This intelligence is not possible without an integrated, systems approach that ties together all aspects of a business solution. Further, integrated systems help to reduce overall complexity and costs by providing faster deployment and usability of services.

Finally, policies based on business rules enable organizations to uniquely adapt this systems-level intelligence for their own infrastructures. By providing flexible policy controls with the application awareness of intelligent networking, infrastructure managers can optimize their networks and implement actions that directly improve business operations.

“This is the first time Cisco has offered such a vision for the future,” Mazzola says. “But we’re already working on the software and the hardware in systems-level implementations.” Elements of intelligent networking are already embedded in solutions the company sells today. Because investment protection is a major component of the intelligent networking approach, customers can add increasing levels of intelligence by building on the foundation of their current Cisco network. The benefits include a secure infrastructure, faster deployment of services and applications, reduced complexity, and lower ownership costs.

For more information, visit cisco.com/go/intelligentnetworking. ■

Turbo-Charged TAC

“Virtual contact center” for Mercedes-Benz USA accelerates auto diagnosis and puts the brakes on telephony costs.



Tim Boyle/Getty Images

STAR SERVICE A new IP communications-based virtual contact center has helped Mercedes-Benz USA technicians strengthen the company's reputation for service behind the “three-pointed star.”

By Rhonda Raider

Mercedes-Benz automobile owners around the world are among the most loyal, and to maintain that loyalty, the company is serious about service. Mercedes-Benz USA (MBUSA), a wholly-owned subsidiary of DaimlerChrysler AG, is responsible for the sales, service, and marketing of Mercedes-Benz products in the US. When a car is brought in to any of the company's more than 325 dealerships in the US, technicians who are stymied can call the company's technical assistance center (TAC) for advice from several dozen of the company's star technicians, who take turns manning the phones on a rotational schedule.

But until recently, the company's four regional call centers—at headquarters in Montvale, New Jersey, as well as in Florida, Illinois, and California—operated as isolated silos, so TAC technicians in different regions couldn't readily share expertise.

“Let's say a dealership in New York took three days to solve a problem,” says Thomas George, lead network analyst for MBUSA. “If a dealership in Los Angeles encountered that same problem a week later, the technician didn't have access to the

information from the New York dealership that could accelerate problem resolution. If we could find a way to share knowledge, we would provide better customer service, improve technician productivity, and reduce dealership costs for loaner cars.”

MBUSA (mbusa.com) achieved this vision by using Cisco IP Contact Center (IPCC) Enterprise Edition to create a “virtual TAC”—an integrated, more efficient Customer Interaction Network—from its four physical call centers.

“Cisco IPCC transformed our isolated regional hubs into a single, unified TAC,” says George.

Creating a Knowledge Base

George was tasked to create the contact center in November 2001. Its threefold goal: to improve service quality, disseminate troubleshooting and repair knowledge to technicians, and collect trending information for MBUSA and the Mercedes-Benz corporate offices in Germany. A centralized TAC would accomplish these goals by creating a unified repository of information and expertise

pertaining to service, support, warranties, and parts, for all dealerships and technicians throughout the US. The resulting knowledge base would also set the stage for business intelligence and analytics. In fact, Mercedes-Benz maintains multiple knowledge bases from its operations around the world, all of which are integrated to the master knowledge base at the global headquarters in Stuttgart, Germany.

“By monitoring repair patterns in the US and around the world, the company would be able to identify problems on the production line before they enter any of the global market showroom floors. MBUSA has no higher priority than assuring quality,” says George.

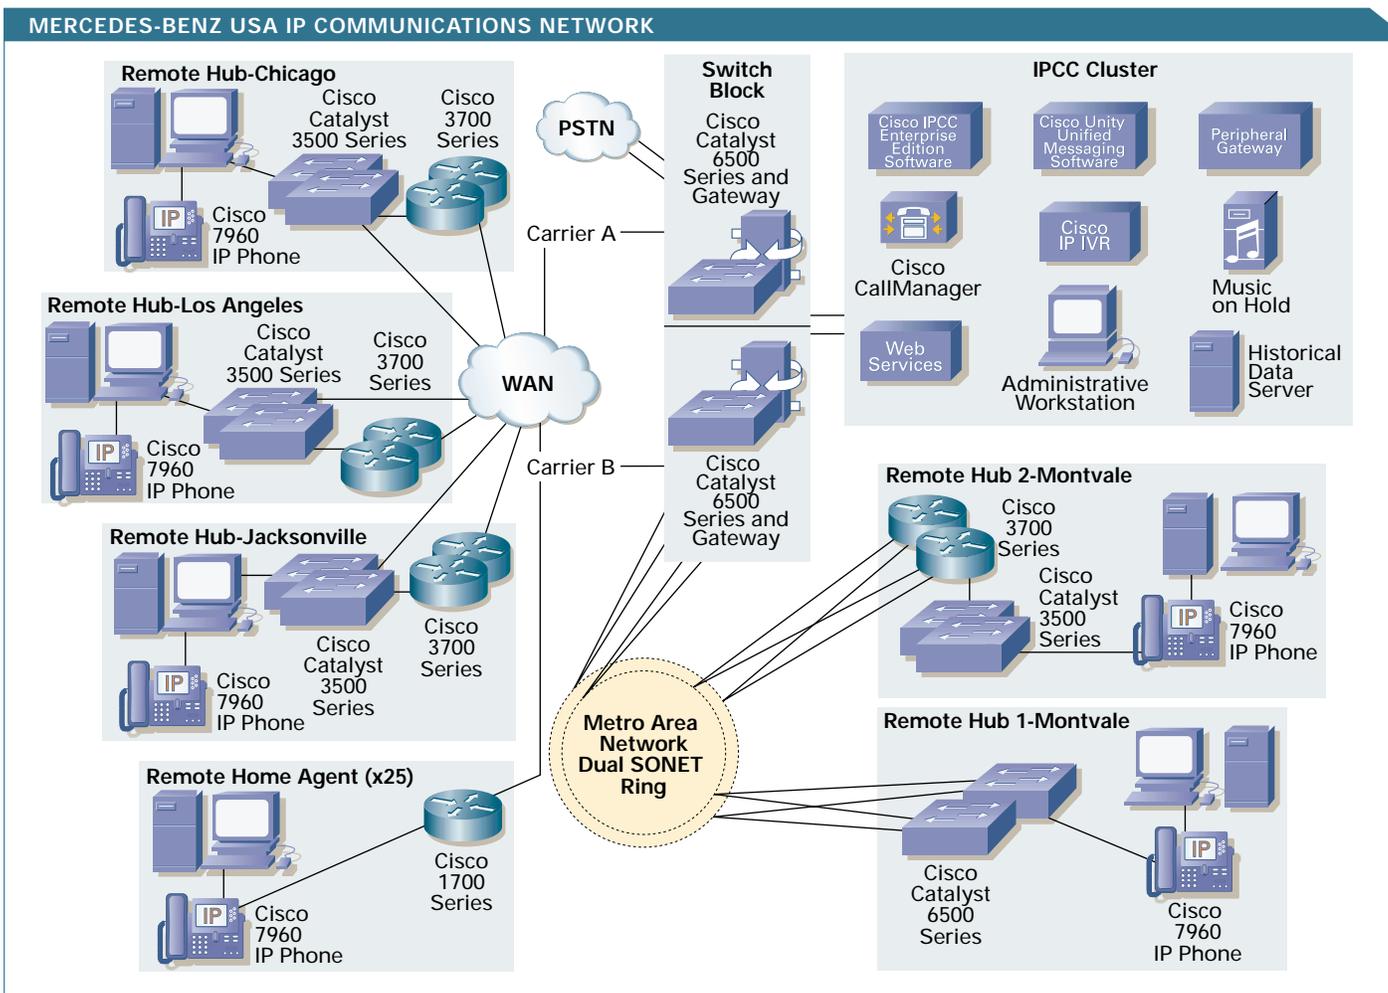
Consolidating Dealership Silos

The chief technological challenge of a centralized call center was creating a common network connecting the more than 325 dealerships nationwide. “Each dealership had its own technical body, but they were working independently until we created this solution

to leverage and share their knowledge bases through a centralized tool,” says George.

A centralized configuration of Cisco IPCC Enterprise Edition overcame the challenge. At the heart of the solution is Cisco CallManager 4.0, deployed on two redundant servers at MBUSA’s Montvale headquarters and connected to the corporate WAN via a Cisco Catalyst® 6500 Series Switch. The servers also run Cisco IPCC for contact routing, Cisco IP Interactive Voice Response (IP IVR) for contact queuing and prompting, and Cisco Unity™ for unified messaging. The three regional offices tap into the same IP telephony and contact center features available at headquarters over the WAN, via redundant Cisco 3745 multiservice access routers. They need no special contact center servers or software—only Cisco IP 7960 IP phones and a Cisco Catalyst 3524-PWR XL Switch to power the phones (see figure).

“Using a single Cisco CallManager deployment to



DIAL “E” FOR EFFICIENCY At the heart of Mercedes-Benz USA’s IPCC solution is Cisco CallManager 4.0, deployed on two redundant servers at the company’s Montvale headquarters and connected to the corporate WAN via a Cisco Catalyst 6500 Series Switch.

serve all four regional hubs is a great savings because it eliminates the need to buy separate telephony switches, which we'd need if we had a TDM [time-division multiplexed] network," says George.

Rapid Deployment

The production call center opened in June 2002, a scant nine months after planning began. All four hubs were operational just three months later. George attributes the rapid deployment to coordination between Cisco and the MBUSA application and network teams. "We spent a substantial amount of time planning the network architecture, call center routing architecture, and business logic," he says. "Mercedes-Benz USA led the effort on the network and call center components, and our systems integrator implemented it."

Turbo-Charged Service

A typical call proceeds like this: Technicians anywhere in the US call a single toll-free number, and are prompted to enter their access code. Cisco IP IVR captures the code and sends it to the Cisco IPCC software, which authenticates the code, and then checks a database to see which dealer hub is closest. "It's the Cisco IPCC routing engine that provides the appearance of a single call center," says George, who created the routing logic in cooperation with engineers in the Cisco Customer Contact Business Unit.

The logic provides "high-touch" service by intelligently routing the call to the technician who's likely to be the most knowledgeable. For instance, the call is routed to the closest hub if an agent there is available—and, if possible, to the agent who previously took a call pertaining to the same car. If someone in the local hub is not available, the call is routed into a national queue, which agents from any of the four hubs can answer. And if no agents are available, the IP IVR queues the call and plays music and technical announcements until an agent is available to answer the call.

George took advantage of the converged voice-data network by using the Cisco CTI Option to integrate the application with the company's internally developed customer relationship management (CRM) software. When an agent receives a call on his Cisco IP Phone, the CRM application pops up a list of cars matching the VIN digits provided. The agent simply selects the correct car to instantly see a vehicle history screen. "If an agent in Montvale wants to collaborate with someone in the Los Angeles hub, they can conduct the call at the same time they use our CRM application to view account history," says George. "This kind of collaboration improves our service edge."

Operational Reporting and Trending

A centralized contact center application also makes it easier for MBUSA to collect and interpret contact center statistics. The company uses Cisco Enterprise

Engineering a Smooth Ride

To ensure that the phone system in hubs remains operational even if the WAN link to its Montvale headquarters becomes unavailable, MBUSA takes advantage of the Survivable Remote Site Telephony (SRST) feature in its Cisco 3745 multiservice access routers at each hub. "In the event of an outage in the WAN link, regional offices can still dial local numbers and emergency services," says George. "If the WAN goes down, call center functionality is somewhat restricted, but the hubs are never without basic telephony services." The network itself is fully redundant and resilient, with redundant WAN lines and backup lines to the routers.

Another tactic the company uses to ensure business continuity is to protect the IP telephony network against threats such as intrusion or viruses, using the built-in integrated security features of its Cisco IP communications solutions. For instance, Cisco Call-Manager 4.0 includes Cisco Security Agent for IP Communications, which detects anomalous behavior that could signal an intrusion and prevents the action from executing. Similarly, the Cisco IPCC system is protected by the Host Intrusion Detection System, which intercepts potential attacks on the server.

Reporting with its Cisco IPCC Enterprise application to monitor and report on agent status, agent states, current calls, and which technicians are available. And managers can use Enterprise Reporting to identify trends that help continually improve automobile quality. MBUSA uses the reporting product to extract TAC call statistics to a data warehouse, where IT staff uses other tools to monitor call center statistics, such as service levels, regional call frequency, repeat occurrences, and productivity.

Home Agent Setups Slash Travel Costs

In the past, TAC agents had to fly to the regional call center during their once-a-month stint. "The trouble with this arrangement was that the heavy travel schedule created job stress," says George. "And airplane and hotel costs for the rotational period were substantial."

To provide superior service to its customers, cut employee travel costs, and relieve technicians of the heavy travel burden, MBUSA now provides a home tech support setup for any TAC technician located more than an hour away from a TAC hub. With a Cisco 1760 Router and a Cisco IP 7960 IP Phone at home, technicians can receive live calls just as if they were in the physical call center. "The cost of leased

lines is minimal compared to monthly travel costs,” says George. “In fact, setting up our technicians as home agents cost roughly 30 percent of what it used to cost to fly them to a hub and put them up in a hotel for one week a month. We achieved ROI [return on investment] in just a few months.”

Toll Cost Avoidance

More cost savings arise from toll bypass. Calls are routed from dealerships to the TAC over the WAN instead of the public switched telephone network (PSTN). A call from Los Angeles to Montvale, for example, travels entirely over MBUSA's IP network, with no cost for the long-distance call.

“The cost of leased lines is minimal compared to monthly travel costs. In fact, setting up our technicians as home agents cost roughly 30 percent of what it used to cost to fly them to a hub and put them up in a hotel for one week a month. We achieved ROI in just a few months.”

—Thomas George, Lead Network Analyst, Mercedes-Benz USA

Capital costs have plummeted, as well. The purchase price of a centralized Cisco CallManager system was just 25 percent of what MBUSA would have spent for a PBX system, according to George. “And the maintenance for the Cisco CallManager solution is about 20 percent of traditional TDM maintenance costs,” he adds. “Throw in the fact that we’re maintaining one IP system instead of four TDM systems, and the savings are significant.”

Foundation for Easy Expansion

The value of the centralized Cisco CallManager configuration is extending beyond the TAC itself, because it allows MBUSA to rapidly add new locations and additional specialized call centers with little or no additional equipment purchase. “A few months after the virtual TAC went live, the home office told us we needed a new building in New Jersey fully equipped for voice in just one month,” says George. “We did it with time to spare, using the CallManager configuration in Montvale to deploy voice over IP to the remote office. Not only did we meet the aggressive time frame, we saved the cost of a TDM switch.”

The company enjoyed similar cost savings after replacing the Centrex system in its Chicago regional training center with IP telephony. “The scalability of Cisco CallManager is where the equity is. We started with four locations and now have nine plus home agents, with the same configuration,” says George.

And still more possibilities for IP communications applications, such as voice over IP (VoIP), are on the horizon. In the works is a new call center dedicated to providing enhanced diagnostics for expensive repairs, such as engine replacements or repairs to navigational systems. And the company is building a new master parts distribution center that leverages the same investment in Cisco IP communications infrastructure. “Instead of buying new switches and PBXs, we’re leveraging our initial investment to deliver more and more services, which will result in continual growth in customer service,” says George.

To provide disaster recovery and ensure business continuity as more critical voice services are delivered over its IP network, MBUSA plans to add a set of redundant Cisco CallManager servers at the Montvale campus, where buildings are interconnected with dual SONET rings.

A Good Call

At its two-year anniversary, the MBUSA TAC has met all the company's original goals of enhancing superior customer service by improving knowledge sharing, cutting travel costs, reducing voice costs, and facilitating real-time and operational reporting. Dealerships have responded very favorably to the new virtual call center. Its technicians answer calls from dealerships every day, strengthening the company's reputation for service.

“Now our dealerships have a venue to more easily reach the most experienced technicians,” says George. “It’s a concrete gesture that corporate is behind the dealerships and committed to their success. Ultimately, our customers benefit from the strength behind the three-pointed star.” ■

FURTHER READING

- Cisco contact center
cisco.com/go/cc
- Cisco IP communications
cisco.com/go/ipc
- Cisco IP communications security
cisco.com/go/ipcsecurity

Routed Radio

Cisco Catalyst module blends wired, wireless networks.

By Gail Meredith Otteson

Large organizations gained the ability to conduct enterprise-class Layer 3 wireless networking with Cisco's recent introduction of the Cisco Catalyst® 6500 Series Wireless LAN Services Module (WLSM). The module integrates wired and wireless network services in very large enterprises. It also enables fast secure inter-subnet roaming, which is particularly important for latency-sensitive applications such as wireless voice.

The module joins enhancements to the CiscoWorks Wireless LAN Solution Engine (WLSE)—which manages and secures the radio-frequency (RF) airspace—to deliver the scalable management, security, and RF control enterprises require to deploy very large, stable wireless networks.

The Cisco Catalyst 6500 WLSM, which requires a Cisco Catalyst Supervisor Engine 720 for operation, and the CiscoWorks WLSE Release 2.7 are the latest enhancements to the Cisco Structured Wireless-Aware Network (SWAN) product portfolio.

Inter-Subnet Roaming

Why is the ability to construct “routed” wireless LANs important?

“As a best practice, enterprise network designs should avoid a single campus-spanning VLAN [virtual LAN],” advises Bob Beliles, manager of product marketing for the Cisco Catalyst 6500 in Cisco's Internet Switching Business Unit. “Flat networks don't scale in extremely large environments, whether they are wired or wireless.”

But wireless LANs have often historically required a flat Layer 2 topology so that mobile users could avoid reauthenticating their Layer 3 credentials when roaming among access points. Reauthentication can take as long as 200 ms, which can cause application sessions to time out.

The Cisco Catalyst 6500 WLSM, however, allows users roaming at Layer 3 to be authenticated by a different Cisco Aironet® access point and a new data path to be established in less than 50 ms. Switchover is undetectable, even during wireless IP phone calls.

Consistent Services Span Wired, Wireless Nets

The Cisco Catalyst 6500 WLSM enables Cisco Catalyst customers to configure and manage user-specific security and quality of service (QoS) policies



ENTERPRISE-CLASS WIRELESS The new Cisco Catalyst 6500 Series WLSM integrates provisioning and policy setting for wireless and wired networks while bringing Layer 3 services and roaming to mobile users.

for both their wireless and wired networks from a single system.

“Network managers want consistent policies for all users, whether they connect through a wire or a radio,” says Beliles. “It's easier to configure policies on a single module than on multiple access points or appliances.”

By separating control and data forwarding planes across the WLSM and Catalyst Supervisor Engine 720, the Layer 3 solution scales to support up to 300 access points and 6000 users. “The WLSM's industry-leading scalability translates into very cost-effective implementation. Other vendors will easily require ten times the number of appliances,” Beliles says.

“Enterprises can also use all the intelligent services available in Cisco Catalyst 6500 services modules and in Cisco IOS® Software on the Supervisor Engine 720 across both portions of their networks,” says Ann Sun, Cisco's senior manager of wireless and mobility marketing.

A Cisco Catalyst 6500 WLSM-centered network requires no design changes. The module slips into an existing Catalyst 6500 chassis located anywhere in the network. Cisco recommends the network distribution layer or data center for optimal placement. But the location is flexible, because the module communicates with access points located using multipoint Generic Routing Encapsulation (mGRE) tunnels. With tunneling, any number of switches can reside between the module and the Cisco Aironet 1100 or 1200 Series access point it controls.

Learn more about WLANs from Cisco experts and your peers at the Networking Professionals Connection “Wireless/Mobility” forum: cisco.com/discuss/wireless.

Banking on Wireless

Sovereign Bank in the eastern US is already using several features in the CiscoWorks WLSE Release 2.7 for managing its RF.

Part of the US\$47 billion financial institution's production network uses Cisco Aironet access points to securely connect laptop PCs distributed across 535 offices throughout the northeast, including New York, New Jersey, and Pennsylvania. As a result, says senior network engineer Todd Dierksheide, the bank relies upon the central RF management capabilities of the CiscoWorks WLSE Release 2.7.

The self-healing network feature helps his team keep banking services available to maintain high customer satisfaction—Sovereign Bank's number one stated priority. Dierksheide says he finds the assisted site survey, AutoConfig, and mass configuration features useful for rapid network turn-ups, such as when his team brought equipment in 50 new branches on line in one week.

And the CiscoWorks WLSE provides much-needed vigilance for detecting suspicious access points. "We find rogue access points all the time," says Dierksheide. "They're usually not on our network, but in our [air] space."

The access point suppression feature is handy in a network where a suspect access point may be several states away, Dierksheide says. Access points geographically close to the rogue radio perform the actual detection. Those access points, deployed in scan-only mode, can disable any switch port to which the foreign access point is connected and report its existence to the centralized CiscoWorks WLSE.

"We chose mGRE because it provides optimal bandwidth efficiencies, lower latency, and greater flexibility than other options," says Ajit Sanzgiri, Cisco Catalyst 6500 WLSM lead designer. "The module can support any Layer 3 protocol and IPv6."

Segregating Users into Mobility Groups

The Cisco Catalyst 6500 WLSM supports wireless mobility groups, defined by service set identifiers (SSIDs) or VLANs, that enable administrators to base users' network access rights on their work profiles. Each mobility group has its own tunnel that transports traffic from the access point to the

WLSM. From there, the Cisco Catalyst 6500 can apply policies to each mobility group for enforcement across the wired network.

"The ability to securely segment the user population is powerful," says Beliles. "Network managers can allocate network resources on a per-group basis with different security policies applied to each."

Users can be sorted into 16 mobility groups. This feature is useful, for example, where a landlord operates a common network for multiple, segregated tenants. Or guests can log into an enterprise network but get routed to public Web services only, with no access to the intranet.

RF Enhancements

Cisco has also bolstered the feature set of CiscoWorks WLSE software. The engine manages wireless-specific services and collects and evaluates radio measurements. It also sends configuration changes through the module to specific access points. Among the CiscoWorks WLSE 2.7 enhancements:

- Rogue access point suppression
- Self-healing capabilities, such that access points adjust their power levels to compensate for performance degradation in nearby radios
- Support for Cisco single- and dual-mode IEEE 802.11a and 802.11g access points
- AutoConfig of new access points, which are automatically added to the CiscoWorks WLSE list of managed devices using Dynamic Host Configuration Protocol (DHCP)
- Single-command mass access point configuration
- Cisco Aironet access points

◆ ◆ ◆

Together, the Cisco Catalyst 6500 WLSM and CiscoWorks WLSE provide better integration of wired and wireless networks plus improvements in security and performance of the airspace. Both are important as organizations attempt to scale their wireless networks ever larger and gain common services and management across both types of networks. ■

FURTHER READING

- Cisco Catalyst 6500 Series WLSM
cisco.com/packet/163_7b1
- CiscoWorks WLSE Software 2.7
cisco.com/packet/163_7b2
- Cisco Structured Wireless-Aware Network
cisco.com/go/swan

Radio Meets Multicast

Multicast VPN replaces satellite, data, and voice networks for lower costs and more creative programming.

By Rhonda Raider

The sonorous voice of the disc jockey takes an intricate journey before arriving at your radio receiver. Consider GWR Group (gwrgroup.com), which has more radio licenses and a larger audience than any other commercial radio broadcaster in the UK. Until recently, GWR managed a medley of three separate satellite networks in addition to leased line and ISDN for its broadcasts, not to mention a Frame Relay data network and time-division multiplexed (TDM) voice network. Now, since migrating to a Multiprotocol Label Switching (MPLS) virtual private network (VPN) with Multicast support GWR has dramatically simplified network management, cut costs to the tune of £1 million annually, and gained unprecedented programming flexibility.

Less Complexity, Less Expense

With 12 million listeners per week, GWR is the largest commercial radio broadcaster in the UK. Its 50 radio services provide both national and local programming and include AM, FM, digital radio, and Internet radio. "Like a lot of companies, GWR grew through acquisitions of smaller groups of stations," says Aidan Hancock, network manager for GWR Group. "We ended up with an unmanageable mess of technologies. The business case for consolidating our networks was very clear: savings of £500,000 a year in capital and operations costs, and even more than that in toll bypass savings."

Why Multicast?

The prospect of multicast had tantalized GWR for some time. "We liked the idea of sourcing a show to all of our stations without having to pay a huge amount for bandwidth," Hancock explains. For instance, the company sources one service in Dunstable that feeds 21 AM stations. With Multicast, bandwidth requirements from Dunstable are equal to one stream alone, rather than 21 streams. This results in far less bandwidth being needed.

Multicast also appealed because of its greater flexibility for network programming—when a show is created one place and then broadcast to others. (Local programming, in contrast, is when a show is broadcast from the same location that creates it.) With GWR's satellite network, shows could be broadcast only from those locations with a physical uplink, while with a Multicast VPN (MVPN), in contrast, a show could be broadcast from any location with IP connectivity. "Multicast liberates the creative side of the business," says Hancock.



BREAKING GROUND Aidan Hancock, network manager, GWR Group, outside Classic FM, one of the UK's top radio stations.

While GWR already employed multicast, its use was restricted to nonbusiness-critical applications such as conferencing and video. "Adopting MVPN for our core business would be a radical advance, both in terms of manageability and broadcast programming flexibility. Migrating to MVPN would reduce bandwidth requirements for national programming, give GWR more flexibility to dynamically change sources and receivers, and replace our disparate broadcast, data, and voice networks with a single converged network. We were motivated: our only alternative would be to continue using our unscalable networks and resign ourselves to the management headaches and lack of flexibility," says Hancock.

Managed MPLS VPN Service

MPLS VPN is typically offered by a service provider as a managed service, so GWR evaluated service providers based on their ability to meet stringent service requirements. One such requirement was quality of service (QoS), which is essential to ensure audio quality. Another requirement was the ability to dynamically change sources and receivers so that a station could creatively assemble a program tailored to its listeners' tastes. "This ruled out a network based on GRE [Generic Routing Encapsulation] tunnels over an MPLS core," says Richard Moir, systems engineering manager for Cisco in Scotland. In addition, GWR wanted sub-ten-second failover times for catastrophic

failures of customer edge, provider edge, or provider routers. "Achieving this service goal would not be an easy task, especially because all unicast topology changes would have to stabilize and routing tables would have to update before the multicast streams could recommence," Moir adds.

Following its evaluation, GWR selected THUS plc, of the UK, as its service provider. "We engaged with THUS at a technical level and involved them in our planning and discussions from the earliest stages, both for Multicast and voice," says Hancock. THUS would build an MPLS core that used the newly introduced MVPN functionality in the Cisco IOS® Software.

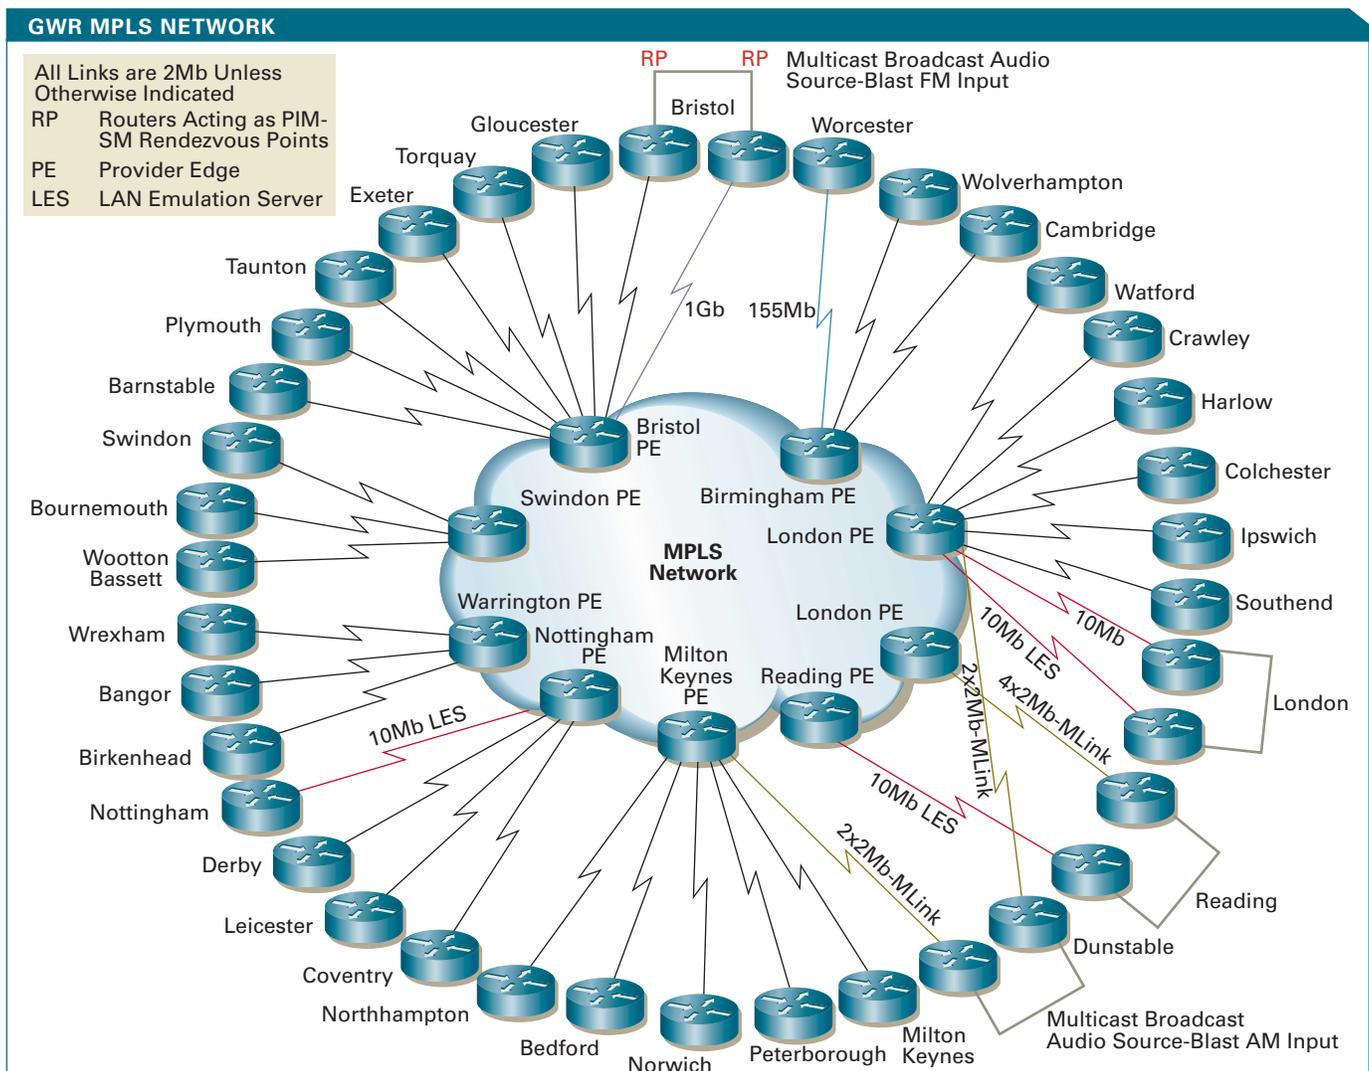
"GWR Group is known for innovation, and its idea to use Multicast VPN for radio broadcast made perfect sense," says Falk Bleyl, MPLS IP VPN product manager for THUS. "Because Cisco offers MVPN

capabilities built into its Cisco IOS Software, we were able to meet our customer's needs by rigorously testing and then deploying multicast capability for GWR's groundbreaking application. Now we've enhanced our managed MPLS VPN service and expect to attract new customers as a result."

The Test

Before going into production, THUS created a test lab in its own facilities that replicated the GWR network. The goal: to ensure the MPLS VPN would deliver the expected network performance and failover and recovery capabilities. GWR provided THUS with its home-grown, Windows-based Multicast application, called Blast (Broadcast Live Audio Streaming Transport), which THUS used to put the simulated network through its paces over the course of several weeks in mid-2003.

Cisco, THUS, and GWR collaborated closely during



BOOST IN FLEXIBILITY Migrating to a MVPN has reduced bandwidth requirements for national programming, giving GWR more flexibility to dynamically change sources and receivers.

VoIP over MPLS VPN

“Because we designed the network to tight specifications for broadcast, we were able to throw in voice as an aside, without any further preparation,” says Hancock. “The network was completely ready from day one.” At publication time, GWR had deployed VoIP at 10 sites and was planning to cut over 30 more in the next few months. “We’re seeing enormous cost savings,” says Hancock. “We no longer need expensive maintenance for every PBX site, even the smallest. By adopting VoIP on a converged network, we’ve reduced complexity and avoided costs for hardware, maintenance contracts, and TDM circuits.” In addition, by aggregating ingress and egress to the PSTN at a couple of core sites, GWR has been able to negotiate better rates from its service carriers. The Bristol site, for example, sends more than one million call minutes per month to the PSTN, resulting in very large annual savings per year in volume discounts.

In addition to cutting costs, VoIP has given GWR the ability to use telephony in more creative ways. For instance, because GWR’s broadcast streams are already available on the network as multicast sources, program controllers and sales managers can monitor output by listening live to any station on any Cisco IP Phone.

the proof of concept. GWR network staff frequently visited the test lab, and Cisco supplied technical expertise, becoming “virtual members of the GWR team,” according to Hancock. As the service provider or GWR noted areas for improving failover in a multicast environment, Cisco responded with successive experimental releases of the Cisco IOS Software. “We were impressed with the very rapid response—sometimes the same day or next day,” says Hancock. “You don’t tend to experience that level of interaction with a vendor.”

The GWR network uses the Open Shortest Path First (OSPF) routing protocol, and on top of that is Protocol-Independent Multicast Sparse Mode (PIM-SM). During testing, GWR and THUS fine-tuned the timers for OSPF, PIM, and BGP within the Cisco IOS Software to achieve the fastest performance while still maintaining stability. “We had to make certain that no combination of failures could take us off-air,” says Hancock. “The ability to cope with failure is critical in our environment, because otherwise the slightest problem might result in silence coming out of millions of radios across the UK.”

The tweaking yielded the desired result: the failure of a single router that’s actively in the forwarding path, either in the backbone, edge or LAN, results in an outage of only two to five seconds. “Had we not tweaked the OSPF timers, the outage could have been as long as 40 or 45 seconds,” Hancock notes, adding that the most common outage, a remote site failing over to ISDN backup, never takes longer than six seconds. “I don’t believe any other software besides the Cisco IOS Software would have given us this kind of control,” he says. “It exposes nearly every parameter so that we can adjust it if we choose.” For instance, GWR and THUS were able to make failover results more predictable by modifying the Reverse Path Forwarding (RPF) check backoff timer to actually take longer than usual, in order to cope with the typical delay of two to four seconds for ISDN to dial.

GWR used end-to-end Cisco gear for the multicast broadcast application as well as for voice over IP (VoIP). At its radio stations, GWR uses 30 Cisco 2651 XM routers, four Cisco 3700 Series routers, and two Cisco 7200 Series routers, each receiving one or more multicast feeds. Cisco 2600 and 3600 series routers are used for backup. They’re connected to the THUS MPLS core using nine Cisco 7206 VXR routers running MVPN. The GWR broadcast and data-center LANs comprise four Cisco Catalyst® 6500 Series switches and two Cisco Catalyst 4500 Series switches. “I’m not certain we could have done it with another manufacturer’s workgroup switches,” says Hancock.

Three-Day Implementation

GWR began using the MPLS VPN for data and voice in October 2003, initially continuing to use its old satellite network for broadcast. Over three days, all GWR sites were cut across to the MPLS VPN network. “At several sites we made the transition without dropping even a single packet,” says Hancock, noting that, with Blast, several packets can be dropped before the loss is audible to listeners.

In December 2003, GWR began using the MPLS VPN for multicast broadcast, in parallel with its old network. Then, in February 2004, with confidence in the network’s performance and reliability, GWR permanently turned off its old collection of networks and began relying entirely on the MPLS VPN.

Unprecedented Flexibility for Radio Programming

MVPN has boosted GWR’s flexibility in station programming by enabling stations to switch between source groups dynamically. “Instead of being restricted to one broadcast source, stations can subscribe to Source A for two hours, Source B for five minutes, and so on,” says Hancock.

GWR also has become more efficient because it can source a programming stream once and then deploy it to multiple partners and multiple services such as

AM, FM, and Internet radio. Previously, GWR had to set up completely separate hardware and connections for each partner or channel. Now, with IP connectivity, the partner or channel subscribes to the sources it needs and there's no need to install new servers.

Stable, Resilient, and Reliable

The single network for data, voice, and broadcast is "stable, resilient, and reliable," according to Hancock. "With our MPLS VPN we now have much finer visibility at the network level. Most of the time we don't lose a single packet at any receivers. On the satellite network, in contrast, we couldn't easily measure packet loss unless it persisted for more than 20 seconds." Hancock notes that THUS is generally beating its guaranteed service level agreements (SLAs) of less than 0.01 percent packet loss on broadcast streams. "We're very pleased and attribute the performance to careful design."

More Nimble Growth

Now able to respond faster and more flexibly to its customers' shifting programming tastes, GWR has gained a competitive advantage over its rivals. It can also more quickly enter new markets. "Our managed MPLS VPN service gives us a jump start to deploy new radio stations more quickly, creatively, and flexibly, whether those stations are digital, analog, or

Internet-based," says Hancock. That's an important edge in the radio industry, which is consolidating to fewer, larger media groups. "It is not uncommon for IT to find out that we've suddenly acquired ten more stations," says Hancock. "Integrating them used to be very labor-intensive because a whole raft of technology goes into every single site, including data, voice, and broadcast links, and backup links for all three. Our MPLS VPN has made it far easier to absorb other stations; in fact, although we haven't had the opportunity to test it yet, we believe the time to integrate a new station will drop from around six months to just one or two."

The company plans to further exploit MVPN by adopting source switching and mixing. Currently, GWR uses distribution-based systems, in which a single source multicasts to tens, hundreds, or thousands of receivers. Next up will be contribution-based systems, in which content from two sources—such as on-air personnel located in different sites—can talk live to the same Multicast group.

Hancock sums it up: "MVPN has created notable competitive advantages in terms of creativity. At the same time our converged network has cut costs to make our bottom line healthier. We really are doing more with less." ■

Virtual Firewall Management

Managing Multiple Security Contexts Using Cisco PIX Device Manager Version 4.0

By David Baum

Today's firewalls do more than just insulate a corporate network from unauthorized access. Besides protecting the perimeter of the network from external threats, firewalls can also prevent users from accessing a particular subnet, workgroup, or LAN within a corporate network.

To achieve this type of segmentation, many customers have implemented structured campus networks that are shared by multiple independent business units. Unfortunately, this sometimes results in completely open networks that have very few security mechanisms to prevent malicious or even unintentional attacks.

Segmenting the campuswide network into *security zones* is a good way to solve the problem. Traditionally, organizations have accomplished this by deploying individual firewalls spread out across the campus network at the distribution layer. As the network grows and the requirements change, however, supporting and maintaining these individual devices can prove to be costly.

Cisco simplifies these configurations with an integrated firewall module designed for Cisco Catalyst® 6500 Series switches and Cisco 7600 Series routers called the Firewall Services Module (FWSM). Now, FWSM version 2.2 enables network administrators to configure multiple virtual firewalls, called *security contexts*, within the same hardware appliance.

A security context is a virtual firewall that has its own security policies and interfaces. When properly configured, security contexts enable the same capabilities as multiple independent firewalls, with fewer management headaches. In essence, these contexts provide completely independent security domains (see Figure 1).

FWSM version 2.2 allows any port on the switch to operate as a firewall port, integrating firewall security inside the network infrastructure. Up to four FWSMs can be installed in a single chassis, providing scalability to 20 Gbit/s per chassis. Network administrators can use this infrastructure to create up to 100 separate security contexts per module (depending on the software license).

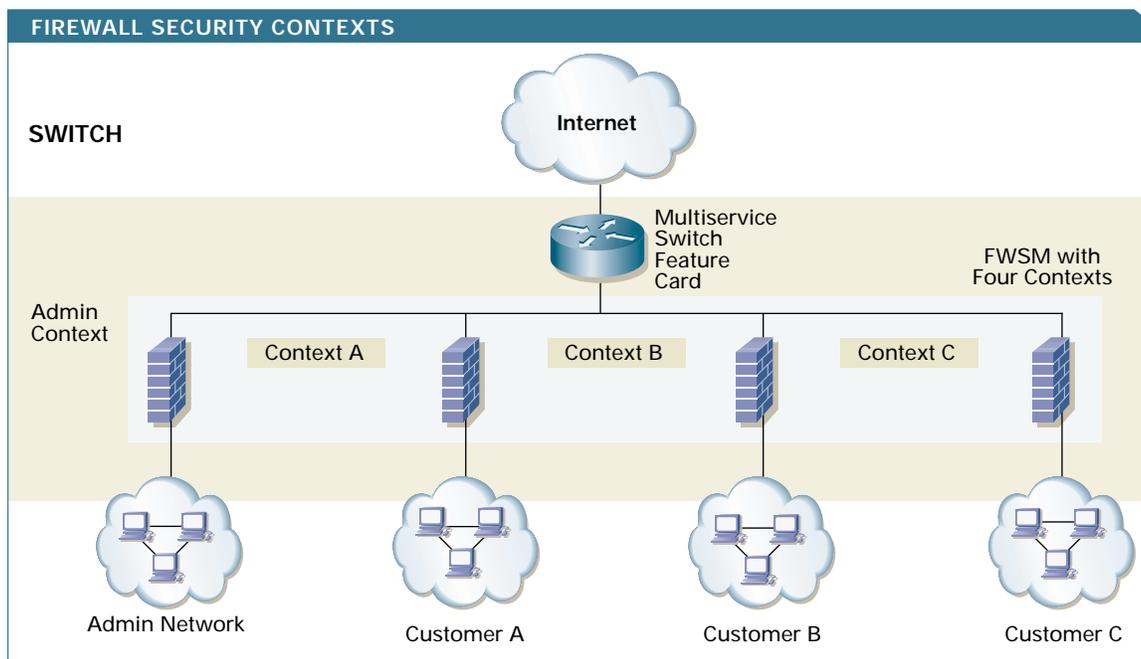


FIGURE 1 Version 2.2 of the Cisco Firewall Services Module (FWSM) enables multicontext firewall deployment within a hardware appliance.



FIGURE 2 Cisco PDM version 4.0 uses an intuitive, highly visual interface to simplify the management and monitoring of virtual firewalls.

Security contexts are functionally similar to a collection of independent physical firewalls but are much easier to manage. Because they are virtual devices, it is easy to add or delete security contexts based on subscriber growth. This reduces management costs, because organizations do not need to deploy multiple devices, yet they can achieve the same capabilities and maintain complete control over the firewall infrastructure from one consolidated platform.

“FWSM provides many of the key firewall and networking features that security managers need to implement multiple security zones or contexts throughout a switched campus network or enterprise data center,” says Iqlas Ottamalika, technical lead, Cisco Security Appliance Group. “Instead of having hundreds of small firewalls spread around the network you can install one hardware platform that will manage everything. This can represent tremendous administrative savings.”

Simplified Configuration and Management

The administrative savings become even more pronounced with the release of version 4.0 of the Cisco PIX® Device Manager (PDM), which includes a Web-based graphical user interface (GUI) and intelligent wizards to simplify the deployment, configuration, and management of security contexts within FWSM 2.2. Cisco PDM 4.0 for FWSM comes bundled with the FWSM operating system, and no hardware upgrade is necessary for customers that already have the FWSM.

“With PDM 4.0, network administrators can configure, deploy, and manage multiple firewalls on one physical hardware platform as if they were separate devices,” says Ottamalika. “PDM for FWSM streamlines the configuration, operation, and monitoring of

security contexts, making it a highly effective productivity tool.”

In testing the newest release of PDM at one of the largest research centers in the US—the Argonne National Laboratory in Argonne, Illinois—Network Manager Corby Schmitz has already seen vast performance improvements. “With past versions of PDM speed and complexity issues made the interface less desirable for us than using the CLI (command line interface). With PDM 4.0 the lag associated with making changes is diminished overall and is almost nonexistent at times,” says Schmitz.

PDM makes it easier to add, create, and delete contexts and their associated policies. It includes real-time and historical reports about firewall usage trends, performance, and security event information. Network administrators can use the GUI to quickly review the status of all firewalls, as well as to gather information on resource utilization and traffic statistics.

Adds Schmitz: “The PDM user interface is very intuitive and allows us to easily change between firewall contexts, gather information, and make changes quickly. The initial trial has caused me to reconsider using the CLI exclusively for management, and now I actually use both interfaces interchangeably to manage our FWSM devices—and we use the PDM exclusively for initial setup of the FWSM.”

By default, the FWSM is available with three contexts: one special administrative context and two customer contexts. Independent management allows different business entities to maintain their own specific policies. Administrators can see all contexts, while customers can just see their own contexts (see Figure 2).

Network administrators log on to a context, which appears as a single firewall. From there, they can change to the system execution space to create contexts and perform global system tasks. Each context can be assigned its own management, syslogging, AAA server, authentication, and URL server policies.

Accurate and Efficient Management

The new GUI in PDM 4.0 is more accurate than using the CLI, because cross-checking is automatically imposed and certain key values are populated automatically.

“With the CLI, you can make unintentional mistakes such as providing a static command but forgetting to include an access list,” says Steven Lee, PDM software development manager in Cisco’s Security Appliance Group.

“The GUI includes a startup wizard to help administrators put the basic elements in place—defining an

interface, applying an IP address, enabling communications services such as HTTP and Telnet, and determining who can access the device. At the command line, you don't get this kind of guidance, which means you can create an incomplete context if you don't initiate the right sequence of commands."

The new version of PDM provides a significant jump in performance and scalability over the previous version and supports all of the new configuration features in FWSM 2.2. Because PDM is a Java applet, it runs as fast as an application and provides live monitoring of the FWSM through the use of graphs and tables.

Types of Security Contexts

Service providers can use PDM 4.0 to set up security contexts for individual customers, while enterprises might use it to assign security contexts to distinct departments or divisions. For example, a service provider deploying a managed firewall service could set up security contexts with distinct policies for each customer.

"For all practical purposes, the impression given to managed-services customers would be that they are using independent physical firewalls, even though they are actually using virtual firewalls residing on the Firewall Services Module," says Ajay Gupta, manager of product marketing for FWSM.

Contexts can be in routed mode (Layer 3) or in transparent mode (Layer 2). Each context can support 256 interfaces in routed mode, or two interfaces in transparent mode. In transparent mode, the FWSM acts like a bridge, enabling network managers to deploy security contexts without changing any of their IP addresses or subnets.

"Cisco's firewall solution now supports multiple contexts for either IP-based [Layer 3] or Ethernet-based [Layer 2] security services," explains Gupta. "The deployment of these multiple contexts provides for full separation of security policy, including authentication, Network Address Translation, stateful access control, and syslog/statistics recording."

Resource Allocation

According to Ottamalika, having multiple contexts is similar to having multiple standalone firewalls. To use them effectively, however, administrators must carefully allocate the FWSM resources.

"Network administrators can limit the resources allocated to any security context at any time, thus ensuring that one security context does not interfere with another," Ottamalika explains. "If you have several contexts, sometimes all the resources will be consumed by the main one, leaving none for the smaller ones. With the Cisco solution, network administrators can allocate resources in an equitable fashion and manage them all through PDM. This

avoids situations in which too many resources are consumed by one context."

Using the Resource Manager in FWSM 2.2, network administrators can allocate firewall resources to individual contexts or classes of contexts: 5.5 Gbit/s of throughput, 1 million connections, 260,000 translations, 100,000 connections per second, 100,000 fix-ups per second, and 27,000 syslogs per second.

For example, if all contexts are to be treated equally, then a single FWSM could be divided into one hundred 50-Mbit/s firewalls using a default resource class. If some firewalls require greater throughput or more resources, then a separate class could be created and limited to a fixed percentage of performance so that the rest of the contexts are not affected.

Summing Up the Benefits

Today's diverse organizations need more than just perimeter security. They need to connect business partners and provide campus security domains that serve multiple groups and constituents. With FWSM 2.2 and the associated PDM 4.0 interface, Cisco is providing the tools network administrators need to deploy multiple security layers across the enterprise network.

"With security contexts, customers can easily segment their structured campus networks into several network-wide trust zones and continue to enjoy the benefits of a distributed network that extends to all parts of the organization," sums up Lee. "Whether adding a new entry or modifying an existing one, PDM 4.0 makes it a lot easier to establish security domains with different policies within the organization." ■

FURTHER READING

- Cisco Security and VPN products
cisco.com/packet/163_4a1
- Cisco Firewall Services Module product information
cisco.com/packet/163_4a2
- Cisco PIX Device Manager product information
cisco.com/packet/163_4a3
- Cisco SAFE Blueprint
cisco.com/go/safe

Wholesale BLISS

Cisco BLISS solution offers carriers unique wholesaler/retailer opportunity.



Riku + Anna

OF MICE AND MEN Allen Sims, chief operating officer at ILEC Smart City, on the grounds at Walt Disney World, Florida. Smart City provides 18,000 phone lines to this famous resort and nearby communities.

By Sam Masud

All employees at Smart City, a Florida-based incumbent local exchange carrier (ILEC), undergo an intensive three-day customer-service training program conducted by the Ritz-Carlton organization. For managers, the rigorous training lasts a full week. As the ILEC that provides 18,000 phone lines to Walt Disney World in Orlando, Florida, as well as to the nearby communities of Celebration and Lake Buena Vista, Smart City's laser focus on the customer is easy to understand.

"If on a day-to-day basis we meet Walt Disney World's expectations of customer service, that speaks volumes about who we are as a phone company," says Allen Sims, chief operating officer at Smart City.

For Sims, the migration to next-generation voice is not only inevitable but a means for enhancing customer service. "We're interested in VoIP [voice over IP] because we see this as a technology that is beginning to sweep across the telecom landscape," says Sims. "We view it as both a threat to our business and as an opportunity to provide enhanced services to customers that we cannot offer with TDM [time-division multiplexing] switching technology."

With VoIP in mind, Smart City is now looking to target small- and mid-sized businesses in the greater Orlando area, but as a competitive local exchange carrier (CLEC). To accomplish this, Smart City is presently considering two options: undertake all the expense and effort entailed in deploying its own VoIP solution or partner with Z-Tel Communications, Inc., a CLEC based in Tampa, Florida.

Such an arrangement would make Z-Tel the first service provider to leverage Cisco's Broadband Local Integrated Services Solution (BLISS) for wholesaling VoIP and high-speed Internet services to a retail partner.

TDM Roots, IP Future

Z-Tel holds the distinction of being the only CLEC serving 49 states and the District of Columbia. As with many other CLECs, Z-Tel has earned its success in the traditional TDM world. The provider uses unbundled network elements (UNEs) from RBOCs to offer local voice services to small and mid-sized businesses along with its long distance service.

Popular RBOC custom calling features such as call waiting, caller ID, and call forwarding are bundled with Z-Tel's proprietary, award-winning Personal Voice Assistant (PVA) technology (referred to as Trinsic Center for the VoIP offering). PVA gives Z-Tel's customers several productivity tools. Among them are Find Me, which allows calls on a line to be forwarded to any three phone numbers, and Notify Me, which alerts users of their voicemail messages via e-mail, text-enabled cell phone, or pager. Users can also dial a number by saying the party's name and store all their contact information in virtual Address Books accessible by phone or the Web.

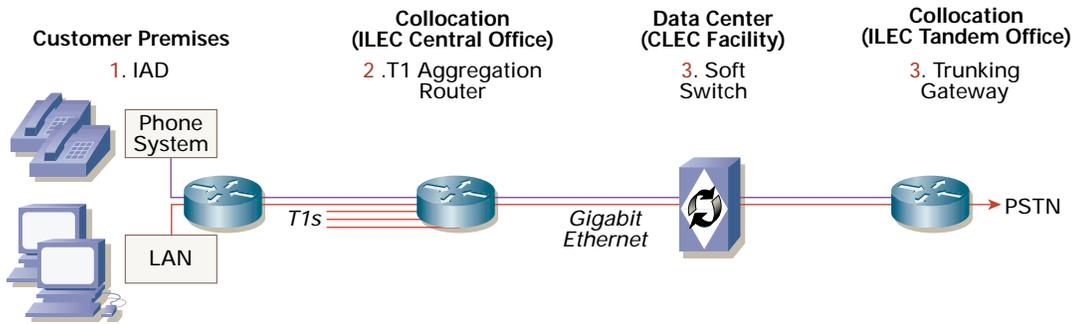
A UNE-based services platform for voice has served Z-Tel well in targeting the low end of the small and mid-sized business market, typically businesses that need five or fewer lines. But Z-Tel also serves much larger businesses with small sites scattered across the US, such as Darden Restaurants, which operates the Olive Garden and Red Lobster restaurant chains.

Recently, however, Z-Tel began offering Cisco BLISS-based voice and Internet services in Tampa, Florida, and Atlanta, Georgia. Z-Tel views BLISS as the right framework for addressing a broader range of customers, those needing five to perhaps as many as 30 lines. With BLISS, Z-Tel can use a single T1 line to the customer for both voice and Internet access, with the bandwidth dynamically allocated to ensure that data traffic does not adversely affect quality of the voice calls. By contrast, Z-Tel's UNE-based customers who want fast Internet access get a separate Internet-only T1 or DSL connection.

The modular architecture of BLISS—currently the only end-to-end solution for delivering affordable, integrated VoIP and data services to small and mid-sized businesses—makes it a very flexible framework for a service provider to extend its coverage as a retailer or as a wholesaler (see sidebar, "BLISS Architecture forms POP-in-a-Rack").

VALUE PROPOSITION: NETWORK SIMPLICITY VERSUS COMPLEXITY

BLISS T1



TDM Network

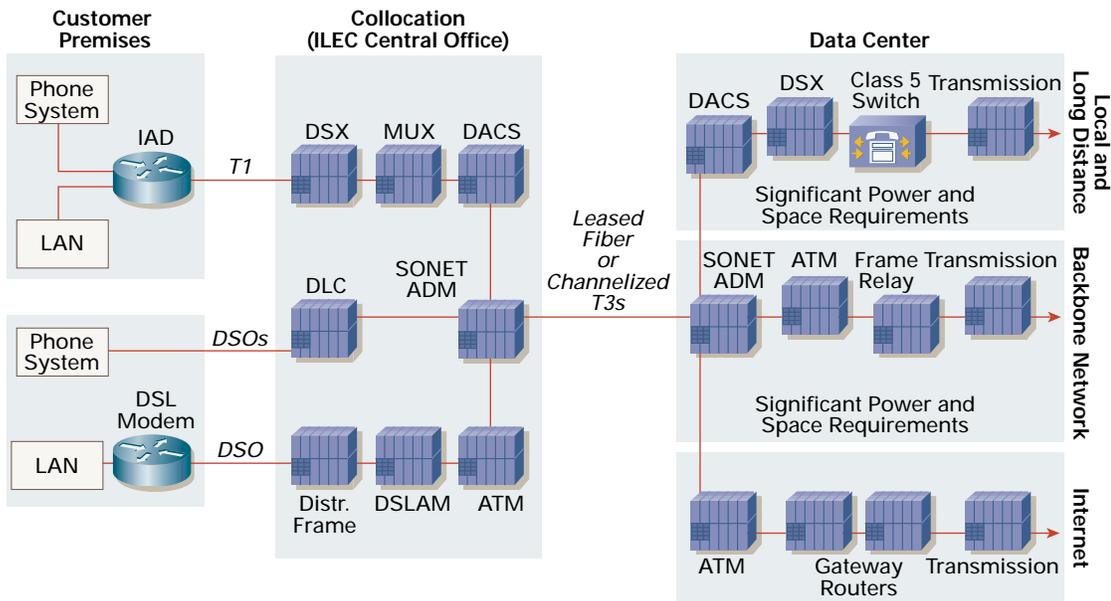


FIGURE 1 Cisco BLISS offers a packet-based, simple, and cost-efficient framework. Compare that to a TDM network, which requires several separate voice and data network components.

Unlike a complex TDM network that requires separate, distinct network elements for voice and data delivery, BLISS supports voice and data services over the same IP connection, a simple framework that translates into lower capital and operations expenditures (see Figure 1).

Retail Scenarios, Revenue Sweet Spots

In a retail arrangement such as the one that could be formed by Smart City and Z-Tel, a key area of negotiation is divvying up how—and by whom—services will be delivered to small and mid-sized businesses. While there are several scenarios that can be worked out, following are three typical retail partner arrangements:

- **Fully outsourced**—Z-Tel does it all to provide the retail partner's customers bundled local, long distance, and data services
- **Shared operation**—Z-Tel integrates the retail partner's local telephony operations with its softswitch, manages the interconnection to the PSTN, or both. In this model, the retail partner owns and maintains the integrated access devices (IADs) and/or the voice gateways.
- **Back office only**—The retail partner is responsible for all network investment but leverages Z-Tel's interconnection and back-office capabilities.

According to Sims, Z-Tel's deployment of BLISS combined with its back-office expertise are strong reasons for entering into a retail partnering arrangement.

"We're already a CPN [Cisco Powered Network] service provider, so going forward [with Z-Tel/BLISS] would simply be a continuation of our relationship with Cisco," says Sims, noting that the high-speed Internet service offered by Smart City in approximately 10,000 hotel rooms in the Orlando area is delivered using all Cisco equipment. "Secondly, we would not have to duplicate efforts because we could use Z-Tel's provisioning, customer service, and billing—their OSS [operation support system]—which is beyond what we have today."

As providers well know, back-office systems at times require more upfront capital than the network equipment for delivering a service. Z-Tel's OSS—used for both UNE- and BLISS-based services—has already proven itself in several large and small retail partnerships. For example, on the UNE side, Z-Tel helped launch and support the first year of MCI's "The Neighborhood," a package of unlimited local and long distance voice service that signed up more than 800,000 subscribers in its first four months. Z-Tel currently also provides local phone service for customers of "Sprint Complete Sense," a competitive package to MCI's The Neighborhood.

Z-Tel is looking to BLISS to increase revenue from wholesale. "On the UNE side, the biggest portion of our revenues comes from consumers, followed by wholesale and then business, which is our fastest grow-

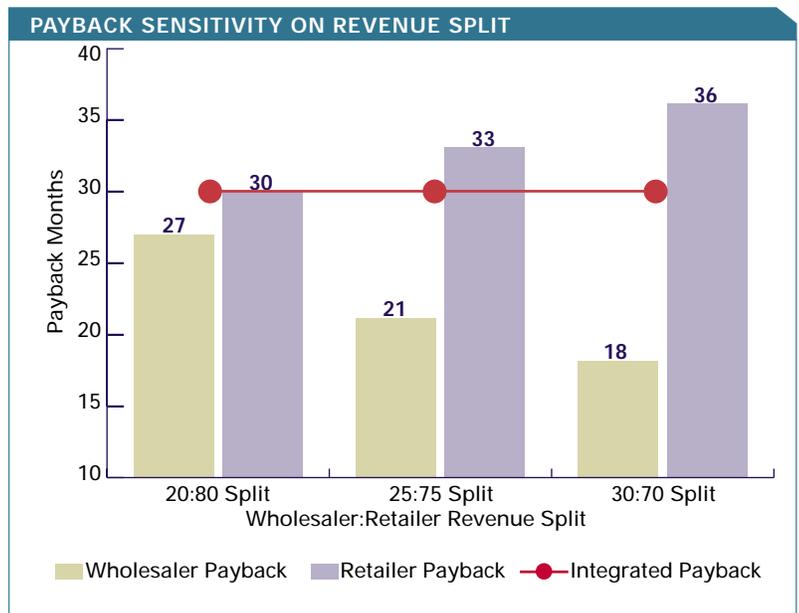


FIGURE 2 The revenue split between wholesaler (service provider) and retailer significantly impacts value allocation between the parties. The best scenario for the wholesaler is 18-month payback at 30:70 revenue split; for the retailer, it is a 30-month payback at 20:80 revenue split.

ing business unit," says Frank Grillo, senior vice president of business services at Z-Tel. "BLISS will provide significant growth opportunities in all revenue areas."

In addition to CLECs, Grillo expects smaller cable operators and particularly second- and third-tier domestic inter-exchange carriers (IXCs) as well as international IXCs—carriers that might be interested in adding VoIP-based local phone service to their long-distance service—to be interested in a retail partnering arrangement.

How quickly the retail partner recovers its investment in BLISS depends on how the wholesaler (service provider) and retail partner agree to split revenues. Cisco estimates that the system payback ranges from about 18 to 36 months, for either the service provider or retail partner (see Figure 2). By teaming with retail partners, Z-Tel would increase its reach faster without having to assume all of the capital costs; the retailer would benefit from accelerating the go-to-market time for integrated voice and data services while reducing operational expenditures.

Numbers Help Tell the Story

Numbers based on real-life service provider experience, combined with Cisco's own estimates, make a compelling case for service providers to retail BLISS to small and mid-sized businesses. Over a five-year period, a service provider can realistically expect to sign up as many as 6000 customers, representing 95,000 subscribers and 52,000 business lines. After a service provider has recovered its investment in BLISS, the payback period for each new customer is estimated to be about nine to 13 months.

BLISS Architecture Forms POP-in-a-Rack

The BLISS distributed architecture with centralized call control makes it an ideal framework for service providers who want to retail and/or wholesale integrated voice and data services, regardless of whether the services are targeted at a single city, a region, or the entire country.

The centerpiece of BLISS is the Cisco BTS 10200 Softswitch, housed in the BLISS-based service provider's data center. It provides the call-control intelligence for establishing, maintaining, routing, and terminating voice calls, and is the interface to enhanced service and application platforms such as a unified messaging server for one-stop access to voice/fax/e-mail or an interactive voice response (IVR) server for automated phone attendant.

Other key product components of the BLISS solution include Cisco 2400 Series IADs for customer premises; Cisco ESR 10000 Series Edge Services Router for T1 aggregation; Cisco MGX 8800 Series Carrier Voice Gateway or Cisco AS5850 gateway access servers for connection to the PSTN; and Cisco Catalyst® switches for switching traffic among different components such as the edge router and voice gateway. In a nutshell, these devices function as a "POP-in-a-rack" residing in an RBOC collocation facility.

In the case of Z-Tel and Smart City, the POP-in-a-rack would reside in the RBOC's collocation sites in Orlando, with local voice service provided by the Cisco BTS 10200 Softswitch in Z-Tel's Tampa data center. The modularity of BLISS eliminates the expense of installing traditional Class 5 switches in relative close proximity to customers.

Based on an acquisition of 6000 customers in five years, following are some additional estimated numbers:

- US\$17.1 million in total capital expenditures
- Earnings Before Interest, Taxes, Depreciation, and Amortization (EBITDA) positive in seven months
- EBITDA margin of up to 54 percent by fifth year
- 101 percent Internal Rate of Return (IRR)
- US\$172 million net present value
- US\$1 in sales generated for every 11 cents in capital investment

"Service providers recognize the need to migrate to VoIP to lower their OPEX [operating expenses]. They also recognize that the challenges to deliver these services go far beyond just changing their infrastructure and require transformations related to their organizations and processes," says Sameer Padhye, vice president of service provider marketing at Cisco. "For some providers, taking advantage of wholesale VoIP opportunity, such as that offered by Z-Tel, will help them accelerate this business transformation with less risk and less capital outlay."

In working with a provider such as Z-Tel, a retail partner would not be limited to offering a one-size-fits-all service. It could provide competitively priced voice and data packages tailored to the needs of individual small and mid-sized businesses. For example, a low-end package might consist of six phone lines, 2000 long distance minutes, and up to 1.5 MB of data for customers with T1 access. A larger business might get as many as 24 phone lines, 6000 long distance minutes, and up to 3 MB of data over two T1 connections. What's more, in addition to a number of custom calling features, customers would also get the capabilities of Z-Tel's PVA at no additional charge.

"PVA is certainly a plus because everybody is looking for something to differentiate themselves from competitors," says Sims.

Although Smart City's current workforce is capable of taking on additional responsibilities, its rapid success as a CLEC could strain the provider's resources. "I hope to be in a position where I begin to outsell my resources quickly and need additional customer-service resources, provisioning resources, or something else along those lines," says Sims. "So, if I'm successful in the marketplace with VoIP, I would either have to add people or use Z-Tel," he says.

"I liken the opportunity [a partnering deal with Z-Tel] to going into franchising. . . . You can take advantage of certain things that somebody else already has done—plowed the ground, hit the bumps, solved the problems. You're learning from their experience," notes Sims, "and, therefore, getting a headstart in the market." ■

FURTHER READING

- BLISS for Cable
cisco.com/packet/163_8a1
- BLISS for T1/E1
cisco.com/packet/163_8a2
- Z-Tel's VoIP Offering (Trinsic)
trinsicnow.com

Taking to the ROADM

Optical technology, ROADM, poised to spur metro DWDM market.

By David Barry

As rapidly increasing bandwidth demands and new types of services such as Gigabit Ethernet, Fibre Channel, and SONET OC-192 affect metro networks, service providers are poised to evolve their metro infrastructures to accommodate these changes. Recent developments in reconfigurable optical add/drop multiplexer (ROADM) technology, which is now available on the Cisco ONS 15454 Multiservice Transport Platform (MSTP), are providing operational simplicity and lower operations costs, spurring service providers to give powerful and flexible dense wavelength-division multiplexing (DWDM) technology a second look for their metro networks.

Although early expectations for metro DWDM were great, broad global deployment of these solutions has not been realized. Traditional DWDM technology has not seen widespread adoption based on two primary factors:

- Expense—especially because of the need to adapt the wavelength to the correct ITU frequency
- Operations inefficiency—DWDM has been difficult to operate and required a significant level of manual installation, measurement, and provisioning

The majority of metro DWDM deployments today are being implemented as storage-area and local-area network extensions (SAN/LAN). However, recent and future growth in this market is likely to be increasingly derived from cable MSOs and service providers building internal infrastructures. A recent Dell'Oro optical report indicates that the worldwide metro DWDM market was US\$575 million in 2003 (based on manufacturers' revenue), and this market is expected to grow 21 percent to US\$698 million in 2004.

The promise of metro DWDM was that it would solve the exploding bandwidth demands created by the rapid adoption of high-speed technologies such as Gigabit Ethernet and Fibre Channel. As SONET/SDH rings became congested and service providers built costly overlay networks with new fiber installations, metro DWDM was intended to deliver the wavelength services for unlimited metro scalability.

But metro DWDM solutions based on fixed optical add/drop multiplexers (OADMs) have not been attractive for mass deployments because they do not

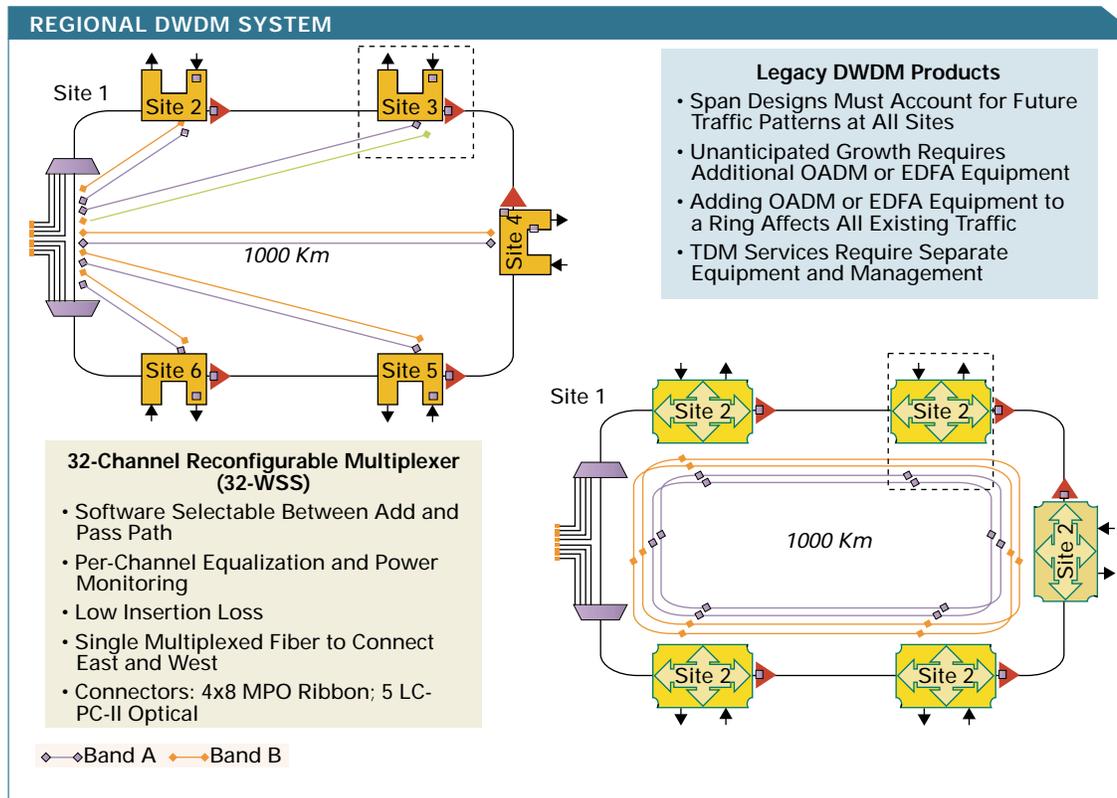
provide the operational simplicity and flexibility service providers have grown to expect with SONET/SDH, leading to high operations costs or simply the unwillingness to take on the operational model. While fully mature SONET/SDH systems allow service providers to easily determine the status of current connections (DS-1, DS-3, OC-3, OC-12, Gigabit Ethernet), evaluate service requests, and provision new connections remotely, these capabilities are highly complex and time consuming in most metro DWDM systems today.

Metro DWDM networks are commonly based on fixed OADMs, in which optical paths are typically predetermined when the system is engineered and unused capacity cannot be reallocated without reengineering the system. To drop a wavelength at a node requires physical visits by design engineers to essentially redesign the network—changing filters, retuning existing optical interfaces, or possibly adding new ones. Other working services are disrupted during this process. In a dynamic metro carrier environment, where customer churn is constant and demand for new service turn-up is constant, the high-touch model of fixed OADMs is not a feasible solution.

The power of ROADMs is that they enable service providers to add, drop, or pass through any combination of the available wavelengths by remote control without having to deploy DWDM experts on site to manually configure or adjust node configurations.

“There is a lot of operational pain when service providers want to add services to the network,” says Ron Johnson, product manager for the Cisco ONS 15454. “Turning up new channels beyond the pre-deployed filters that have been initially deployed causes the network to be re-engineered. This is something that most service providers have chosen not to introduce to their operational models. ROADM stands to alleviate that situation and allow DWDM to be a much more widely used technology.”

For the most part, metro DWDM has seen limited point-to-point applications where one customer will



RECONFIGURABLE NETWORKING ROADMs allow individual wavelengths to be dropped off at sites without demultiplexing an entire band of wavelengths, bringing great operational flexibility. Granular SONET/SDH services can also be accommodated on the Cisco 15454 platform.

purchase an entire DWDM system, and has mostly been used by service providers for niche applications for fiber relief.

The Promise of ROADMs

ROADM technology is an exciting new advance that brings SONET-like manageability to metro DWDM. The power of ROADMs is that they enable service providers to add, drop, or pass through any combination of the available wavelengths by remote control without having to deploy DWDM experts on site to manually configure or adjust node configurations. An important benefit of this remote capability is that it reduces the likelihood of human errors caused when a technician must touch or change the network configuration. Another benefit is that because the ROADM handles any combination of pass, drop, and add there is only one ROADM equipment configuration regardless of the A-to-Z traffic demand for the system. This introduces an entirely new approach to designing and managing networks for ever-changing forecast models that start with modest bandwidth requirements and are projected to provide growing demand and revenue streams.

In a dynamic DWDM transmission system, changes to the multiwavelength optical signal from provisioning or restoration activities affect the integrity of other optical channels. ROADMs include adaptive compensation capabilities that allow the system to

compensate for the power changes caused by adding or dropping wavelengths. In static OADMs, compensation is typically performed in bands across all wavelengths to get the lowest cost. But static DWDM systems are not faced with the added complexity of adapting to changing route lengths and its impact on optical transmission parameters—including optical attenuation and dispersion—and on total system cost. With ROADMs and the dynamic network they help create, vendors are using per-channel variable optical attenuators (VOAs) or dynamic gain equalizers (DGEs) to manage power levels (and signal attenuation) and are creating any-to-any dispersion maps for dispersion control.

Reducing Sparring Requirements

One of the most compelling aspects of ROADMs is that, because of their inherent reconfigurability and use of tunable lasers, fewer card types are required in the network. Without ROADMs, a service provider might need to select from more than 50 different card types—single-channel devices with 32 different wavelengths, for example, or 2-channel cards with 16 different wavelengths. Tracking these cards and providing sparring for each is a major cost.

ROADMs on the Cisco ONS 15454 MSTP are made up of two different blades, and each provides visibility into all 32 channels at every location. Service providers, therefore, only have two part numbers to

deploy throughout the network, greatly reducing sparing requirements.

Cisco created its ROADM cards through greater silicon integration using Planar Linear Circuit (PLC) technology. First-generation PLCs were simple, often including only a 32-channel demultiplexer or multiplexer without any power measurement capability or use of VOAs. Today's PLCs can demultiplex all 32 channels with power taps for measurement of individual channel powers and offer a per-channel VOA. The ROADM packages in a very small space the ability to demultiplex, multiplex, attenuate, and switch on a per-channel basis.

The other challenge with static OADMs is that of cascading amplifiers used to boost all the power signals. This creates power variation, or *gain tilt*, in the channels. ROADMs act as channel equalizers and can continually flatten the spectrum automatically. Traditionally, this has been a manual function, requiring manual per-band equalizers to attenuate bands of channels.

On the Cisco ONS 15454 MSTP, this power control is handled automatically by software protocols with the intelligence needed to constantly optimize the DWDM system, eliminating the need to send a technician to set amplifier gains or VOA levels.

"With ROADM, it's one fiber in, one fiber out," says Johnson. "We demultiplex all 32 channels and have power visibility, and we can regulate attenuation on all 32 channels and choose either the pass-through path or the add path."

"Even more impressive is that this tremendous increase in functionality comes at a price that is comparable to fixed OADMs," adds Johnson. "And Cisco has developed a highly repeatable manufacturing process that will bring down the price points of ROADM much more quickly than labor-intensive processes such as blocker-based ROADM. What's more, unlike competing systems, our ROADM provides customers with all 32 channels on day one, eliminating the need to upgrade from 8 to 16 to 32 channels."

Accommodating Granular Services

Handling service granularity is another key challenge in metro networks. According to a report from RHK Inc: "In some regions, a large ROADM ring could easily satisfy demand for many years: 32 or 64 wavelengths at

Cisco ONS 15454 MSTP Release 4.7

The latest release of the Cisco ONS 15454 MSTP provides ROADM and DWDM functions along with a comprehensive suite of transparent wavelength service interfaces and subwavelength interfaces, including:

- DS-1/E-1, DS-3/E-3, OC-n/STM-n
- SAN: 1- and 2-Gbit/s Fibre Channel, ESCON, FICON
- Ethernet: Gigabit Ethernet, 10 Gigabit Ethernet (LAN/WAN)
- Optical: SONET/SDH: OC-3/STM-1 to OC-192/STM-64
- Intelligent DWDM: D1, HDTV, and wavelength-based data services

10 GB is a lot of capacity for most metro areas. In metro, the focus is on delivering services at a much finer granularity, such as T1/E1 or 10/100/1000 Ethernet or others, thus metro equipment must support the underlying services in addition to transport." ■

FURTHER READING

- Cisco ONS 15454 Multiservice Transport Platform cisco.com/packet/163_8b1
- Article on multiservice transport capabilities of the Cisco ONS 15454 (*Packet* Third Quarter 2003) cisco.com/packet/163_8b2

Calculating New Routes Faster

Software enhancements speed IS-IS network convergence.

By Clarence Filsfils

As an engineer who works regularly alongside Cisco's service provider customers, I have learned hands on that attaining high availability in IP networks requires more than ensuring that physical circuits and devices are always operating.

One day, one of my carrier customers expressed satisfaction over our joint success at improving his service-level agreements (SLAs) by deploying a backbone differentiated services design. Then he remarked, "If only we could do something to speed up the performance of the routing protocols!"

That got me to thinking. Soon thereafter, a set of capabilities that we call *Fast IS-IS Convergence* made its way into Cisco IOS® Software. This suite of software enhancements shaves quite a bit of time off the convergence process in networks that run the Intermediate System-to-Intermediate System (IS-IS) interior gateway protocol. Fast IS-IS Convergence is already being deployed, and some member service providers (see accompanying table) have described their experience with it at the RIPE Network Coordination Center (the Regional Internet Registry responsible for assigning IP addresses in Europe and Northern Africa).

Similar enhancements have also been made to several other routing protocols in the Cisco portfolio, such as Open Shortest Path First (OSPF); however, this article focuses on those available for the IS-IS algorithm.

Why Fast Convergence Matters

Convergence, of course, is the process whereby all networked routers agree on optimal routes in a network. When a network event—such as the loss or addition of an interface, router, or circuit—causes routes through the network to become available or unavailable, routers exchange update messages. These messages prompt routing algorithms to recalculate optimal routes that take these network-topology changes into account.

Eventually, all the routers agree on the "best" routes, based on the new information. They then update



CLARENCE FILSFILS, Distinguished Systems Engineer at Cisco, is an industry-recognized expert in advanced routing, QoS, capacity planning, and admission control technologies, and is active in the standardization of protection techniques for IP and MPLS networks. He can be reached at cfilsfil@cisco.com.

IS-IS CONVERGENCE OPTIMIZATION

	IS-IS Before	IS-IS Optimized
Failure Detection	2000 ms	<=20 ms
Delay Before Flooding LSP	33 ms	<=5 ms
Delay Before SPF Computation	5500 ms	1 ms to max. time (avg. 20 ms)
SPF Computation Time and RIB Update	40 to 250 ms	40 to 250 ms

Source: France Telecom

their routing tables with the revised hop-by-hop route information. Once these tables are updated, packets begin to be getting forwarded over the new routes.

The faster a network of routers can conduct these best-route calculations, the greater the benefit to network performance. By contrast, routing algorithms that converge slowly might cause routing loops or network unavailability. This could drag down a performance metric that is integral to a customer SLA or even affect the ability to deploy a particular premium service.

Starting with Cisco IOS Software Release 12.0(27)S, the IS-IS routing protocol can adapt to a network change in less than one second without compromising stability. Typically, other networks running interior gateway protocols such as IS-IS require up to 30 seconds to recover from a change. (Note: While this IS-IS solution has been pioneered based on service provider requirements for Cisco IOS Software Release 12.0(27)S and the Cisco 12000 Series Router, it will also be available for enterprise designs involving Cisco 7000 Series platforms and IOS 12.2S and 12.3T releases.)

Reducing convergence times from 30 seconds to under 1 second is a significant stride, and it is very important for successfully supporting real-time services such as voice over IP (VoIP). First of all, the human ear notices latency of 200 ms or more, and convergence time can figure into the overall latency budget. More importantly, calls can be dropped altogether if an outage lasts for more than a few seconds.

How the Optimization Works

The Cisco Fast IS-IS Convergence benefits are achieved within an autonomous system—when source and destination nodes are both IS-IS routers—as well as between an IS-IS node and an upstream Border Gateway Protocol (BGP) node. The convergence benefit also extends to that BGP node's next-

hop BGP neighbor. Cisco has accelerated the IS-IS convergence process in several ways:

- **Quicker detection of failures.** The Cisco packet-over-SONET (POS) implementation promptly detects failures and immediately signals the IS-IS protocol, keeping the detection process to less than 10 ms.
- **Faster announcement of failures throughout the network.** Once a failure has been quickly discovered, it is described in an IS-IS packet that is flooded across the network. The flooding code has been optimized to limit the process to less than a few milliseconds per hop.
- **Prioritized update of the routing table.** Service provider backbones often carry several thousands of IS-IS prefixes. However, only a small percentage of those prefixes are important to network operations. It is often accepted that the largest number of important prefixes will be 500, even though the total number of prefixes in routing tables might be as high as 5000.

To further hasten convergence, Cisco IOS Software uses a routing table process that prioritizes the importance of the packet destination. VoIP destinations are updated first, then VPN and Internet destinations, and then others. Convergence now involves updating a much smaller number of routing table entries. The reduction can be as high as ten- or twenty-fold, thereby accelerating routing table updates and decreasing downtime.

- **Caching of redistributed routes.** Cisco's IS-IS implementation caches routes that are redistributed from other routing protocols or from another IS-IS level into a local redistribution cache maintained by IS-IS. Caching occurs automatically and requires no configuration.
- **Accelerated computation of the new network topology.** Upon failure notification, all routers must run the Dijkstra algorithm on the entire topology to compute the "shortest path tree" (the set of best routes to any remote locations). With the latest Cisco route processors (for example, PRP-2 for the Cisco 12000 Series), a good rule indicates that the time to compute the shortest path tree is roughly $n \times 40\text{ms}$, where "n" is the number of nodes in the topology. For a network of 1000 nodes, we thus see that each network change requires a CPU computation lasting approximately 40 ms. Using the IS-IS Incremental Shortest Path First (Incremental SPF) algorithm in Cisco IOS Software, routers preserve as much of the previous convergence computation as possible and recompute only the part of the tree affected by the topology change. Depending on how far the topology change is from the computing node, the incremental SPF algorithm might reduce the computation by as much as 80 to 90 percent and, in several cases, smart logic might be used to avoid any computation. This indicates that the shortest-path computation itself is becoming negligible compared to

the subsecond convergence objective, and that more CPU is available for other tasks—providing even better overall network performance and stability.

- **No compromise in stability.** Together with these convergence enhancements, Cisco IOS Software delivers a set of intelligent capabilities that automatically adapts between network events that require immediate reaction and network events that need to be dampened to preserve stability. The root of these mechanisms is called *exponential backoff timer*, which controls how fast Cisco IOS Software reacts to an event. Initially, the reaction is almost immediate, but the software logs these iterations and, if they occur too often, it automatically delays the IS-IS reaction to preserve stability.

Several factors make use of this stability mechanism: IP event dampening (controls the reaction to link failures and filters out the negative impact of flapping links), IS-IS reaction timer for SPF computation, and IS-IS reaction timer for local link-state packet (LSP) generation.

Business Implications

Reduced network convergence time with no compromise in stability is an important enabler that allows IP networks to mature beyond best-effort data networking to truly converged networks. Such networks require consistent, high-performance service metrics for network availability (uptime), latency, packet loss, and jitter to help enable carriers to meet business goals, including deploying premium services such as VoIP, whose success depends on the ability to guarantee specific performance metrics; offering tighter SLAs for competitive differentiation; and conserving network operations and management costs.

A large factor in meeting these objectives is making sure that the network remains available and that packets continue to be forwarded for the greatest percentage of time possible. Fast IS-IS Convergence in Cisco IOS Software provides another network performance optimization tool by reducing convergence times in IS-IS networks to less than one second without compromising stability.

Enabling routers to converge at these rates improves overall network service levels and is especially helpful in supporting real-time network services such as voice. ■

FURTHER READING

- "Reducing Route Calculations in IS-IS and OSPF Networks Speeds Convergence" (from the Cisco *Beyond Basic IP* archives)
cisco.com/packet/163_8d1
- "High Availability in IP Routing," *IP Journal*, March 2004
cisco.com/packet/163_8d2

IP VPNs Gain Momentum

By out-tasking their IP VPNs to a managed services provider, SMBs can save time and money.



Rodney Davidson

Large enterprises have been at the forefront in the adoption of IP virtual private networks (VPNs). They are using the security, scalability, extensibility, and multiservice capabilities of IP VPNs as a lower cost alternative to legacy data networks such as Frame Relay and ATM. But service providers such as BellSouth see an even broader market opportunity for IP VPNs within the million-strong, small- and mid-sized business (SMB) market.

“Our VPN solution is generally targeted at BellSouth customers looking to simplify their WAN management. Many of these businesses are using BellSouth’s VPN service to cost-effectively add smaller locations to their corporate networks,” says Amy Hollister, a senior marketing manager for BellSouth.

These businesses can cut costs by out-tasking their IP VPN needs to service providers. By opting for a managed IP VPN, SMBs find that they can save capital and are better able to focus on their core business—and it frees up their IT staff to do the sort of the work that makes employees more productive. Moreover, they can make the transition to VPN in steps by starting, for example, with out-tasking remote access or by using an IP VPN as a backup or alternative to a site’s Frame Relay or private-line connection.

One of the ways in which BellSouth has made its IP VPN service more attractive to SMBs is by using DSL as a last-mile access

technology for telecommuters and branch offices to securely share corporate resources over its IP/Multiprotocol Label Switching (MPLS) core. The price point is low enough to prompt customers that have never considered a WAN before to cost justify the business operations impact of having real-time connectivity. “We have customers with existing Frame Relay networks that previously could not afford to extend Frame Relay to their small or remote offices. Now they are able to provide those sites with more bandwidth for less money using our Network VPN DSL service, and that’s very attractive to SMBs that want to add locations to their WAN,” Hollister says.

The Market

Although defined in various ways, there are slightly under 1 million SMBs in the US. The Yankee Group estimates there are 810,000 small businesses employing from 21 to 99 people and almost 92,000 mid-sized businesses with 100 to 499 employees. These SMBs, moreover, have emerged as an important customer segment when it comes to spending on information and telecommunications technologies.

Another analyst firm, AMI Partners, estimates that small businesses in the US spent about \$142 billion on IT/telecommunications in the 12-month period that ended October 2003, while spending by mid-sized businesses during the same period was

US\$57 billion. AMI Partners defines a small business as having from one to 99 employees, and a midsize business as having 100-499 employees.

Among small businesses that have been aggressive in adopting technology—or what the AMI Partners firm terms Tier 1 small businesses—more than half use a WAN. The analyst firm notes that Internet access for this group has reached “saturation,” with nearly half of the businesses using DSL for Internet access. Of the top-tier, midsize businesses, 92 percent have a WAN and all have Internet access (typically T1/T3 connections), according to AMI Partners.

Many of these businesses use a combination of legacy WAN technologies, such as private lines or Frame Relay, to link company locations to corporate headquarters, and IP VPNs to enable telecommuters, mobile workers, a small branch office, and even a business partner to exchange information. Although no one example paints a complete picture of the networking needs of SMBs, the case of Arizona State Savings and Credit Union is illustrative: The credit union has about 350 people in almost two dozen locations. Five sites acting as hub sites connect through leased lines to a data center in Glendale, Arizona, with smaller locations in turn linked by leased lines to the hubs in a traditional hub-and-spoke network topology.

Kim O’Connor, senior network engineer for the credit union, says the credit union uses IP Security (IPSec) VPNs for remote access, typically for LAN-to-LAN large file transfers with other organizations and for giving the credit union’s IT staff and senior managers access to the company’s network from their laptops.

The credit union is hardly alone in its use of an IP VPN. A recent Cisco global survey of network professionals found that, for SMBs, the top application for IP VPNs is to replace the dialup infrastructure used by teleworkers with secure Internet access to corporate resources. More broadly, SMBs also see IP VPNs as an intranet to securely communicate between sites as well as an extranet for facilitating business with partners.

Though the amount that many SMBs spend on their IP VPNs might be small, it is expected to increase significantly over time. According to the Cisco survey, although between one-third and one-half of SMBs spend less than US\$1000 per month on hardware, software, network connectivity, and self-managing their IP VPNs, another one-third spend substantial amounts—as much as from \$1000 to \$4999 every month—and about 10 percent spend as much as \$5000 to \$9999 per month.

Out-Tasking VPNs

Instead of spending money and effort on self-serviced VPNs, SMBs that out-task will find that service providers deliver VPNs in two ways, depending upon the technology. IPSec, the dominant IP VPN technolo-

Powering Scalability and Flexibility with IP/MPLS

For SMBs that want encryption, IPSec is currently the preferred building block for IP VPNs. Aside from Point-to-Point Tunneling Protocol (PPTP), which is generally used by teleworkers over dialup or broadband connections, other IP VPN technologies are MPLS, L2TPv3, and Generic Routing Encapsulation (GRE). Other than MPLS, which does not provide encryption, all of the aforementioned IP VPN technologies offer encryption and use tunneling methods across an IP network to establish point-to-point connections. Hence these are termed *overlay networks*, and it is because of their overlay nature that they present the same scalability problem inherent in scaling Frame Relay/ATM networks for customers who want to directly connect each location with every other location through Frame Relay/ATM virtual channels. This is a key reason why IP VPNs—other than MPLS-based VPNs—are configured in a hub-and-spoke topology in which branch locations connect either directly to the data center or to hub sites, and the hub sites to the data center. Unlike overlay networks, IP/MPLS enables any-to-any connections.

The scalability problem of point-to-point IP VPNs is further compounded when new sites are added to an intranet or an extranet. However, Cisco has developed a mechanism called Dynamic Multipoint VPN (DMVPN) that alleviates the scalability issue in point-to-point IP VPNs by simplifying operations and management of point-to-point IP VPNs to provide effects that are similar to the any-to-any connectivity of IP/MPLS.

gy in North America, and the new version of Layer 2 Tunneling Protocol version 3—L2TPv3—can be provided as customer premises equipment (CPE)-based managed services where the service provider manages and configures the VPN on the CPE router at each customer and business partner location. The other option is to deliver the VPN from the edge of a service provider’s network, or what is known as a network-based service. The choices here are MPLS, L2TPv3, and IPSec. One other IP VPN technology gaining popularity is Secure Sockets Layer (SSL), which enables per-application (for example, e-mail, Telnet, or FTP) VPNs using a Web browser.

The leading technology today for a managed, network-based VPN is IP/MPLS. This rapidly growing technology offers scalability because it can

support tens of thousands of VPNs across a common network; it allows traffic engineering to increase network availability; and, very importantly, it can provide quality of service (QoS). Because of these attributes, more than 200 carriers worldwide have deployed IP/MPLS in their network cores. They are using their IP/MPLS backbones as a single, converged platform for data, voice, and video traffic as well as to transport other Layer 2 services such as Frame Relay, ATM, and Ethernet.

Unlike other IP VPN technologies, IP/MPLS VPNs provide inherent full meshing or any-to-any connectivity between locations, thus enabling IP traffic to run over an IP/MPLS infrastructure without having to build point-to-point tunnels to connect every location to every other location (see sidebar, “Powering Scalability and Flexibility with IP/MPLS,” on page 82). The scalability problem of point-to-point technologies is not just an issue that affects large enterprise customers; it even has an impact on smaller businesses that might want secure any-to-any communications between as few as a handful of locations. Additionally, IP/MPLS can be combined with IPSec to offer SMBs an integrated architecture for secure site-to-site and remote access. This, in fact, is exactly what BellSouth has done.

BellSouth offers network-based IP/MPLS for site-to-site IP VPNs or “on-net” remote access and provides a fully integrated IPSec gateway for any “off-net” connections via the Internet. Because IP/MPLS is technology agnostic, the carrier’s customers can use various BellSouth services such as DSL, Frame Relay, and private lines for secure communications between their own locations and with business partners across BellSouth’s private IP/MPLS network. IP/MPLS completely separates one customer’s IP traffic from another customer’s IP traffic in a way analogous to Frame Relay/ATM networks. But off-net customers, or those customers not directly using various BellSouth connectivity services, can also become part of BellSouth’s IP/MPLS-based VPN. This is because the VPN can be extended across the Internet to a mobile worker or telecommuter who has IPSec running on the PC, or a branch office with IPSec-based CPE. In the case of these off-net users, their IPSec traffic would terminate at a BellSouth gateway, and from there would be carried across BellSouth’s IP/MPLS network. On-net users, on the other hand, would get secure, centralized access to the Internet through BellSouth’s network-based firewall and would also receive the benefit of the carrier’s intrusion detection system (IDS).

“By putting together a WAN, which is really what we are providing with our network-based VPN, we offload the day-to-day management of a network for a business that doesn’t have the staff to do it,” says BellSouth’s Hollister. “So if you look at a customer that’s taking advantage of both our site-to-site and remote-access VPN capabilities as well as our centralized Internet access [via the firewall], these definitely represent a



Make your businesses work smarter with insights, strategies, and news for decision makers. Subscribe today to get *iQ Magazine*, a free quarterly publication from Cisco for small and mid-sized businesses. cisco.com/go/iqmagazine/subscribe/packet.

cost saving for customers. And even if you were to take just the remote user VPN services, the management and administration time it would take to deploy a remote-access solution alone would be pretty high.”

Managed VPNs: a Partnership

Although eventually SMBs will want to move all of their data traffic to IP VPNs, later they will want to leverage their IP connections to support their voice and video needs. They can do this by taking advantage of the different classes of service (CoS) enabled by IP/MPLS and backed by very granular service-level agreements (SLAs) because of the traffic engineering and QoS capabilities of IP/MPLS networks.

SMBs that out-task their VPNs will find that the cost benefit of their relationship with the service provider will increase over time. This is because IP VPNs can serve as foundation for other managed services such as firewall, Internet gateway, IDS, telecommuter access—as BellSouth is doing—as well as IP telephony, managed LAN, storage, etc. When a service provider is able to extract more value from its network by offering more services, the additional revenues generated from these services enables the service provider to offer customers a multiservice package for less cost.

Managed IP VPNs offer SMBs more capability for less cost while also saving time and effort in order for them to focus more on their core business. Such benefits make managed IP VPNs a true win-win proposition for SMBs. ■

FURTHER READING

- *Power Up Your Small-Medium Business: A Guide to Enabling Network Technologies* (Cisco Press, 2004, ISBN 1-58705-135-4)
cisco.com/packet/163_9a1
- Cisco Security and VPN solutions for SMBs
cisco.com/packet/163_9a2
- Security Technical Implementation Guide for SMBs
cisco.com/packet/163_9a3
- Article on “What You Need to Know About MPLS” (*iQ Magazine*)
cisco.com/packet/163_9a4
- Article on “DMVPN Extends Business Ready Teleworker” (*Packet* Second Quarter 2004)
cisco.com/packet/163_9a5

IPSec & SSL: Complementary VPN Remote-Access Technologies

By Pete Davis

Enterprises are increasingly expected to provide secure remote-access connectivity to a diverse mix of users: remote employees who need access to the same applications and resources as their corporate-based counterparts; and customers, partners, and contractors who require connectivity to select applications.

Whomever the user, virtual private networks (VPNs) are the logical solution for secure, remote-access connectivity. The two prevalent VPN technologies are IP Security (IPSec) and Secure Sockets Layer (SSL). IPSec and SSL are complementary technologies that are best fits for different deployment scenarios. Companies do not have to choose one technology over the other.

In fact, together IPSec and SSL provide a complementary approach that can enable companies to better meet the differing needs of remote-access users—and converged IPSec and SSL VPN products eliminate the additional cost and care required for deploying separate, distinct platforms for IPSec and SSL VPNs.

IPSec VPN

Because of its ability to transparently support nearly any IP application, an IPSec VPN client on the remote system provides a user experience and workflow consistent with an office environment. IPSec allows access to IP applications on the enterprise network based on individual user privileges. It is easy to use and delivers very robust data encryption at the IP packet layer along with authentication, anti-replay, and data confidentiality services.

IPSec is ideal for persistent, or always-on, connections. The technology can detect a lost session and, because connections are generally established from a corporate-owned system, short idle timeouts are not required.

Cisco has optimized IPSec VPN for corporate remote-access security. It is available on many Cisco devices, including the Cisco VPN 3000 Series concentrators, Cisco PIX® security appliances, and Cisco IOS® routers.

SSL VPN

The main advantage of SSL VPN is that users can securely access their corporate network from any supported browser without preinstalled VPN client software. This is a very attractive option, especially for mobile workers. For example, a doctor on call can use a home computer or personal digital assistant (PDA) to access an in-patient record. A partner can upload a quote or work with a customer on a joint project that requires remote access to a particular application. And an employee can access secure content and e-mail from an Internet café.

The limitation of SSL VPN is that not every application is Web-enabled and supported by clientless SSL-based VPN. More complex applications might be supported with a downloadable applet or application that performs proxy or tunneling activities.

SSL/clientless connectivity is supported on existing Cisco VPN 3000 Concentrator devices with Cisco VPN 3000 Concentrator Software version 4.1 or greater. The software, free to Cisco customers with a SMARTnet™ contract, is available for download at cisco.com/kobayashi/sw-center/.

Both IPSec and SSL technologies provide support for strong authentication and one-time passwords. Soft tokens are often more practical for IPSec, because preinstalled software is required, whereas hard tokens can be generated by a separate, handheld device and used more easily with either technology.

Cisco VPN 3000 Concentrator: Best of Both Worlds

The Cisco VPN 3000 Series Concentrator supports IPSec VPN (Cisco VPN Client), SSL VPN (WebVPN), and Microsoft embedded clients. With the Cisco VPN 3000 Series Concentrator, SSL VPN and IPSec VPN connectivity can be achieved through a single device and management framework—allowing an enterprise to choose the most appropriate remote-access security solution for each user, as needed.

The Cisco VPN 3000 Series Concentrator can scale to meet the needs of the largest organizations and supports many authentication types, including RADIUS, Kerberos/Active Directory, Microsoft NT Domain, RSA SDI, X.509 digital certificates, and an integrated authentication server. The centrally configured, granular access controls ensure that users have access to only their designated corporate resources.

Twingo Acquisition Enhances Cisco SSL VPN

Cisco's acquisition of security firm Twingo earlier this year will bring enhanced SSL VPN and endpoint security features to Cisco's WebVPN (SSL VPN) solution. Twingo's Secure Desktop will provide a consistent and reliable technology for eliminating all traces of sensitive data (for example, history files, temporary files, caches, cookies, e-mail file attachments, and other downloaded data) at the close of each SSL VPN session.

In addition, the Twingo technology will allow access decisions to be made based on the presence of key system software or system watermark. The Twingo technology can operate on systems without the need for administrative privileges. ■



PETE DAVIS, product line manager for remote-access VPNs at Cisco, is a resident expert in IPSec and SSL technologies, and is responsible for driving new VPN-related products and features at Cisco. He can be reached at psd@cisco.com.

SPOTLIGHT ON:**Cisco Aironet 1300 Series Outdoor Access Point/Bridge**

Based on the IEEE 802.11g wireless standard, the Cisco Aironet® 1300 Series Outdoor Access Point/Bridge connects multiple fixed or mobile networks and clients in a metropolitan-area deployment. The product's flexible design allows it to operate as a wireless bridge, access point, or a workgroup bridge in campus, mobile, outdoor access, and temporary networks. Wireless LAN service providers and a variety of enterprises can benefit from the Cisco Aironet 1300 Series, including organizations in education, government, healthcare, military, public safety, and transportation.



As a WiFi-certified access point, the Cisco Aironet 1300 Series is suitable for indoor deployment but is also engineered specifically for harsh outdoor environments. As a bridge, the device supports either point-to-point or point-to-multipoint configurations and effectively provides simultaneous bridge and access point capabilities. As a workgroup bridge, the Cisco Aironet 1300 Series quickly connects any Ethernet-enabled device, supporting up to 255 device connections via a standard Ethernet hub or switch. The unit covers a range of 20 miles (32 km), delivers up to 54-Mbit/s data rates with security, and is compatible with legacy IEEE 802.11b wireless devices.

Based on Cisco IOS® Software, the Cisco Aironet 1300 Series supports advanced features such as fast secure roaming, quality of service (QoS), and virtual LANs. Maintenance and installation for the units is simplified through integration with a wired network via the Cisco Structured Wireless-Area Network (SWAN) solution. The Cisco SWAN solution is covered in greater detail on page 61.

For more recently released wireless products from Cisco, see page 86.
cisco.com/go/aironet

Core Routing**Cisco CRS-1 Carrier Routing System**

Designed for service providers and research organizations, the new Cisco CRS-1 Carrier Routing System supports scalable capacity up to 92 terabits per second. This capacity enables large-scale delivery of high-bandwidth applications such as video on demand, online gaming, and real-time interactive services. The Cisco CRS-1 introduces several innovations including the Cisco IOS® XR Software for continuous system operation, the programmable 40-Gbit/s Cisco Silicon Packet Processor, the industry's first OC-768c/STM-256c packet interface, a visual management tool, and the Cisco Intelligent ServiceFlex Design for service flexibility and rapid deployment. The multi-shelf CRS-1 system architecture supports up to 1152 40-Gbit/s slots in 72 line-card shelves that interconnect with eight fabric shelves, all operating as a single system. For more on the new Cisco CRS-1 Carrier Routing System, see "Reinventing the Router," page 41.

cisco.com/go/crs

Security and VPNs**Cisco Guard XT 5650 and Cisco Traffic Anomaly Detector XT 5600**

The new Cisco Guard XT 5650 and Cisco Traffic Anomaly Detector XT 5600 security appliances detect and mitigate distributed denial-of-service (DDoS) attacks. Cisco Guard XT 5650 recognizes anomalies that might be a security threat by comparing individual traffic flows to profiles of normal traffic patterns, behavior, and protocol compliance. The appliance also provides source verification and anti-spoofing capabilities, which block individual attack traffic flows while helping to ensure delivery of legitimate transactions. Cisco Traffic Anomaly Detector XT 5600 quickly and accurately identifies a broad range of known and previously unseen DDoS attacks and automates activation of the Cisco Guard XT 5650. These products are covered in greater detail on page 28.

cisco.com/packet/163_npd1

Cisco Catalyst 6500 Series Switches and Cisco OSR 7600 Router: Firewall Services Module Version 2.2

The Firewall Services Module software Version 2.2 offers a managed virtualization capability that allows a single physical Firewall Services Module in a Cisco Catalyst® 6500 Series Switch or Cisco 7600 Series Router to act as many virtual devices or “contexts.” With this capability, enterprises and service providers can deliver differentiated firewall service by customer, user, or application type and provide each with separate network management. The Resource Manager feature allocates performance and resource availability for each virtual firewall, enabling multiple service-level definitions and guarantees. The Firewall Services Module also includes new Layer 2 transparent firewall support, which flexibly segments the network into multiple Layer 2 security trust zones while preserving the existing IP addressing scheme.

cisco.com/packet/163_npd2

Cisco VPN Products: FIPS-Certified Models

New models of several Cisco virtual private network (VPN) and security products are available for companies that require products certified for FIPS 140-2 compliance. These products include selected models of Cisco VPN 3000 Series concentrators, the Cisco VPN 3002 Hardware Client, the Cisco VPN Software Client, the IPsec VPN Services Module for the Cisco Catalyst 6509 Switch and selected Cisco 7600 Series routers, and the Cisco 7206 VXR Router with the VPN Acceleration Module (VAM) and VPN Acceleration Module 2 (VAM2). Cisco FIPS 140-2 products comply with these stringent security requirements for government and enterprise customers in the US.

cisco.com/go/securitycert

Wireless

Cisco Catalyst 6500 Series Wireless LAN Services Module

The new Cisco Catalyst® 6500 Series Wireless LAN Services Module (WLSM) enables secure, campus-wide Layer 3 roaming by supporting up to 300 Cisco Aironet® access points and 6000 users. Because the Cisco Catalyst 6500 WLSM integrates wireless capabilities into the switch, no changes are needed in the underlying wireline infrastructure. This design simplifies deploy-

ment of a large-scale wireless LAN for an enterprise or service provider. Cisco Catalyst 6500 Series Switch capabilities are available transparently while fast, highly secure roaming gives users seamless access even in times of high network utilization. To control user access, the Cisco Catalyst 6500 WLSM can establish up to 16 logical mobility groups across different subnets. The module also supports centralized configuration and policy enforcement as well as nonstop forwarding and stateful switchover capabilities. For more on the Cisco Catalyst 6500 Series WLSM, see page 61.

cisco.com/packet/163_npd3 



Cisco 3200 Series Wireless Mobile Interface Card

The Wireless Mobile Interface Card (WMIC) for Cisco 3200 Series wireless and mobile routers provides integrated IEEE 802.11b/g wireless WAN or LAN capabilities. Based on the same ruggedized, compact PC/104-Plus architecture as the router's other interface cards, the WMIC eliminates the need to deploy external 802.11 bridges or access points to support in-vehicle networks and outdoor wireless infrastructure. The Cisco 3200 Series WMIC supports data rates up to 54 Mbit/s and can be configured as a root bridge, a non-root bridge, or as an access point. In a vehicle network, such as a police car, the Cisco 3200 Series Router with 802.11 capabilities seamlessly connects the moving vehicle network to 802.11 coverage areas and other wireless technologies such as cellular. A WMIC can also be used in the same configuration to add access point functionality in and around the vehicle, giving users with a wireless PDA or laptop the ability to stay connected wirelessly to the vehicle network. As an infrastructure device in an outdoor enclosure, the Cisco 3200 Series Router with the WMIC can be used in Cisco metropolitan mobile networks as an outdoor wireless base station. The router can connect multiple WMICs and other

wireless technologies in a single configuration combining Layer 2 and Layer 3 functionality, allowing overlapping coverage areas to scale across a city environment with link redundancy, security, and the manageability of Cisco IOS Software.

cisco.com/go/3200

Networked Home Linksys Wireless-G Presentation Player

The Linksys® Wireless-G Presentation Player enables wireless laptop and desktop PC users to show presentations on multimedia projectors, monitors, LCD panels, or other VGA-compatible devices using a remote control. Designed for conference rooms and auditoriums, the player connects via a standard VGA connector to a projector or can act as an IEEE 802.11g network access point. The Wireless-G Presentation Player is compatible with displays of up to 1024x768 resolution and 24-bit color, provides 32-MB flash memory for downloading presentations, allows file upload to a USB memory disk, and supports wireless Internet access with enhanced security and encryption.

cisco.com/packet/163_npd4

Linksys Wireless-B Media Link for Music and Wireless-B Music System

The Linksys® Wireless-B Media Link for Music uses IEEE 802.11b wireless networking to send audio content stored on a PC or other device to a home stereo. The product supports MP3 files and play lists as well as Internet radio, jukebox, and music streaming services. The media link device connects to a stereo system by standard RCA cables or an optical Sony/Philips Digital Interface (SPDIF) cable. Users select songs, play lists, and music services directly on the embedded LCD screen or via an infrared remote control. The Linksys Wireless-B Music System includes attached speakers for a complete, portable stereo player.

cisco.com/packet/163_npd5 



ABOUT NEW PRODUCT DISPATCHES

Keeping up with Cisco's myriad new products can be a challenge. To help readers stay informed, *Packet*® magazine's "New Product Dispatches" provide snapshots of the latest products released by Cisco between May and July 2004. For real-time announcements of the most recently released products, see "News Archive, News Releases by Date" at newsroom.cisco.com/dlls/.

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

Cisco IOS Software **Cisco Optimized Edge Routing** **Version 1.0**

Cisco Optimized Edge Routing (OER) provides multihomed enterprises the ability to enable intelligent network traffic load distribution and dynamic failure detection of data paths at the WAN edge. Whereas other routing mechanisms can provide both load sharing and failure mitigation, Cisco OER is unique in its ability to make real-time routing adjustments based on criteria other than static routing metrics, such as response time, packet loss, path availability, traffic load distribution, and cost minimization. Cisco OER is supported on Cisco 1700 Series, Cisco 2600 Series, Cisco 3700 Series, Cisco 7200 Series, Cisco 7300 Series, and Cisco 7500 Series routers. The software can be configured on a router with the command-line interface (CLI) or on an external Cisco CNS 2100 Series Intelligence Engine using a GUI interface. Cisco OER supports multiple routing protocols including Border Gateway Protocol (BGP) and static routing.
cisco.com/go/oer

Threat Detection, Continued from page 16

action. Frames sourced from 172.168.100.165 arriving at serial0/0 and failing the uRPF check are logged by the ACL log statement and forwarded by the ACL permit action.

The **show ip interface <>** command displays uRPF statistics for dropped or suppressed packets for a specific interface and can be used with the **show access-list** command to detect IP address spoofing. If ACL logging is enabled, the data logs can be viewed to gather additional information about the network attack.

```
Router# show ip interface serial0/0
  Unicast RPF ACL 101
  15 unicast RPF drop
  2581 unicast RPF suppressed drop
Router# show access-lists
  Extended IP access list 101
    deny ip 172.168.101.0 0.0.0.127 any log-input (15
  matches)
    permit ip 172.168.101.128 0.0.0.127 any log-
  input 2581 match)
```

Packet Capture and Sink Holes

Other popular techniques for detecting and classifying threats using Cisco IOS Software include *packet capture* and deploying *sink holes* in the network. Packet capture is one of the most helpful and com-

monly deployed techniques for detailed analysis of new threats. Because packet capture can generate huge amounts of data, it must be combined with a data analysis tool such as NetFlow to effectively identify and classify threats. Packet capture can also be done on routers using the **debug ip packet detail** command along with access lists to restrict debugging to packet streams of interest only.

It is necessary to perform detailed analysis of the packet capture using protocol analyzers; however, insertion of such external tools into the network might not always be possible. The *RAW IP Traffic Export* feature using the **ip traffic-export profile <>** command exports IP packets to and from the router via a designated LAN interface to a designated device in the network that can perform further analysis. This feature can be used in conjunction with access lists to export only packets of specific interest.

A sink hole is a part of the network topology that is deployed by manipulating the BGP routing tables to redirect attacks away from the network to a router or subnet designed to withstand the attack. Once the offending flows have been diverted, security managers can perform packet analysis, traceback, and other diagnoses on the miscreant packets with minimal impact to the production network.

◆ ◆ ◆

Early detection of spurious traffic patterns and their classification is a crucial first step toward effectively securing your network from various threats. Appropriate network architecture and traffic baseline profiling goes a long way in providing administrators with the ability to accelerate threat detection and classification with minimal impact to network operations. Tight integration among various Cisco IOS security features and IOS IP services such as QoS, ACLs, NBAR, intrusion protection systems, and firewalls can be used to deploy intelligent security policies that effectively inspect traffic, analyze and classify threats, mitigate risks, and help you defend your network. ■



PACKET ADVERTISER INDEX

ADVERTISER	URL	PAGE
ADC - The Broadband Company	www.adc.com/performance	D
AdTran	www.adtran.com/info/wanemulation	2
Aladdin Knowledge Systems	www.eAladdin.com/Cisco	IFC
American Power Conversion (APC)	http://promo.apc.com	F
BellSouth Business	www.bellsouth.com/business/netvpn	OBC
Boson Software	www.boson.com	A
Cisco Press	www.ciscopress.com/firststep	B
Cisco Systems	www.cisco.com/securitynow	32/33
Cisco Systems	www.cisco.com/go/ipcnow	70
Dalhousie University	www.dal.ca/internetworking	14
eiQ Networks	www.eiqnetworks.com	87
GL Communications	www.gl.com/packet	10
NIKSUN	www.niksun.com/packet	22
NIKSUN	www.niksun.com/pro	66
OPNET Technologies	www.opnet.com	80
Panduit	www.panduit.com/vip09	IBC
Pulver.com	www.von.com	52
RedSiren	www.redsiren.com/packet.htm	4
SMARTS	www.smarts.com	40
Solsoft	www.solsoft.com/packet	8
The Siemon Company	www.siemon.com/go/10Gip/pa	12
Websense	www.websense.com	30

Outdoor Wireless LAN Infrastructure

The Cisco Networking Professionals Connection is an online gathering place to share questions, suggestions, and information about networking solutions, products, and technologies. Following are excerpts from a recent *Ask the Expert* forum, “Outdoor Wireless LAN Infrastructure,” moderated by Cisco’s Jon Leary. To view the full discussion, visit cisco.com/packet/163_10a1. To join in on other interesting live online discussions, visit cisco.com/discuss/networking.

Q: *Our campus is using Cisco WLSE [Wireless LAN Solution Engine] version 3.1 with Cisco Aironet® 350 Series access points (APs) and bridges and Cisco Catalyst® 2950 Series switches. While using the Web interface of WLSE, we are not able to discover the devices using CDP [Cisco Discovery Protocol]. When using the connectivity tools, we can ping the APs, but the SNMP [Simple Network Management Protocol] reachable tool gives a SNMP timeout error.*

A: Is CDP enabled on the AP? Is SNMP enabled on the AP? Does WLSE have the correct community string? Is a firewall between the WLSE and AP that is blocking SNMP traffic? WLSE 1.3 is an older release, so you might want to consider upgrading to a newer one. If you do, note that the following bug (CSCsa04400 - GUI does not save snmp community string unless rebooted) was discovered in WLSE 2.5 and fixed in WLSE 2.7.

Q: *I have a problem with a solution between two Cisco Aironet 350 Series devices working as bridges for interconnecting two remote sites. Sometimes I don't have connectivity with the Cisco Aironet access points (pings don't respond), and then the other side connected through the bridge doesn't respond. I think the problem is with the STP [Spanning Tree Protocol] negotiating with the core of the LAN. Can I disable STP on the Cisco Aironet 350 when it works like a bridge?*

A: There are two areas to focus on for troubleshooting: the wireless interface and the network interface. On the wireless side, I would run a link test as well as a carrier busy test to determine the quality of the link—fade margin, interference energy, etc. On the network side, if you disable STP on the BR350, it becomes an access point so you can't take that approach.

Q: *Can anybody describe how collocated access points running IOS® provide load balancing?*

A: The load-balancing feature optimizes aggregate bandwidth with intelligent user associations, resulting in a better load distribution. At initialization, the client polls all access points within range for the device load information and selects the one with the lightest load. The access point interprets the request and provides loading information to the client. For additional load-balancing configuration information, visit cisco.com/packet/163_10a3.

Q: *What's your opinion on implementation of wireless access (802.11b) on a fast-moving target such as a train? The train will*

be moving from point A to B. Passengers will require wireless access all the way between A and B.

A: There are a number of challenges with trying to provide wireless access to a fast-moving object such as a train. The two major ones are RF penetration of the train and the channel impairments that the high relative speed induces. It has been our experience that while it is possible to get an 802.11 signal from a transmitter to a laptop inside the train, this becomes much more difficult as the train moves further and further away from the station. The train can effectively act as a Faraday cage, thereby screening the RF energy from reaching passengers inside the train. Instead, we suggest the use of an outdoor or window-mounted antenna that communicates with the station-mounted bridge or access point. Then a separate AP can be deployed in the train for providing the passengers with access.

Q: *I would like to set up a wireless connection between a storage facility and our main clinic approximately ¼ mile away, to cut out the need for a T1 line. If I use a Cisco Aironet 1400 at the warehouse to connect to a Cisco Aironet 1200 at the clinic, will I be able to connect to the 1400 with 802.11a wireless cards, or will I need an additional component at the warehouse? Also, can the Aironet 1200 be mounted outside and connected to an internal LAN via Cat5 cable, or will I need an external antenna?*

A: The BR1410 only supports the Root BR and Non-Root BR roles, and does not support client associations. Only other BR1410s can associate to a root BR1410. You could go with two BR1410s, but I'm guessing that would be a lot more power than you need for this relatively short link. The AP1200 is not weather hardened, but you have a couple of options. You could place the AP1200 indoors and run RF cabling out to an external antenna through a lightning arrestor. Or you could put the AP1200 in a NEMA enclosure with an antenna mounted outside the enclosure. You'll need to ground the Cat5 cable at the building entry point. Alternatively, you could use the Cisco Aironet 1300 Series, which serves as either an AP or a bridge. It is already IP56/NEMA4 rated so there is no need for an extra enclosure, and you can go with the integrated antenna version so no external antenna is needed. You could use the BR1310 on both sides or just one. For the remote side, you could either use the AIR-PCI35x with external antenna, the WGB352, or the BR1310.

Do you have a question about outdoor wireless LAN infrastructure? Ask the NetPro Expert. Send your question to packet-netpro@cisco.com, with the subject line “Outdoor Wireless LAN.” ■



JON LEARY, product line manager in the Wireless Networking Business Unit at Cisco, leads a team of product managers who focus on outdoor unlicensed wireless, public WLAN access, and voice over WLAN. He can be reached at leary@cisco.com.

CACHE FILE

Snippets of Wisdom from Out on the Net

CYBER QUOTE

“Looking at the proliferation of personal Web pages on the Net, it looks like very soon everyone on earth will have 15 Megabytes of fame.”

—M.G. Siriam

Cisco CRS-1 Makes Guinness World Records

In July of this year, the Cisco CRS-1 Carrier Routing System became the first networking technology to be recognized by Guinness World Records—as the highest capacity Internet router ever developed, with 92 terabits, or 92 trillion bps, of bandwidth capacity. When augmented by adequate network and transmission capacity, the Cisco CRS-1 would enable the following:

- The entire printed collection of the US Library of Congress could be downloaded in 4.6 seconds
- 1 billion people at the same time could play an online game, using real-time voice and chat
- The entire global population (6.4 billion) could have a simultaneous phone call using voice over IP

Filename Extensions

Got a file with the extension .emf or .nap? Don't know what it means? Check out “Every File Format in the World” at <http://whatis.techtarget.com/fileFormatA/>. Hundreds of filename extensions with brief descriptions for each are listed alphabetically.

Global Mobile Population on the Rise

The number of global mobile users will reach 1.8 billion by the end of 2007, according to estimates from The Yankee Group (yankeegroup.com). The research firm expects that Europe, the Middle East, and Africa (EMEA) will account for 40 percent of mobile users, with Africa's wireless market penetration alone reaching 13 percent by 2008, or 125 million subscribers and \$25 billion in revenue worldwide. The highest growth rate in this market, 13.6 percent, will be in Asia Pacific, which is expected to contribute 38 percent of the total revenue by the end of 2007, reports The Yankee Group.

Scandinavia Leads in E-Readiness

Denmark climbed to the top spot in this year's e-readiness ranking of 64 countries, conducted by IBM and the Economist Intelligence Unit. Other countries that made the top ten, in descending order, were the UK, Sweden, Norway, Finland, US, Singapore, Netherlands, Hong Kong, and Switzerland.

THE 5TH WAVE



“You know, this was a situation question on my Network+ exam, but I always thought it was hypothetical.”

©The 5th Wave, www.the5thwave.com