

PACKET

CISCO SYSTEMS USERS MAGAZINE

FIRST QUARTER 2005

BUSINESS RESILIENCE

**IT Strategies to Keep
Your Business Moving** 30

**Adaptive Security for
a Changing World** 19

**When Data and Voice Teams
Converge** 59

**SPECIAL REPORT:
Broadband Explosion** 49

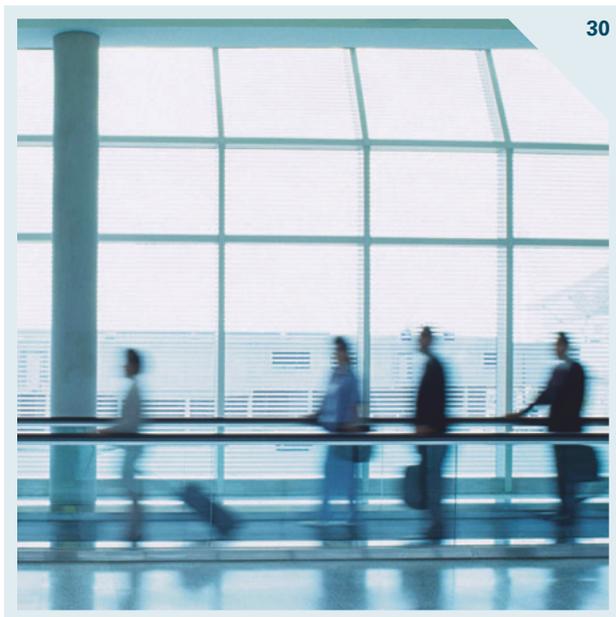


CISCO.COM/PACKET

PACKET

CISCO SYSTEMS USERS MAGAZINE

FIRST QUARTER 2005
VOLUME 17, NO. 1



ON THE COVER

How Resilient Is Your Business?

30

Successful business operation depends on the continuity of IT systems. And because everything upon which a business depends is part of an interconnected system, the entire system must be resilient. Read how network strategies that enhance operations can also protect your business.

Recipe for Resilience

35

An effective resilience strategy for your business involves building resilience into campus, branch office, and teleworker IT systems. Learn about the Cisco products and technologies that increase business agility and flexibility.

Ready for Anything

40

Enterprises that incorporate resilience into consolidated data centers can build secure, virtualized frameworks that adapt to future applications and technologies.

Adventures in Resilience

45

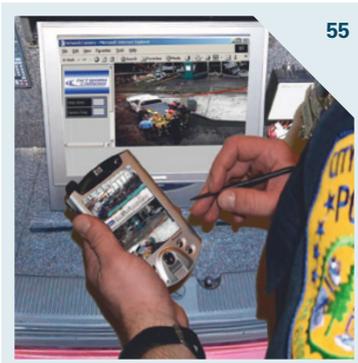
Four Cisco customers tell how they successfully maintained business continuity in the face of floods, fires, and power outages.

SPECIAL REPORT

Broadband Explosion

49

Broadband is proving to be the ideal platform on which to converge consumer applications for voice, video, and data. This report looks at the state of broadband technology worldwide—what it means for consumers and service providers alike.



55



69



77

IN EVERY ISSUE	
Mail	3
Calendar	6
Acquisitions	5
Networkers	6
Tech Tips	17
Advertiser Index	89
Cache File	90
The 5th Wave	90

TECHNOLOGY

SECURITY: Looming Security Challenges	19
Zombies, Trojans, "bots," and worms: What have we wrought?	
WIRELESS: Fixed-Mobile Convergence	23
Bringing wireless LAN and wireless WAN radio access technologies together for a seamless service offering.	
SECURITY: Cisco MPLS VPN over IP	27
Extending MPLS VPN to operate over IP networks with the same overall architecture and service experience.	

ENTERPRISE SOLUTIONS

Wireless Patrol	55
Two cities in the US and the UK demonstrate how mobile networks can enhance public safety and law enforcement.	
When Organizations Converge	59
Lessons learned by companies and employees that have successfully adopted converged IP networks.	
Technology Wakeup Call	65
Hotel operators look to IP networks to vanquish fierce competition and indulge sophisticated guests.	

SERVICE PROVIDER SOLUTIONS

Metro Ethernet Coming Your Way	69
Metro Ethernet gives service providers the flexibility and quality of service to deliver <i>real</i> broadband services—and customers are buying.	
The IP NGN Journey	72
On the journey toward IP next-generation networks (IP NGN), Cisco innovation and product advancements are helping service providers make the trip.	
Spinning IOS into Gold	75
How to turn IOS technologies into profitable services.	

SMALL AND MIDSIZED BUSINESSES

Passing the Test	77
Early Field Trials participants discuss the challenges and benefits of working with pre-release products.	
Who's at Your Service?	81
Setting up support contracts for your technology requires flexibility on both sides of the negotiating table.	

DEPARTMENTS

From the Editor	1	Technically Speaking	83
Business Resilience and Network Relics		Peter Lei and Randall Stewart on new Rate-Based Satellite Control Protocol feature in Cisco IOS Software.	
User Connection	5	New Product Dispatches	84
New CCIE Storage Certification • CCNA Prep Center • Worldwide Technology News		What's new from Cisco over the past quarter.	
Tech Tips & Training	9	NetPro Expert	88
Enterprise Segmented Extranets • Reader Tips		Advice from Cisco's Gilles Dufour on load balancing with Cisco CSM, CSS, and their SSL modules.	

PACKET MAGAZINE

David Ball
Editor in Chief

Jere King
Publisher

Jennifer Redovian
Managing Editor

Susan Borton
Senior Editor

Kim Austin Peterson
SMB Editor

Karen Dalal
Staff Editor

Joanie Wexler
Contributing Editor

Robert J. Smith
Sunset Custom Publishing
Project Manager

Amy Mackey, Nicole Mazzei,
Mark Ryan, Norma Tennis
Sunset Custom Publishing
Production

Jeff Brand
Art Director

Emily Burch
Designer

Ellen Sokoloff
Diagram Illustrator

Bill Littell
Print Production Manager

Cecelia Glover Taylor
Circulation Director

Valerie Marliac
Promotions Manager

Getty Images
Cover Photograph

Advertising Information:

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

Publisher Information:

Packet magazine (ISSN 1535-2439) is published quarterly by Cisco Systems and distributed free of charge to users of Cisco products.

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

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

Packet copyright © 2005 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

FROM THE EDITOR

Business Resilience and Network Relics

IP networks are no longer merely a way to interconnect computer resources; they are now a mainstream method of communicating and interacting with customers, employees, and partners. Organizations worldwide rely on their IT infrastructures to be more efficient, productive, agile, and competitive.

Mobility solutions such as virtual private networks (VPNs) and wireless LANs (WLANs) provide anytime, anywhere access to business-critical applications such as customer relationship management, sales force automation, and enterprise resource planning. IP telephony not only helps slash costs, it fundamentally changes the way business communicates. In short, the role of IP networks has shifted from a largely opaque support mechanism (read: cost center) to an integral part of business operations (read: competitive advantage).

If the network is so important to the business, then it follows that network resilience is vital to business resilience. What is business resilience, you ask? Business resilience refers to the *operational and technological readiness* that prepares organizations to make day-to-day operations efficient and cost-effective, respond quickly to opportunities with the potential to increase competitive advantage, and react appropriately to unplanned events.

But network resilience is only part of the story. An organization's business resilience strategy should consider how IT systems interact with each other. You can't just look at point products or single systems; you have to consider how all IT systems interoperate to achieve your goals for business agility and continuity. How resilient is your business? Turn to our feature articles on business resilience, beginning on page 30, to find out.

While an aspect of overall business resilience has to do with longevity of the individual products that make up the network—continuously having to pull out and replace boxes can wreak havoc on network availability and stability—we have a hunch that some of you network folks can take the notion of investment protection to the extreme. Do you have a Cisco museum piece still running on your network? An AGS router or a Catalyst 2500 Series Switch, perhaps? If so, we want to hear about it.

We're looking for the oldest, continuously running Cisco equipment in a production network across a variety of product series. Visit cisco.com/packet/museum for contest requirements, a list of product series that qualify, and details on how you can enter the contest. You could be featured in an upcoming issue of *Packet*, be included in other Cisco PR activities, or even receive a free upgrade to the "latest and greatest" replacement equipment.

While you're searching for that ancient artifact hiding in your network, ask yourself how it could be contributing to or hindering the resilience of your business.

David A. Ball

David Ball
Editor in Chief
daball@cisco.com



Rob Brodman

MAIL

Deciphering AES

The article “Bundled Security” [Fourth Quarter 2004, page 38], states that “AES supports 128-, 192-, and 256-bit block cipher lengths and encryption key sizes.”



The actual Advanced Encryption Standard (AES), which is available at csrc.nist.gov/publications/fips/fips197/fips-197.pdf, only supports 128-bit block cipher lengths. The original Rijndael algorithm (on which AES is based) supported additional block sizes, but they are not part of the AES standard.

—Jim Burtoft, Blair Technology Group, Altoona, Pennsylvania, USA

The AES standard is 128-bit cipher length and 128, 192, and 256 key length. The original Rijndael algorithm had a variable cipher length of 192 and 256, but they were not adopted by the National Institute of Standards and Technology (NIST). The Federal Information Processing Standards (FIPS) 197 standard you mention validates that point: “using cryptographic keys of 128, 192, and 256 bits to encrypt and decrypt data in blocks of 128 bits.” The sentence you pointed out in Packet should actually state: “AES supports 128-, 192- and 256-bit encryption keys in 128-bit block ciphers.”—Editors

Bullish on Torrus

Packet readers might be interested in a free open-source product for collecting and monitoring data series. For more than two years I have been working on the Round Robin Database Framework (RRFW) platform, soon to be released with the new name Torrus (torrus.org).

Most users use Torrus as a Simple Network Management Protocol (SNMP) collector, but its functionality is much broader. A single server can handle a few

hundred network devices, and its SNMP discovery engine supports many vendors and device types, including Cisco routers and switches and some Cisco-specific features such as SAA agents, temperature sensors, and CPU and memory statistics. It is also the first complete open-source implementation of Cisco class-based quality of service (QoS) monitoring. The threshold monitor daemon can periodically check if the values and rates are in accordance with specified limits. The Web interface provides great flexibility in organizing the data sources, and in managing user access rights.

Torrus is distributed under GNU General Public License and commercial support is provided by a Swiss company.

—Stanislav Simyagin, CCIE No. 5478, GPS-Technik AG, Schlieren, Switzerland

Looking for the Penalty Box

Sometime ago I read a very interesting Packet article about a network administrator who had created scripts to reduce bandwidth usage in a university campus environment. He devised some policies that he put on the switches to prevent bandwidth overuse by students and others. Where can I find this article?

—Claus Bernitsen, CCNP, TDC Totalløsninger, Odense, Denmark

The article you are looking for is “The Penalty Box,” from our Second Quarter 2004 issue. You can find it on our Website at cisco.com/packet/171_2a1.—Editors

Why No Picture?

Your article, “The VT Advantage,” [Packet, Second Quarter 2004] mentions body language, but you overlooked the obvious fact that body language includes the eyes. In describing Cisco’s video telephony solution you failed to include a picture of the video capability. A picture speaks for itself and can be accompanied

by a precise, short description, making extra words unnecessary. If your readers want more information, why not refer them to another Web link for the details? The eyes are the most important aspect of body language, so please include a photo of the video telephony solution.

—Egno F. Reyes, Verizon

Thank you for your suggestion. Although the article on Packet Online does not include a photo of the VT Advantage solution, the print version of the magazine does. You can download a PDF of the article, including the photo, from our Website at cisco.com/packet/171_2a2. Scroll to the end of the article and click on the PDF icon. For ease of use and readability we have redesigned Packet Online and the PDF’s of current articles are now easier to find at the top right side of article Web pages.—Editors



Send your comments to Packet

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

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

New CCIE Storage Certification

Building on the success of three specialist certifications that validate storage networking design, support, and sales skills, Cisco now offers an expert-level CCIE certification in storage networking.



Although there are no formal prerequisites for CCIE certification, candidates are encouraged to have 3–5 years of experience with end-to-end storage networking. The certification process includes a written exam covering theoretical concepts and a hands-on lab exam featuring devices such as storage area network (SAN) switches and directors, routers, and storage management solutions.

Candidates are expected to demonstrate expertise configuring intelligent storage solutions using transport technologies such as Fibre Channel, Small Computer System Interface over IP (iSCSI), Fibre Channel over IP (FCIP), and IBM Fiber Connection (FICON).

Business Value of Storage

“New storage demands are rapidly outstripping the way in which data storage is managed at most organizations today,” says Mike Quinn, vice president, Cisco Technical Support

Group. “Cisco is offering storage networking certifications that are evolving in step with what the market and our customers are demanding for optimizing productivity and performance.”

Companies are using storage networking to improve disaster recovery and network performance, and to take advantage of network services such as volume management, data replication, and enhanced integration with servers and storage appliances (see “Ready for Anything” on page 40).

Exam Availability

The CCIE Storage Networking written exam is expected to be available by March 2005. The corresponding lab exam is expected to be announced soon after. Initially, the lab exam will be available at Cisco offices in San Jose, California; Research Triangle Park, North Carolina; and Brussels, Belgium.

Cisco also offers CCIE certifications in routing and switching, network security, service provider networks, and voice.

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

Recently Announced Cisco Acquisitions

Acquired		Employees	Location
Airespace, Inc.	Provider of wireless local area network (WLAN) systems, including WLAN controllers, access points, management software, and intrusion detection system capabilities. Its employees will join Cisco's Data Center, Switching, and Wireless Technology Group.	175	San Jose, California, USA
BCN Systems, Inc.	Provider of networking software infrastructure design that will contribute to the evolution of Cisco's next-generation IP technology. Its employees will join the Cisco Routing Technology Group.	45	Santa Clara, California, USA
Jahi Networks, Inc.	Provider of network management appliances designed to simplify device deployment, configuration, and management. Its employees will join Cisco's Network Management Technology Group.	20	San Jose, California, USA and Hyderabad, India
Perfigo, Inc.	Developer of network access control solutions that analyze endpoint devices, scan for vulnerabilities, and enforce network access policies. The Perfigo team will join Cisco's Security Technology Group.	31	San Francisco, California, USA
Protego Networks, Inc.	Provider of security monitoring and management appliances that can detect, correlate, and mitigate network security threats. Its employees will join the Cisco Security Technology Group.	38	Sunnyvale, California, USA

CCNA Help Is Only a Click Away



For anyone starting on the path to Cisco certification, the CCNA Prep Center is an online resource available to all candidates preparing for associate-level CCNA certification exams.

The site provides sample questions, labs and simulations, e-learning modules, tips and advice from CCNA certified professionals and

other networking experts, candidate success stories, peer discussion forums, and more. For additional training and support, candidates are encouraged to contact Cisco Learning Partners or companies that are authorized by Cisco to administer career certification training courses.

To access the site, you must be a Cisco.com registered user. Visit Cisco.com for registration information. For more information about the CCNA Prep Center, visit cisco.com/go/prepcenter. ■

CISCO WORLDWIDE EVENTS

April 4–5, 2005	Cisco Partner Summit, Vancouver, B.C., Canada
April 12–14, 2005	Storage Networking World, Phoenix, Arizona, USA
April 25–27, 2005	Cisco Powered Network Marketing Summit San Diego, California, USA
May 3–5, 2005	Networld + Interop, Las Vegas, Nevada, USA
June 7–9, 2005	Supercomm, Chicago, Illinois, USA
June 19–24, 2005	Networkers, Las Vegas, Nevada, USA
Sept. 19–22, 2005	Networkers Australia, Gold Coast, Australia
Nov. 1–3, 2005	Networkers Korea, Seoul, Korea

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

Cisco Technology News from Around the World

Cisco Invests in Japan-Based R&D Center

Cisco opened a new research and development facility in Tokyo, Japan, in February 2005, with plans to invest US\$12 million over the next five years. With a projected staff of 10 engineers, the new center will focus on the development of IP-based networking technologies, including routers, Cisco IOS XR Software, IP version 6 (IPv6), IP Multicast, wireless, network security, and quality of service.

Japanese service provider networks carry loads that are five times higher than in the US, and broadband access is growing at more than 500 percent a year. "Products and technologies produced to meet Japan's demand for intelligent bandwidth will be robust enough to handle any other market in the world," says Mike Volpi, senior vice president and general manager of Cisco's Routing Technology Group. For more information, read a Q&A with Volpi at cisco.com/packet/171_3b1.

Cisco and NetHope Deliver NetRelief Kits to Tsunami Areas

Cisco is working with the NetHope consortium, which includes companies and nongovernmental associations, to make "NetReliefKits" (NRKs) available in the disaster areas stricken by the Asian tsunami of December 2004. The kits make it easy for non-technical people to set up and operate a communications hub where normal communications infrastructure is absent or destroyed. The NRK is a rugged, suitcase-sized, wireless voice and data communications device, with access to the Internet through a mobile or fixed satellite station. For more information about NetHope and deployment of the kits, visit nethope.org.

Hewlett-Packard and Cisco Worldwide Support Services

Hewlett-Packard (HP) and Cisco will deliver co-branded support services through a single point of contact to help customers maintain their enterprise networks. The agreement is part of Cisco's Global Services Alliance program, which was established by Cisco in June 2004 to focus on support services.

Initially, HP and Cisco will provide a single source for network support and problem resolution, and ready access to Cisco product knowledge and expertise. Over time, the combination of HP and Cisco offerings will be designed to address the full lifecycle of services, from planning through deployment to management and support, providing greater consistency worldwide and helping customers with global operations to realize greater return on their network investments. For more information about the Global Certified Partner designation, visit cisco.com/packet/171_3b2.

Networking Academy Program Expands Reach in Vietnam and the Philippines

The Cisco Networking Academy Program (cisco.com/go/netacad), which teaches students and others how to design, build, and maintain computer networks, will gain access to underserved areas of Vietnam because of an expanded relationship with the United Nations Development Programme (UNDP).

A UNDP-sponsored United Nations Volunteer (UNV) has been

appointed to widen the availability of the Cisco program in small Vietnamese cities such as Cantho, Dalat, Danang, and Hue. UNVs support Networking Academy expansion in remote areas throughout the Asia-Pacific region. UNVs are working in Bangladesh, Cambodia, China, India, Indonesia, Mongolia, Nepal, Sri Lanka, Thailand, and now in Vietnam. For more information about UNVs, visit unvolunteers.org.

AMA Computer University (AMACU) is the first Networking Academy in the Philippines to offer the professional-level Cisco CCNP curriculum. Successful completion of the 280-hour advanced CCNP curriculum will prepare AMACU students for the CCNP exam, which is a prerequisite for CCNP certification.

"AMACU's qualification as the first CCNP Academy in the Philippines will accelerate the growth of a highly skilled workforce that will provide the foundation of the Philippines' ability to compete in the global economy," said Luichi Robles, country manager, Cisco Systems Philippines.

First launched in 1997, the Cisco Networking Academy Program now has 10,000 Academies in more than 145 countries worldwide. For more information about CCNP certification, visit cisco.com/go/ccnp.

Netherlands Cable Operator to Test Internet Speeds

Up to 30 Mbit/s

Cable operator UPC Nederland started field trials in Almere, the Netherlands, with the goal of offering Internet services with download speeds of 30 Mbit/s. At these speeds, an average 7-GB digital movie could be downloaded in 30 minutes compared to 16 hours with a 1-Mbit/s cable or asymmetric DSL (ADSL) service.

In 2006 UPC plans to run a trial in Amsterdam with speeds as high as 50 Mbit/s. The goal of both trials is to show that the company's fiber-optic cable network is ready to meet future demands for high-speed Internet connections. For the Almere trial, UPC will set the cable modems of 300 subscribers to 30 Mbit/s for downloads and 1 Mbit/s for uploads. Various other speeds will be tested as well.

For more information about these announcements, visit cisco.com/go/news and enter a relevant search term. ■

Is There a Museum Piece in Your Network?

Do you have an AGS router, a Catalyst 2500 switch, or any other antique Cisco products still running on your network? If so, we want to hear about them. We're looking for the oldest, continuously running Cisco equipment in a production network across a variety of product series. Visit cisco.com/packet/museum for contest requirements, a list of the products that qualify, and details on how you can enter the contest. You could be featured in an upcoming issue of *Packet*, be included in other Cisco PR activities, or even receive a free upgrade to the latest replacement equipment.

Enterprising Extranets

Providing Enterprise Segmented Extranet Services Using IP-Based VPN

By Laure Andrieux, Zaheer Aziz, and James Kline

As enterprises grow, gaining a competitive advantage most often entails building strategic partnerships. One increasingly popular way for business partners to exchange information is via extranets.

Organizations use the extranet—a limited part of a company's intranet that is extended to users outside the company—to exchange data, share product information, collaborate with other companies, develop and implement training programs, and provide or access services. Extranets require security and privacy, including such mechanisms as firewall server management, digital certificates or similar means of user authentication, and message encryption. An *extranet VPN* is a virtual private network that enables companies to securely share some information or operations with suppliers, vendors, partners, customers, or other businesses.

Because extranet VPNs allow external traffic to traverse the same links as a corporate intranet, enterprises understandably raise several questions. For example, how will routing and data separation be handled? What about address changes? And security? A Multiprotocol Label Switching VPN (MPLS VPN) addresses most of these enterprise concerns.

Routing and Data Separation

An MPLS VPN achieves routing separation in two ways. One way is by assigning each extranet VPN to a Virtual Routing and Forwarding (VRF) instance. Each VRF on the customer router is populated with routes from a unique VPN, either through statically configured routes or through routing protocols that run between the customer routers.

The second way is by adding a unique VPN identifier (a route distinguisher) to Multiprotocol-Border Gateway Protocol (MP-BGP). MP-BGP exchanges VPN routes between associated extranet edge routers, which keep routing information in VPN-specific VRFs. Using the route distinguisher ensures that routing across the customer intranet network remains separate for each extranet VPN and separate from the intranet global routing table.

No Address Changes

An extranet VPN service should not require major changes to an enterprise's internal IP network, desktops, or servers. For cost and security reasons,

enterprises want to retain their existing addressing scheme. IP VPNs allow for overlapping address spaces and, as such, the hosting enterprise might need to provide and control Network Address Translation (NAT) services and proper routing.

Segmentation and Encrypted Communications

Enterprises can exercise tradeoffs between security and cost when provisioning an extranet with a partner. In all cases, the enterprise maintains full control over VPN separation. The enterprise treats the extranet customer as untrusted and accepts only pure IP packets from them. Not only is the extranet traffic kept separate from the enterprise intranet network, extranet partners can also be kept separate from each other. Firewalls and intrusion detection devices may be used with the VPN in the shared services network for host protection as well.

An enterprise has the option of sending encrypted traffic through a properly configured extranet service, enabling regulatory compliance and enhancing data security. Encryption operates between enterprise routers. MPLS VPN and IP Security (IPSec) encryption work well in combination, but proprietary or application-level encryption schemes are also compatible if packet payload is transparent to the enterprise network.

IP Tunneling

Using IP tunneling technologies, you can build an extranet with all the benefits of MPLS VPN over an IP infrastructure. Tunneling bridges disparate network segments. In the case of extranets, tunneling adds segmentation policies that need to be preserved. When using any tunneling technology, note the maximum transmission unit (MTU) size. Adding tunnel headers and VPN labels adds size to the IP packets. Because the encapsulation layers add overhead to the original data payload, MTU must be considered.

Besides the IP and higher-level layer payload, you also need to account for a VPN label and possibly a Label Distribution Protocol (LDP) label, which are both 4 bytes in length; a Generic Routing Encapsulation (GRE) header, which is 24 bytes; and Layer 2 Tunneling Protocol version 3 (L2TPv3), which can be up to 16 bytes in length. This information is relevant

when deciding whether fragmentation is needed in the network and where to perform it.

MPLS VPNs, Extranet VPN Solutions

MPLS VPNs allow the separation of customer address space (extranet VPN customers) on the extranet edge router (using what would be the normal functionality of the provider edge, or PE, router). When extranet traffic enters the extranet edge router, a unique VPN label is appended to the incoming IP packet, and the packet is then switched across a tunnel label switch path (LSP). In this case, the LSP is composed of tunnels that traverse the enterprise intranet.

The VPN labels, along with VPN routes, are distributed between extranet edge routers using the MP-BGP extensions, and the label is appended to customer traffic prior to traversing the core network. Then the IP tunnel header (which replaces the outer label) is used to route traffic between the enterprise intranet routers until the traffic reaches the destination extranet edge router. The tunnel IP header and the VPN label are removed just before sending the traffic to the destination network. Label and tunnel header imposition and disposition are completely transparent to the extranet customer. This service requires that the enterprise is participating in the extranet customer routing at the extranet edge-to-enterprise edge demarcation.

Following is a look at two of the options for providing extranet VPN services: *MPLS VPN over GRE tunnels* and *MPLS VPN over L2TPv3 tunnels* (see figure, page 11). Both approaches are based on the Internet Engineering Task Force (IETF) RFC 2547.

MPLS VPN over GRE Tunnels

GRE takes packets or frames from one network system and places them inside frames from another network system in a peer-to-peer configuration. GRE consists of a packet header with components that allow it to identify data for processing when it arrives at the associated peer. These components include an IPv4 tunneling or delivery header; a GRE header with optional fields that include tunnel key, checksum, and sequencing fields; and the payload (or tunneled Layer 3 packet).

MPLS VPN over L2TPv3 Tunnels

L2TP is primarily used by service providers to deploy VPNs directly to their business customers, by ISPs in a

wholesale dial scenario, and by enterprises to support remote users. L2TP is an industry standard that combines components of the proprietary Microsoft Point-to-Point Tunneling Protocol (PPTP) and Cisco Layer 2 Forwarding (L2F) protocols.

L2TPv3, currently an IETF draft standard, expands L2TP to include several new service models. It also supports tunneling alternate Layer 2 protocol data units (PDUs) rather than just PPP. The primary application of L2TPv3 is for service providers to consolidate multiple Layer 2 networks onto a single high-speed IP network infrastructure for operational savings, or to allow traditional Layer 3 service providers to offer conventional Layer 2 services without building out an entirely separate core. In the configuration example on page 11, we will transport the MPLS VPN service over the L2TPv3 tunnel instead of a Layer 2 service.

Note: Cisco IOS Software Release 12.0.(29.4).S2 was used for proof-of-concept testing in the following GRE and L2TPv3 tunnel configuration examples.

Sample Configuration: MPLS VPN over GRE Tunnels

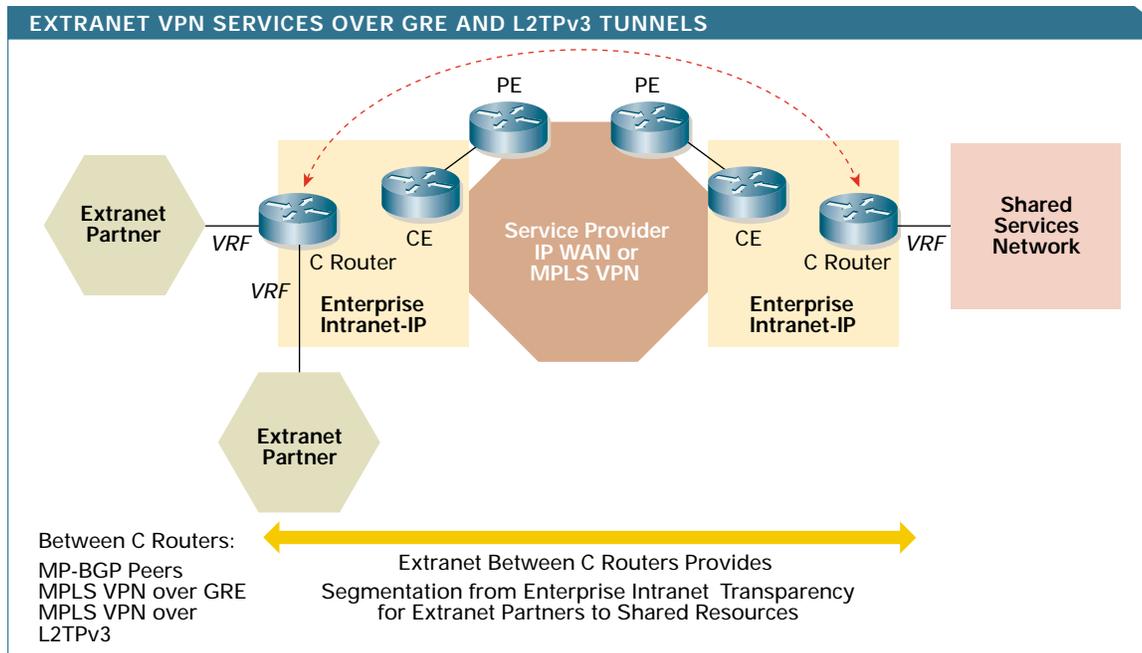
The sample configuration below shows the extranet routers as MP-eBGP peers using the tunnel interfaces as BGP peers. LDP is not needed in this configuration. If running MP-iBGP, it might be necessary to run LDP on the tunnel interfaces, and transport LDP label will be needed if the tunnels are not fully meshed. In this example, only one side of the tunnel is given. The router at the other side of the tunnel would be a mirror of this example.

```
Enterprise-C-Router#
ip cef
ip vrf extranet-customer
description extranet customer vrf
rd 200:1
route-target export 200:1
route-target import 300:1
!
!
interface Loopback0
ip address 20.1.1.1 255.255.255.255
no ip directed-broadcast
!
interface Loopback1
description Tunnel5 source interface
ip address 20.121.121.1 255.255.255.0
no ip directed-broadcast
!
interface Tunnel5
ip address 120.120.120.1 255.255.255.0
no ip directed-broadcast
mpls bgp forwarding
tunnel source 20.121.121.1
tunnel destination 20.129.129.1!
interface Ethernet0/0
```

LAURE ANDRIEUX is a network engineer focusing on MPLS VPNs and MPLS network management in Cisco's Solution Engineering Group. She can be reached at andrieux@cisco.com.

ZAHEER AZIZ is a technical leader focusing on IP VPNs in Service Provider Solution Engineering at Cisco. He is co-author of the Cisco Press book *Troubleshooting IP Routing Protocols*. He can be reached at zaziz@cisco.com.

JAMES KLINE is a network engineer focusing on MPLS networks in Cisco's Solution Engineering Group. He is a CCIE with specializations in both Routing & Switching and Security. He can be reached at jakline@cisco.com.



TUNNEL VISION
MPLS VPN over GRE tunnels works in a peer-to-peer configuration. The L2TPv3 tunneling option expands L2TP to include several new service models.

```

description Interface Facing the Enterprise IP Core
ip address 21.20.1.1 255.255.255.0
no ip directed-broadcast
!
interface Ethernet1/0
description Extranet Customer Facing Interface
ip vrf forwarding extranet-customer
ip address 30.1.1.1 255.255.0.0
no ip directed-broadcast
!
router eigrp 200
network 21.20.1.0
network 20.1.1.0
network 20.121.121.0
!
router bgp 200
no bgp default ipv4-unicast
bgp log-neighbor-changes
timers bgp 10 30
neighbor 120.120.120.9 remote-as 300
!
address-family ipv4
no auto-summary
no synchronization
exit-address-family
!
address-family vpnv4
neighbor 120.120.120.9 activate
neighbor 120.120.120.9 send-community extended
exit-address-family
!
address-family ipv4 vrf extranet-customer
redistribute connected
no auto-summary
no synchronization
exit-address-family
!

```

Sample Configuration: MPLS VPN over L2TPv3 Tunnels

The following sample shows the extranet routers as MP-iBGP peers using the loopback interfaces as BGP peers. LDP is not needed when running L2TPv3, so there is never an additional label. The L2TPv3 endpoints will be auto-discovered via tunnel address family in BGP. In this example, only one side of the tunnel is given. The router at the other side of the tunnel would be a mirror of this example.

```

Enterprise-C-Router#
ip cef
ip vrf l3tunnel
rd 200:101
!
ip vrf extranet-customer
description extranet customer vrf
rd 200:1
route-target export 200:1
route-target import 300:1
!
!
interface Loopback0
ip address 20.2.2.2 255.255.255.255
no ip directed-broadcast
!
interface Tunnel6
ip vrf forwarding l3tunnel
ip address 121.121.121.2 255.255.255.0
no ip redirects
no ip directed-broadcast
tunnel source Loopback0
tunnel mode l3vpn l2tpv3 multipoint
!
interface Ethernet0/0

```

```

description Interface Facing the Enterprise IP Core
ip address 20.20.2.2 255.255.255.0
no ip directed-broadcast
!
interface Ethernet1/0
description Extranet Customer Facing Interface
ip vrf forwarding extranet-customer
ip address 30.2.2.1 255.255.0.0
no ip directed-broadcast
!
router eigrp 200
network 20.2.2.0
network 20.20.2.0
!
router bgp 200
no bgp default ipv4-unicast
bgp log-neighbor-changes
timers bgp 10 30
neighbor 20.8.8.8 remote-as 200
neighbor 20.8.8.8 update-source Loopback0
!
address-family ipv4
no auto-summary
no synchronization
exit-address-family
!
address-family ipv4 tunnel
neighbor 20.8.8.8 activate
exit-address-family
!
address-family vpnv4
neighbor 20.8.8.8 activate
neighbor 20.8.8.8 send-community extended
neighbor 20.8.8.8 route-map rewriteNH in
exit-address-family
!
address-family ipv4 vrf extranet-customer
redistribute connected
no auto-summary
no synchronization
exit-address-family
!
address-family ipv4 vrf l3tunnel
no auto-summary
no synchronization
exit-address-family
!
ip classless
ip route vrf l3tunnel 0.0.0.0 0.0.0.0 Tunnel6
!
!
access-list 6 permit any
route-map rewriteNH permit 10
match ip address 6
set ip next-hop in-vrf l3tunnel

```

Alternative Solutions for Extranet Services

Following is a brief look at three alternative solutions for offering extranet services: *VRF-Lite*, *Carrier Supporting Carrier (CSC)*, and *virtual LAN (VLAN)-based extranet*.

VRF-Lite. Network segmentation can be provided on the CE router using the VRF-Lite feature (also referred to as Multi-VRF CE), supported in Cisco IOS Software Release 12.2(20)EWA and higher. This feature enables the Cisco Catalyst 4500 Series Switch to support multiple VPN routing/forwarding instances in CE devices. With the VRF-Lite option, it is assumed that the enterprise might combine VRF-Lite with an existing MPLS VPN or Frame Relay service using multiple interfaces/subinterfaces to segment traffic.

With VRF-Lite, logical traffic separation can be extended from service provider PE router to enterprise CE router. However, both the service provider PE and CE must configure separate interfaces or subinterfaces, separate VRF contexts, and separate routing protocols on a per-extranet basis.

Carrier Supporting Carrier. CSC is an attractive solution for offering extranet services at the enterprise CE router but requires label transport service from the service provider. The label transport service assumes that the enterprise has a current MPLS VPN service and the service provider supports the CSC deployment. If CSC deployments increase, enterprises can benefit from extranet provisioning ease because they themselves act like the MPLS VPN provider.

VLAN-based extranet. The enterprise CE can configure separate VLANs for extranets but would have to manage inter-VLAN traffic using access control lists (ACLs) or policy-based routing (PBR). This becomes a less attractive solution, because once the extranet traffic reaches the Layer 3 global routing table, it becomes very difficult to control. Managing ACLs and PBR in large deployments is not an easy undertaking.

◆ ◆ ◆

Cisco offers many solutions for implementing inexpensive, reliable internetwork connections between enterprise global networks and their partners. This article has provided an overview of how enterprises can benefit from MPLS VPN solutions to provide extranet services, regardless of how the enterprises themselves get WAN connectivity from their service providers. ■

FURTHER READING

- IETF RFC 2547 "BGP/MPLS VPNs"
cisco.com/packet/171_4a1
- Designing MPLS Extensions for CE Routers
cisco.com/packet/171_4a2
- MPLS for VPNs
cisco.com/packet/171_4a3
- MPLS VPN Carrier Supporting Carrier
cisco.com/packet/171_4a4

Reader Tips

Packet thanks all of the readers who have submitted technical tips. Each quarter we receive many more tips than we have space to include. While every effort has been made to verify the following reader tips, *Packet* magazine and Cisco Systems cannot guarantee their accuracy or completeness, or be held responsible for their use.

Configuration

TIP Creating Access Control Lists

Our group creates access control lists (ACLs) with a source and destination MAC address as follows:

```
router(config)#access-list ?

<700-799>          48-bit MAC address access list

access-list 711 permit 0002.7df5.4001
0000.0000.0000
access-list 711 deny 0050.dab7.655c
0000.0000.0000
access-list 711 permit 0000.0000.0000
ffff.ffff.ffff
```

—Rodrigo J. Mastropietro, São Paulo, Brazil

TIP Including and Excluding Character Strings in Startup Configurations

In Cisco IOS Software Release 12.0 and later, a nice trick is to use the +, -, and / characters with the **sh run** and **sh start** commands to search for specific strings (or exclude strings) in a configuration. This is similar to “include | exclude” pipe-sentences, but is available in startup or “live” running configurations. (See the “Reader Tips” in *Packet*, Third Quarter 2003, at cisco.com/packet/171_4d1, and *Packet*, Second Quarter 2004, at cisco.com/packet/171_4d2.)

For example, to use **include** to search for “voice” in the current startup or running configuration, type:

```
Router# sh run | incl voice
voice-card 2
voice class permanent 1
no voice hpi capture buffer
voice-port 2/0/0
dial-peer voice 123456 voip
Router#
```

As with **include**, to exclude strings from output, pipe **sh run** with **exclude** and the string you want to filter (**sh run | excl voice**).

With the **sh run** or **sh start** commands, if your terminal is set to display one page at a time, when the

display stops scrolling at the bottom of the screen use the “+”, “-”, and “/” characters to search configuration lines. For example, type **+voice** to get the same output as **sh run | incl voice**:

```
+voice
voice-card 2
voice class permanent 1
no voice hpi capture buffer
voice-port 2/0/0
dial-peer voice 123456 voip
Router#
```

Type a hyphen (-) to skip lines that contain the string you want to exclude. You can add several strings to include or exclude, grouping them with the pipe (|) character. Do not allow spaces between the pipes (for example, **-voice|description|line**). This feature still supports regular expression matching (anchoring, etc.). The forward slash (/) character lets you find the first occurrence of the given string. It stops at the first occurrence of the string to search for the next line that contains the string, so to continue searching, type “/” and the string you want to find again.

—Gabriel Zicarelli, Grupo López León, Buenos Aires, Argentina

TIP Stacking the Cisco Catalyst 3750 Switch

In large-scale deployments, a consistent approach to switch numbering can help in reducing overall downtime due to scheduled or unscheduled maintenance. I have found the following basic, stacking-related considerations for the Cisco Catalyst 3750 Switch helpful.

A switch numbering scheme example: Stack members should always be numbered from top to bottom, e.g., in a three-unit stack. First switch (top) should be “switch 1,” second switch (middle) should be “switch 2,” and the third switch (bottom) should be “switch 3.” When staging the stack, switch power on sequence might affect the switch numbering as the default number of each new switch is set to “1.” Subsequent additions to the stack are allocated the next available number by the stack master. A switch may be renumbered using the following command:

```
cat3750(config)#switch 1 renumber 1
WARNING: Changing the switch number may result in a
configuration change for that switch. The interface
configuration associated with the old switch number
will remain as a provisioned configuration.
Do you want to continue?[confirm]
Changing Switch Number 1 to Switch Number 1
New Switch Number will be effective after next reboot
```

```

cat3750(config)#switch 2 renumber 2
WARNING: Changing the switch number may result in a
configuration change for that switch. The interface
configuration associated with the old switch number
will remain as a provisioned configuration.
Do you want to continue?[confirm]
Changing Switch Number 2 to Switch Number 2
New Switch Number will be effective after next reboot

```

```

cat3750(config)#switch 3 renumber 3
WARNING: Changing the switch number may result in a
configuration change for that switch. The interface
configuration associated with the old switch number
will remain as a provisioned configuration.
Do you want to continue?[confirm]
Changing Switch Number 3 to Switch Number 3
New Switch Number will be effective after next reboot
cat3750#show switch neighbors

```

Switch #	Port 1	Port 2
1	2	3
2	3	1
3	1	2

Interfaces for each switch in the stack would be numbered as per switch number (also referred as interface of a slot number):

```

interface FastEthernet3/0/12
 switchport access vlan 999
 shutdown
interface GigabitEthernet1/0/2
 description IXIA Card 9
 switchport access vlan 60
 speed nonegotiate

```

If you need to upgrade software, do not use the **copy tftp: flash:** command. Instead you can use the following procedure (steps 1 through 4).

Editor's note: The new Cisco Network Assistant greatly simplifies the software upgrade process. It is available for free at cisco.com/go/networkassistant. Also, starting in Cisco IOS Software Release 12.2(25)SE, a new switch device manager allows easy software upgrade via an intuitive Web interface.

Step 1: Verify the flash memory space

```

cat3750#dir
Directory of flash:/
2 -rwx 736 Mar 1 1993 00:00:51 +00:00 vlan.dat
10 drwx 192 Mar 1 1993 00:14:41 +00:00 c3750-i9-
mz.121-19.EA1c
6 -rwx 10018 Aug 12 2004 19:58:47 +00:00 config.text
3 -rwx 1543 Aug 12 2004 19:58:47 +00:00 private-
config.text

```

```

4 drwx 192 Jul 22 2004 15:27:05 +00:00 c3750-i9k91-
mz.122-20.SE
15998976 bytes total (1488896 bytes free)
cat3750#

```

Editor's note: The above output shows the flash memory space on the master stack switch. To verify the flash memory space on the slave switches, use the **dir flash[switch#1]:** global command. For example, for switch #3 in the stack, the command would be **dir flash3:**.

Step 2: Note the existing Cisco IOS Software Release show version command

```

cat3750#sh ver | b Switch
Switch Ports  Mdel  SW Version  SW Image
-----
1 26 WS-C3750-24TS 12.2(20)SE C3750-I9K91-M
*2 26 WS-C3750-24TS 12.2(20)SE C3750-I9K91-M
3 26 WS-C3750-24TS 12.2(20)SE C3750-I9K91-M
< . . .output truncated. . .>

```

Step 3: Set up a TFTP server with new image file with a .tar extension (e.g., c3750-i5k91-tar.122-20.SE.tar)

Step 4: Start image download process (preferably using a terminal server)

Enough space on flash memory:

```

cr1162-ch8#archive download-sw /leave-old-sw
tftp://146.180.60.42/c3750-i5k91-tar.122-20.SE.tar
Loading c3750-i5k91-tar.122-20.SE.tar from
146.180.60.42 (via Port-channel 7):
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
< . . .output truncated. . .>

```

Not enough space on flash memory. You might need to delete old images files and remove corresponding directories:

```

cr1162-ch8#archive download-sw /overwrite
tftp://146.180.60.42/c3750-i5k91-tar.122-20.SE.tar
Loading c3750-i5k91-tar.122-20.SE.tar from
146.180.60.42 (via Port-channel 7):
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
< . . .output truncated. . .>

```

After downloading, the .tar file switch will expand it automatically and create requisite directories. In the case of a multi-unit stack, the stack master automatically upgrades remaining stack members with new image. Once installation is complete, the following message should be displayed by the switch:

```

Installing (renaming):
  flash: update/c3750-i9k91-mz.122-20.SE' ->
  flash: c3750-i9k91-mz.122-20.SE'
New software image installed in flash:
  c3750-i9k91-mz.122-20.SE
Installing (renaming):
  flash2: update/c3750-i9k91-mz.122-20.SE' ->
  flash2: c3750-i9k91-mz.122-20.SE'
New software image installed in flash2:
  c3750-i9k91-mz.122-20.SE
Installing (renaming):
  flash3: update/c3750-i9k91-mz.122-20.SE' ->
  flash3: c3750-i9k91-mz.122-20.SE'
New software image installed in flash3:
  c3750-i9k91-mz.122-20.SE

```

All software images installed

—Aamer Kaleem, CCIE No. 11443, UBS AG, Chicago, Illinois, USA

Troubleshooting

TIP Clearing a CLI Session on a Router

Sometimes it is impossible to log into a router because all the vty lines are engaged in a command-line interface (CLI) session. You can clear one CLI session if you have Simple Network Management Protocol (SNMP) write access on the device. The following command clears vty 0 on “router” with the community-string “private”: **snmpset router private .1.3.6.1.4.1.9.2.9.10.0 i 0**.

This simple UNIX script enables you to clear a session so that your next Telnet attempt will be successful:

```

#!/bin/zsh
if [[ $1 = "" || $2 = "" || $3 = "" ]]
then
  echo ""
  echo "\tUsage: $0 IP_address community-string
vty_line_number\n\n"
  exit
fi
snmpset $1 $2 .1.3.6.1.4.1.9.2.9.10.0 i $3

```

—Alain Moretti, Transpac USC du Midi, Toulouse, France

TIP Troubleshooting LAN Switch Ports on the Cisco Catalyst Platform

To properly troubleshoot, isolate, and fix LAN switch port issues for Cisco Catalyst 3750, 3550, 2970, 2950/2955, 2940, and 2900/3500 XL switches running Cisco IOS Software, use the command:

```
sh controller Ethernet-controller <interface-number>
```

This command allows you to demarcate and fix known common LAN device issues and identify whether the

problem is with the cabling or switch port, or whether it is caused by a faulty network interface card.

If the FCS error fields increment when you use the command, you probably have a physical layer problem and need to change the media. Incrementing alignment errors indicate that the switch port has not processed an even number of frames and is unaligned. This is a network interface issue and needs to be solved by a LAN administrator or help-desk technician.

```
IOSwitch01#sh controller Ethernet-controller fa0/1
```

Transmit	Receive
5755681 Bytes	1335670 Bytes
78090 Frames	16070 Frames
75520 Multicast frames	19456789 FCS errors
1588 Broadcast frames	6011 Multicast
frames	
0 Pause frames	34 Broadcast
frames	
0 Single defer frames	0 Control frames
0 Multiple defer frames	0 Pause frames
0 1 collision frames	0 Unknown opcode
frames	
0 2-15 collisions	1000 Alignment
errors	
1 Late collisions	0 Length out of
range	
0 Excessive collisions	0 Symbol error
frames	
0 Total collisions	8 False carrier
errors	
0 Control frames	0 Valid frames,
too small	
0 VLAN discard frames	0 Valid frames,
too large	
0 Too old frames	0 Invalid
frames, too small	
72639 Tagged frames	0 Invalid
frames, too large	
1 Aborted Tx frames	0 Discarded
frames	
Transmit and Receive	
5017 Minimum size frames	
87837 65 to 127 byte frames	
1030 128 to 255 byte frames	
265 256 to 511 byte frames	
0 512 to 1023 byte frames	
10 1024 to 1518 byte frames	
0 1519 to 1522 byte frames	

—Major J. Ward III, Hollis, New York, USA

Editor’s note: FCS and alignment errors are two of the most common errors seen on an Ethernet port. FCS (Frame Check Sequence) errors are the number of times that an Ethernet frame was received by the switch port from the attached device and the cyclic redundancy check (CRC) for the frame was not correct. The FCS check is used to detect corruption in

the frame. A frame with an FCS error is dropped by the port that received it on the switch, and the FCS counter is incremented for that port.

Alignment errors are the number of times that an Ethernet frame was received by the switch port from the attached device and the frame was not byte-aligned and had a bad FCS. All frames should end on an 8-bit boundary (1 byte = 8 bits); otherwise the frame is dropped by the port that received it on the switch, and the alignment error counter is incremented for that port.

FCS and alignment errors usually indicate a physical problem (cabling, bad port, NIC card, etc.) but can also indicate a duplex mismatch between the port and the attached device. When the cable is first connected to the port, some of these errors might occur. Also, for a switch port configured for half-duplex operation, Ethernet collisions can result in some of these errors being seen by that port.

For descriptions of other counters in the output of `show controller ethernet-controller`, visit cisco.com/packet/171_4d3.

For Catalyst 6000 and 4000 series IOS-based switches, you can use `show interfaces counters errors` to track FCS (CRC) and alignment errors seen by their ports. This command is also supported on other IOS-based Cisco

Catalyst switches running Cisco IOS Software Release 12.1 or higher.

TIP Troubleshooting Firewall Connectivity Problems

When connectivity problems occur with customer-managed firewalls, customers often assume the problem is the service provider's network. To find the actual problem, we obtain a list of allowable IP, TCP, or UDP ports in customers' firewalls.

From the remote router, we use the IOS command `telnet x.x.x.x port_number` to establish a session with the destination server. The Telnet shows an "Open" if it is connected; if not it just hangs. This helps us prove to our customers that client workstations are possibly accessing TCP ports that are blocked by their own firewalls. You can set the source IP address of your IOS Telnet session by typing `ip telnet source-interface Fa0`.

—Choy Wai Yew, AT&T Singapore, Singapore

SUBMIT A TIP

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

Tech Tips

Configure VPN load balancing on the Cisco Content Switching Module (CSM) in Directed Mode. Using virtual private network (VPN) load balancing you can intelligently distribute VPN sessions along a set of VPN concentrators or VPN headend devices. This configuration example takes you through the process step by step. cisco.com/packet/171_4e1

Create a certificate signing request on the Cisco SSL Services Module. Find out how to create a certificate signing request (CSR) on the Secure Sockets Layer Module (SSLM) and import the certificate using copy and paste in privacy-enhanced mail (PEM) format. cisco.com/packet/171_4e2

Troubleshoot Address Resolution Protocol on the Cisco Content Switching Module. This new tech note, "Understanding CSM ARP Behavior," includes information on how CSM handles issues relating to ARP requests. cisco.com/packet/171_4e3

Configure and troubleshoot the CT3 on the Cisco AS5000 Series. This new document describes how to configure and troubleshoot port adapters, multichannel T3 (platforms such as the Cisco 7200 and Cisco 7500 series), and the channelized T3 Trunk Card (CT3) for the AS5800 and AS5400 series. cisco.com/packet/171_4e4

Recover Cisco Guard and Traffic Anomaly Detector passwords. Learn how to recover the password of the root user in a Cisco Guard or a Cisco Traffic Anomaly Detector Distributed Denial of Service (DDoS) mitigation appliance with this new document. cisco.com/packet/171_4e5

Troubleshoot Cisco Unity outbound fax service. The Cisco Unity IP Fax Configuration Wizard enables you to use e-mail to send faxes over the PSTN. Refer to this new document to troubleshoot problems. cisco.com/packet/171_4e6

Looming Security Challenges

Zombies, Trojans, “bots” and worms: What have we wrought?

By David Barry

In June 2004, a large network of zombified PCs, also known as robots or “bots,” attacked Google, Yahoo, and other major Web-sites, blocking access to those sites for two hours. Security experts were able to identify the bot network, or “botnet” that appeared to be operating and managed to shut it down, stopping the attack. However, the attack was just one in what *USA Today* recently described as “wave after wave of infectious programs [that] have saturated the Internet, causing the number of PCs hijacked by hackers and turned into so-called zombies to soar into the millions.” (usatoday.com, September 8, 2004)

Zombie computers are present-day technical versions of the mindless corpses that rose from the grave to terrorize the living in the horror movies of the 1960s. In 2005, these zombies operate in cyberspace, proliferating across both private networks and the Internet. Botnets are a prime example of the power and complexity of today’s security threats.

Rogue developers create such threats by using worms, viruses, or application-embedded attacks. With botnets, for example, rogue developers can use worms or application-embedded attacks, that is an attack that is hidden within application traffic such as web traffic or peer-to-peer shared files, to deposit “Trojans.” Trojans are small executable programs that are left on a user’s computer (see sidebar, “The Trojan Horse: An Old Concept Revisited,” in this article). When an unsuspecting user logs on to the Internet (which happens automatically on a cable modem or DSL connection), the bots log into a server to await commands from the “zombie master.” Similar to what occurred in the June 2004 incident, hackers can launch virus attacks that deposit Trojans on thousands of computers, unbeknownst to the computer owners. A zombie master can then use these applications to flood a particular site with packets in a distributed denial of service (DDoS) attack or to generate large amounts of spam (see figure on page 21).

According to a recent report on Internet threats by Symantec, more than 30,000 computers are “recruited” into botnets every day.

“Botnets illustrate just how complicated and distributed the network threat environment has become,” says Scott Pope, a product marketing manager in the Security Technology Group at Cisco. “And, unfortunately, the situation continues to worsen as hackers have grown more sophisticated and creative in the attacks they generate.”

This combination of attack techniques—a virus or worm used to deposit a Trojan, for example—is relatively new and is known as a *blended* attack. A blended attack can also occur in phases: an initial attack of a virus with a Trojan that might open up an unsecured port on a computer, disable an access control list (ACL), or



Riccardo Stampatori

disarm antivirus software, with the goal of a more devastating attack to follow soon after.

In its 2004 semi-annual *Internet Security Threat Report* (cisco.com/packet/171_5a1), Symantec’s analysis of malicious code—worms, viruses, Trojans, backdoors, and blended threats—indicates that *malware* is increasingly being designed to steal personal data, particularly financial information and passwords. This data-stealing trend contributes to making all firms—but particularly banks and e-commerce companies handling payment transactions over the Internet—ever more vulnerable to compromise.

Evolving Security Landscape

Changes in network architectures and evolving threats create new security challenges. As well, the concept of the network perimeter is changing. In the past, users could only access the network through a few ingress or egress points—usually where the Internet connected to the enterprise network. Enterprises stacked security at the Internet perimeter using firewalls and intrusion detection systems (IDS).

By contrast, many more means of gaining entry to the network exist today. The perimeter has been extended and distributed, so security must be applied at each of these new ingress and egress points to avoid damaging threats, thus complicating security

Continued on page 21

Security Challenges, Continued from page 19

architectures. Virtual private networks (VPNs), for example, allow enterprise users remote access to the corporate network and are much more widely used than just a few years ago. Whereas previously enterprises might have insisted that VPN software run on a specific enterprise-configured computer, today users run VPNs from their own PCs or even from kiosks at copy centers or other businesses. This phenomenon allows many more entryways to the enterprise network and presents a significant challenge to IT departments. Is the computer equipped with virus protection? Is the virus software current? Did a worm become embedded in the computer?

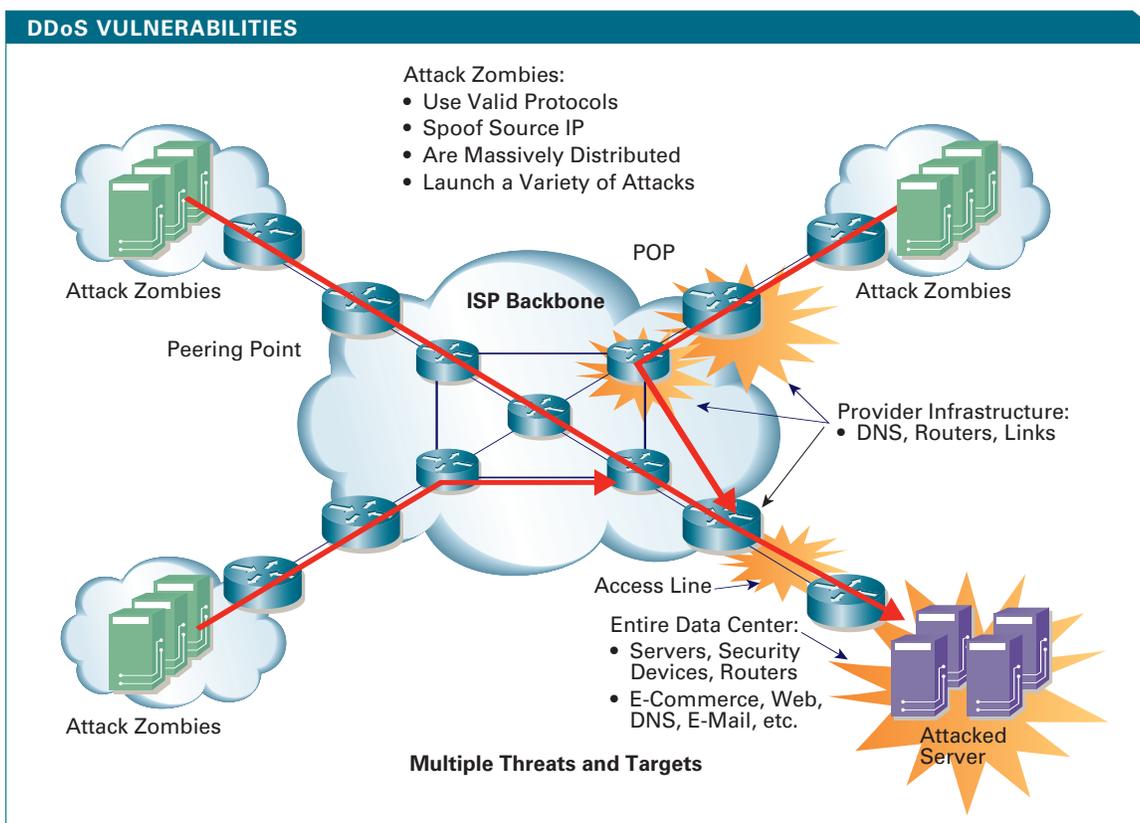
Wireless LANs (WLANs) pose additional security challenges. Users operating on an unsecured wireless network at a local coffee house may be unaware that a rogue PC, also using the same wireless subnet, is depositing a virus on the PC. When that PC is later docked into the corporate network, the virus could gain entry to the network.

At the same time, as the network is becoming more vulnerable to attack because of the expanding number of ingress and egress points, the threats themselves are changing. In addition to Trojans and botnets, newer, even more dangerous threats lurk. Two of the most troublesome are *flash threats* and *self-mutating worms*.

Flash threats are so named because of the speed with which viruses or worms can spread. In 1999, a virus dubbed “Melissa,” one of the earliest and most widespread viruses at the time, took 16 hours to spread globally, according to Network Associates Inc. In January 2003, the Slammer virus managed to infect more than 90 percent of the vulnerable hosts worldwide within 10 minutes using a well-known vulnerability in Microsoft’s SQL Server. New viruses in the coming months and years are expected to spread even faster. According to Pope, “It may be possible that a new type of virus will be able to infect millions of hosts within 60 seconds. So whatever defenses we create must be able to identify the threat and respond much more quickly than ever before.”

The other looming threat is the self-mutating worm. Today’s worms are relatively unintelligent. They are programmed to follow a specific set of instructions, such as to infiltrate one machine through a specific port and once on the machine compromise it in some way, for example, causing a buffer overflow and planting a Trojan. If anything interferes with these planned instructions, the worm lacks the ability to adjust and dies.

Now, however, rogue developers are adding intelligence and logic to worms so that if they can’t complete a specific task worms can mutate and pursue other lines of attack.



ZOMBIE ATTACK Botnets are a prime example of the power and complexity of the security threats prevalent today.

The Trojan Horse: An Old Concept Revisited

In network security parlance, a Trojan horse is a destructive program that masquerades as a benign application. Unlike viruses, Trojan horses do not replicate themselves but they can be just as destructive. One of the most insidious types of Trojan horse is a program that claims to rid your computer of viruses but instead introduces viruses onto your computer.

The term comes from the Greek story of the Trojan War, in which the Greeks give a giant wooden horse to their foes in the city of Troy, the Trojans, ostensibly as a peace offering. But after the Trojans drag the horse inside their city walls, Greek soldiers sneak out of the horse's hollow belly and open the city gates, allowing their compatriots to pour in and capture Troy.

Source: webopedia.com

“The security dilemma is like Moore’s Law in reverse,” says Pope. “Whereas Moore’s Law postulates that processor performance will double every 18 months while costs will decline dramatically, security is moving in the opposite direction—networks are becoming less secure while the cost to defend them is increasing.”

This prognosis is supported by mi2g (mi2g.org), a research firm in the UK that specializes in computer security. Mi2g reports that the economic damage from malevolent network security attacks reached somewhere between US\$157 billion and US\$192 billion worldwide in 2004.

Combating New Threats

The current security defense paradigm is to deploy more and more of the existing security technologies throughout every segment of the network. This includes firewalls and ACLs to block access and perform application inspection, intrusion protection system (IPS) technology to provide very granular traffic inspection and identify known threats, encryption software to counter eavesdropping, anomaly detection to detect worms or DoS attacks, and antivirus software to battle viruses. Many of today’s security technologies were developed to perform their specific function with little context of the overall network threat environment. Operating alone, however, these technologies are less effective in stopping the newer attacks, as well as the changing ways in which users access networks, because of the “security gaps” that exist between each technique’s capability.

With the increased complexity of threats, such as the blended threats that use a combination of techniques to disrupt networks, security technologies must operate in a coordinated fashion to stop attacks and better control network activity and applications.

Unfortunately, over the years, many companies have addressed nagging security concerns by constantly adding devices and software to address each particular problem. This has led to separate antivirus protection, firewalls, VPNs, and intrusion prevention.

While addressing short-term needs, this approach creates an entirely new and bigger problem: managing multiple systems that operate independently of one another. As more advanced threats emerge, Pope and many others believe that network security must become more holistic; security technologies must act in coordination to detect and defend against more sophisticated threats.

“There is a growing need for devices that can assemble the puzzle pieces and lock down the gaps that exist in conventional network security systems,” Pope explains. “Today, a huge problem exists with the misclassification of threats and organizations taking inappropriate action, or even worse, missing threats altogether.”

Adaptive Security for a Changing World

Transforming chaos into clear and manageable security policy is essential, which is why future network security systems need to focus on convergence and consolidation. In network security, a proactive approach is critical. The idea is to accurately identify and stop attacks as early and as far from the destination host as possible, while simultaneously simplifying the security architectures required to do this. Converging numerous security functions into a single adaptive device or system enables these combined functions to operate as a coordinated defense (instead of silos) that stops a broader range of attacks and greatly reduces the number of diverse devices that must be deployed, thereby simplifying security design and management.

Historically, firewalls have generally been considered fairly simple devices, but they are effective at what they do: either block a packet or let it through based on Layer 3 and Layer 4 information and session state. They can provide some level of application inspection but do not perform the detailed inspection of some other technologies. An IPS device can pick up where a traditional firewall leaves off by peering more deeply into a packet’s contents to see whether the data within conforms to company policy. But IPS devices lack the breadth of mitigation actions and resilience of a firewall that network security administrators require. Combined, however, a firewall and an IPS device can be more effective than either one by

Continued on page 87

Fixed-Mobile Convergence

Bringing Wireless LAN and Wireless WAN Radio Access Technologies Together for a Seamless Service Offering

By Steve Hratko

It's a widely held belief that most network access eventually will be untethered. Personal digital assistants (PDAs), cell phones, laptops, and IP-based appliances will all make use of radio technology for connecting to the network. The best solution will stem from a combination of wireless LAN radios based on IEEE 802.11 technologies for high performance in a local area, such as a hotel, airplane, or office building, and a mobile radio service for ubiquitous access when on the move (wireless WAN). This common, seamless blending of wireless LAN/broadband and mobile phone services is called *fixed-mobile convergence*.

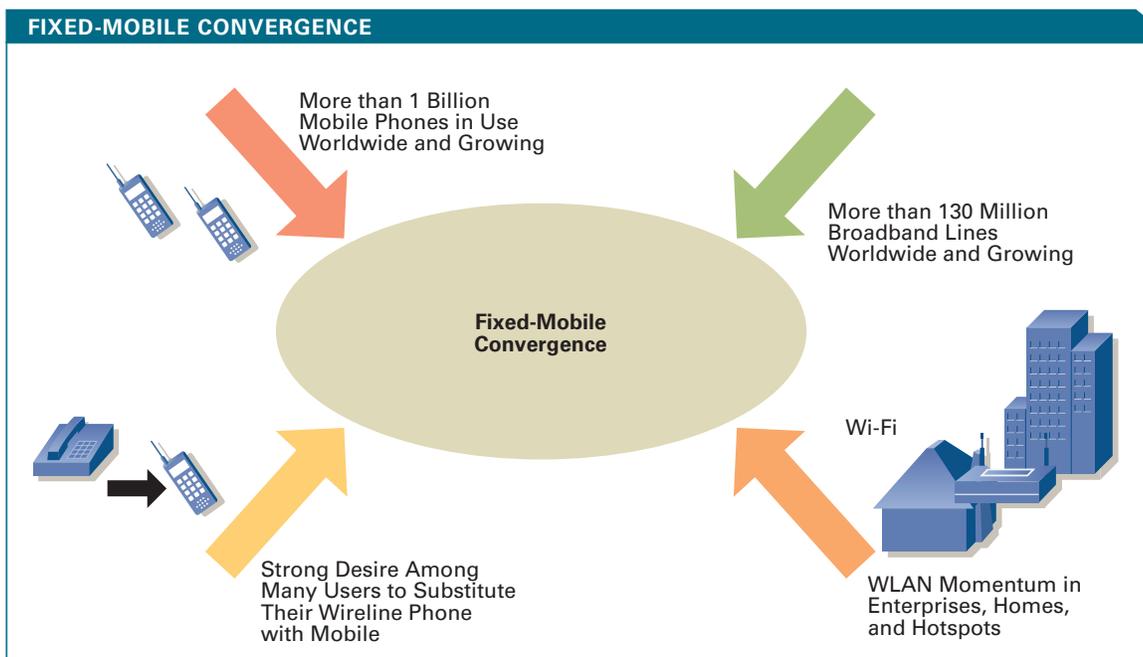
The mobile radio options include CDMA2000 and its high-speed data overlay EV-DO (evolution-data optimized), and GSM/UMTS with its high-speed data overlay HSDPA (high speed downlink packet access). In addition to these options, there are two new mobile technologies coming out of the IEEE that will soon make their way to the market. They are known as 802.20 and 802.16 (WiMAX).

There are a number of other mobile data technologies in use, including General Packet Radio Service (GPRS), EDGE, CDMA 1x, and UMTS, which carry

voice and data together on the same RF carrier; however, this is not the most advantageous data solution in terms of spectral efficiency, because the needs of voice are very different from the needs of data. The new high-performance radio options put voice on one RF carrier and data on another.

The benefits of fixed-mobile service convergence for end users are better connectivity by always utilizing the best available radio signal for that time and place. This approach is especially beneficial when using a mobile phone in-building, because that is where Wi-Fi is at its best and where mobile signals can sometimes be at their weakest.

Among the benefits of fixed-mobile service convergence for mobile operators is that it allows them to pick up a lot of the in-building minutes that often go to wireline operators, thereby accelerating fixed-mobile substitution. It also provides a very strong response to the voice-over-broadband providers that are showing up in all geographies. These providers use voice over IP (VoIP) over broadband to pick up the in-building minutes, and they do it at a much lower price point than do traditional operators. The



A TELECOM PERFECT STORM Several factors on the mobility and broadband fronts are coming together to fuel the trend toward fixed-mobile convergence.

dual-mode phone is a great response to this threat because it provides voice over broadband with full mobility. The bottom line for operators is a much better user experience, which should improve their competitive position and reduce churn.

Stitching Together Different Radio Access Technologies

Mobile operators will use a variety of approaches to allow the different radio access technologies to cooperate in delivering a converged service offering. These include:

- *Unlicensed Mobile Access (UMA)* is a new Layer 2 technology ideally suited to offering seamless voice services using the mobile operator's mobile switching center (MSC) for call control over a GSM or broadband IP access network. With a broadband IP network, the GSM voice signaling and bearer is tunneled across the IP network and back to the mobile operator's domain. As the user moves between Wi-Fi and GSM coverage, the network will seamlessly hand off the call. With a properly engineered IP network, the user experiences no service degradation. This technology is especially suited to deployments where cellular coverage needs to be supplemented with in-building Wi-Fi coverage, and could start to emerge in the next 12 to 18 months.
- *Mobile IP* is a Layer 3 technology ideally suited to laptop-based data services (no voice call control required). Operators such as Swisscom Mobile (swisscom-mobile.ch) are already offering services that make use of tri-mode PCMCIA cards that support Wi-Fi, GPRS, or UMTS radios for connectivity. The laptop selects the best available radio signal, and Mobile IP enables seamless handoff as the user moves across different radio coverage areas. The strength of Mobile IP is that it allows the laptop to keep its IP address as the user moves about. This technology is readily available, and deployments exist in many parts of the world.
- *Session Initiation Protocol (SIP)* application-layer mobility is the future direction for many of the world's service providers. It will enable support for real-time multimedia service in the all-IP world of intelligent endpoints. One of the great advantages of using SIP for service convergence is that it allows the user to transfer an application session between devices. For example, a session could be initiated on a laptop in the home, passed to a PDA as the user leaves home and gets into his car, and passed back to the laptop when the user reaches the office. SIP's requirement for real-time, IP-based multimedia does necessitate substantial investment in the public wireless networks and so will likely roll out over many years.

Stitching together a combination of different radio access technologies into a seamless service offering is

technically challenging, and so it is worthwhile to take a look at just why this undertaking is necessary.

Wireless LANs

The dominant technology for wireless LANs is based on the IEEE 802.11 standard. This technology is widely deployed, cost effective, fast, and uses unlicensed spectrum. The latter has significant implications on how the technology can be deployed—and unlicensed doesn't mean unregulated. The use of unlicensed bands puts limitations on the amount of power that an 802.11 radio can emit. Higher power output risks interfering with other users of that band.

As such, this technology is primarily being used to support wireless LAN services. These services can involve a single access point in a coffee shop or a large number of access points to cover an airport or a hot zone in a downtown area.

Wireless LANs will typically have a performance advantage over mobile services. The reasons for this include the simple physics of radio waves. RF signal strength drops off as the square of the distance (and even faster in some instances). Therefore, the closer users are to the access point the stronger the signal and the higher the performance of the service. Public wireless LANs are usually found in heavily trafficked areas (hotels, airports, and convention centers), and they only need to propagate a signal a few hundred feet. A mobile wireless service must be able to propagate signals over many tens of kilometers.

Another advantage wireless LANs have over mobile services is in the channel width of the carrier. Wireless LANs operate in higher frequency bands and use wider channels. Today's systems use 20-MHz channels, and the future will most likely include both wider and narrower channel options. Wider channels support much higher data rates. While the higher frequency bands don't penetrate structures as well as cellular bands do, this can be an advantage when using unlicensed frequencies because it helps limit interference.

Wireless WANs

Conversely, wireless WAN (mobile) systems operate in lower frequency bands and with narrower channels. Narrower channels mean lower data rates, primarily due to the economics of RF spectrum. Lower



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

frequency spectrum is much more valuable than higher frequencies and is thus auctioned off in narrower channel widths.

For mobile applications, the optimal frequencies are below 1 GHz. In fact, various parts of the world have enjoyed success with mobile services operating at 450 MHz, where one cell tower can cover the same area as more than a dozen towers operating at 1.9 GHz (there is a lot of variability here depending on terrain). In addition to greater propagation ranges, the lower frequencies can also pass through structures more effectively to reach users deep inside buildings. This is very important to operators and subscribers who want ubiquitous connectivity.

Operators of mobile wireless networks face a challenging decision regarding how to evolve their networks to support high-speed data services. It is anticipated that these data-oriented networks will be implemented as overlay networks using dedicated RF spectrum, e.g., using RF carriers devoted to HSDPA, EV-DO, or IEEE 802.16. If high-speed data were made to share RF spectrum with voice, this could degrade the spectrum's ability to support business-critical voice traffic.

All of the data-oriented mobile wireless technologies (HSDPA, EV-DO, 802.20, and 802.16) will offer somewhat similar performance on a bit/sec/Hz basis. As a general rule, users can expect to see about 500 to 600 Kbit/s on the downlink and 100 to 200 Kbit/s on the uplink. The numbers will vary widely depending on such variables as distance from the cell tower, loading on the tower, terrain, and user movement/speed. These rates are only rough approximations and will improve as the technology evolves.

A Closer Look at 802.16 for Mobile Wireless

Of the various mobile wireless technologies, IEEE 802.16 (also known as WiMAX) has been getting significant attention—a result of successful marketing by the WiMAX Forum, strong support from merchant chip vendors such as Intel, and participation from all of the major Radio Access Network (RAN) vendors.

WiMAX originated as a fixed wireless technology that could be used in microwave backhaul applications as well as for fixed wireless access. WiMAX has recently begun adding support for mobility.

The primary advantages of WiMAX-based solutions include the following:

- Very reasonable intellectual property rights licensing that comes with the IEEE's reasonable and non-discriminatory (RAND) licensing policy
- A strong marketing organization (WiMAX Forum) devoted to promoting the technology

- Intel's support, which should translate into inexpensive mobile client devices as WiMAX technology is integrated into laptops and PDAs

This last point is worth emphasizing, because the cost of the mobile client device is a considerable part of the cost of a mobile wireless service and often must be subsidized by the mobile operator.

The primary disadvantage of WiMAX technology is that true standards-based mobile network deployments will probably not occur until at least 2006. In the meantime, solutions based on EV-DO and HSDPA are becoming available. Markets for WiMAX technology include the following:

- *Microwave backhaul*—The genesis of WiMAX, much of this market centers on backhauling voice from cell towers, which is typically done at very high frequencies (>10 GHz) using line-of-sight radios. The vendors in this market all have proprietary solutions, and a standard should reduce costs.
- *Fixed wireless access*—This market is fairly small and focused on areas lacking DSL or cable service. Fixed wireless has had trouble competing with wired solutions, when these solutions are available. WiMAX has the potential to drive fixed wireless technology into laptops and PDAs and make it portable. Users can then enjoy a wireless DSL service that follows them as they move about. But this starts to sound a lot like a mobile service.
- *Huge mobile market*—Worldwide mobile operator capital expenditures exceed US\$80 billion per year, and the client device business is even bigger. If WiMAX can succeed as a high-speed mobile data overlay, it will drive the volumes that will help bring down the cost of the technology for all applications

WiMAX will not be a viable competitor to Wi-Fi in the LAN. It is a WAN technology, and it will compete with the other WAN technologies.

◆ ◆ ◆

Fixed-mobile convergence will be one of the dominant trends in the service-provider market over the next decade. It will involve the use of at least three different service convergence architectures (UMA, Mobile IP, and SIP) and a variety of radio technologies. IEEE 802.11 and extensions to this standard will be the dominant technology in the LAN. In the wireless WAN, a variety of different mobile solutions will be deployed, including HSDPA, EV-DO, and WiMAX. All three should do very well going forward.

You bet, it's going to become a wireless world. ■

Cisco MPLS VPN over IP

Extending MPLS VPN to Operate over IP Networks—With the Same Overall Architecture and Service Experience

By Santiago Alvarez

Multiprotocol Label Switching virtual private network (MPLS VPN) introduced a peer-to-peer model that enables large-scale IP VPN implementations. This model greatly simplifies routing and manageability for the VPN customer and the service provider while guaranteeing proper isolation between VPNs. To implement this model, MPLS VPN relies on extensions to existing IP routing protocols and an MPLS transport network. *Cisco MPLS VPN over IP*, supported in Cisco IOS Software Release 12.0(28)S and higher, reuses the same functionality as MPLS VPN, but replaces the MPLS transport with an IP transport. VPN traffic is carried by an IP tunnel instead of an MPLS Label Switched Path (LSP). This feature enables MPLS VPN services on IP networks that have not been enabled with MPLS.

Applications, Services, and Architecture

Cisco MPLS VPN over IP retains the same application and service characteristics whether implemented over an IP or MPLS backbone. For example, an organization can use this technology to segment an IP network to support different groups within its structure or to provide a private IP service to other parties. Such segmentation would support overlapping addresses and flexible traffic forwarding topologies. In another scenario, a network engineer could use this technology to build a centralized server infrastructure that can be shared across multiple VPN instances.

Cisco MPLS VPN over IP brings new application and service opportunities for providers and subscribers of VPN services. An example: A provider of MPLS VPN services can extend the reach of its offering to networks that are not MPLS-enabled using an inter-autonomous system configuration. Similarly, two providers can make peering agreements for their MPLS VPN services even if IP transport is used. As another example, MPLS VPN subscribers can partition their VPN service to create their own internal VPN services. This application brings increased flexibility in the implementation of hierarchical VPN configurations because little coordination is needed between subscriber and provider. Figure 1 shows two example applications using Cisco MPLS VPN over IP.

Cisco MPLS VPN over IP extends the original MPLS VPN architecture with a collection of multipoint IP tunnels and a separate address space. Each provider edge (PE) has one multipoint tunnel interface that connects the PE to all other PEs that are part of the VPN service. The tunnel forwards the VPN packets to the appropriate destination PE making VPN packets transparent to intermediate nodes. Each PE automatically discovers other PEs reachable through the tunnel (i.e., tunnel endpoints).

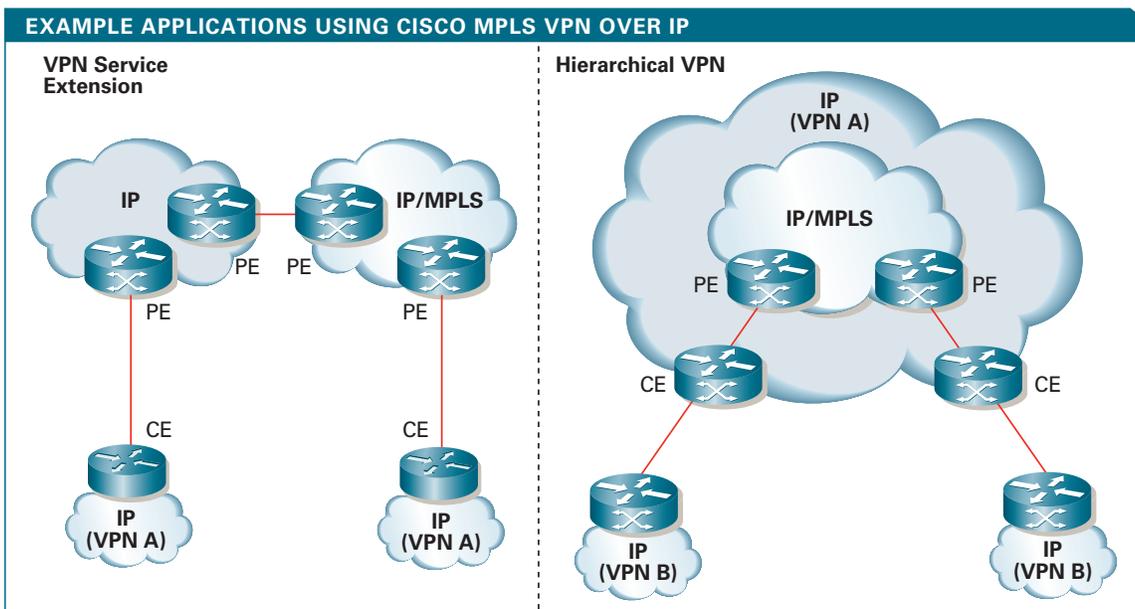


FIGURE 1 Shown are two example applications using Cisco MPLS VPN over IP: VPN service extension and hierarchical VPN.

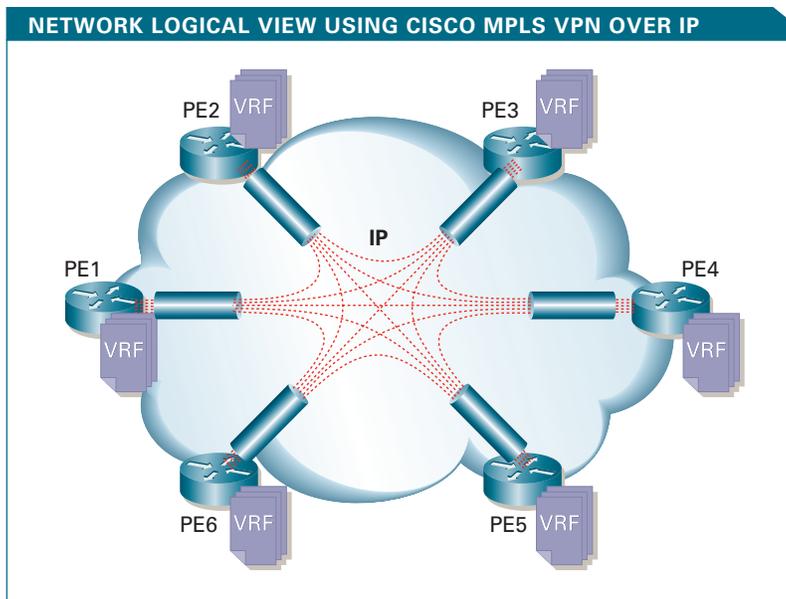


FIGURE 2 A logical view of a network using Cisco MPLS VPN over IP.

The PE discovery process uses simple extensions to Border Gateway Protocol (BGP) multiprotocol that build on top of the BGP extensions already present for any type of MPLS VPN. A separate address space for the multipoint tunnels provides isolation for the VPN traffic. This architecture retains the same scalability characteristics as a classic MPLS VPN service while it is extensible to multiple IP tunneling technologies. See Figure 2.

Traffic Forwarding and Encapsulation

The basics of Cisco MPLS VPN packet forwarding are independent of the backbone transport choice (MPLS or IP). In both cases, VPN traffic is kept separate within a PE by using Virtual Routing and Forwarding (VRF) instances for each VPN the PE supports. However, packet encapsulation varies according to the transport network in Cisco MPLS VPN. There are two encapsulation components when using an IP transport: a tunnel header and a VPN header. The tunnel header transports the packet to the egress PE while the VPN header identifies the appropriate VPN packet processing at that location.

The current implementation of Cisco MPLS VPN over IP uses Layer 2 Tunneling Protocol version 3 (L2TPv3) as the IP tunneling technology. The tunnel header uses the L2TPv3 Session ID field to identify the packet as an IP VPN packet requiring MPLS VPN processing, and uses the Cookie field to provide spoofing protection. The last part of the encapsulation, the VPN header, uses a VPN label identical to that used for MPLS VPNs over an MPLS transport. Figure 3 shows an encapsulation comparison of an MPLS VPN service being offered over different transports.

L2TPv3 provides intrinsic protection for the VPN traffic against external attacks. A malicious user might attempt to inject packets into a VPN by sending

VPN-encapsulated packets toward a PE. This type of attack is generally prevented when using an MPLS transport by rejecting MPLS packets coming from VPN subscribers on customer-facing access interfaces. When implementing MPLS VPN over an IP transport, PE devices are generally more vulnerable to IP spoofing attacks. Network boundaries or the PEs themselves require special configuration and extra processing (e.g., access control lists, or ACLs) to identify and block spoofed VPN packets. L2TPv3 thus provides its own spoofing protection as close to the customer as possible by incorporating strong spoofing protection directly within the PE. So, with or without IP ACLs, L2TPv3 prevents external spoofing attacks on a given VPN because every PE forwards packets using a pre-specified, cryptographically random 64-bit Cookie.

A successful blind spoofing attack against a deployment of Cisco MPLS VPN over IP would require more than 6000 years to accomplish at foreseeable attack rates (100 million pps). A malicious user would need to know the IP addresses of the ingress and egress PE, the L2TPv3 Session ID and Cookie, plus the VPN label, to inject traffic into a VPN from the outside. The effort required to guess a random (64-bit) L2TPv3 Cookie value rules out the possibility of a successful attack. The other fields provide marginal incremental protection beyond what the L2TPv3 Cookie offers, because they generally cannot be chosen in a cryptographically random manner and are not large enough to inhibit a determined attacker guessing values.

MPLS VPN could use other IP tunnel encapsulations such as MPLS over IP or Generic Routing Encapsulation (GRE) when a non-cryptographic solution is required. An MPLS VPN implementation using a plain IP tunneling encapsulation (MPLS over IP) would be the simplest but most susceptible. Assuming a malicious user has discovered the source and destination IP addresses for the PEs, this user only needs to guess a valid VPN label (20 bits). A breach could be expected in seconds even at low attack rates (thousands of pps).

A second alternative involves using the GRE protocol, which has a reserved key (32-bit) field that remains undefined. If this key were used in a manner similar to the L2TPv3 Cookie (e.g., defined to be filled with a cryptographically random value), it would still provide inadequate anti-spoofing protection. A breach would be possible in a matter of hours at relatively low attack rates (less than 12 hours at



SANTIAGO ALVAREZ, CCIE No. 3621, is a technical marketing engineer in Cisco's Internet Technologies Division and focuses on MPLS and QoS technologies. He has been a regular speaker at Networkers and a periodic contributor to *Packet*. He can be reached at saalvare@cisco.com.

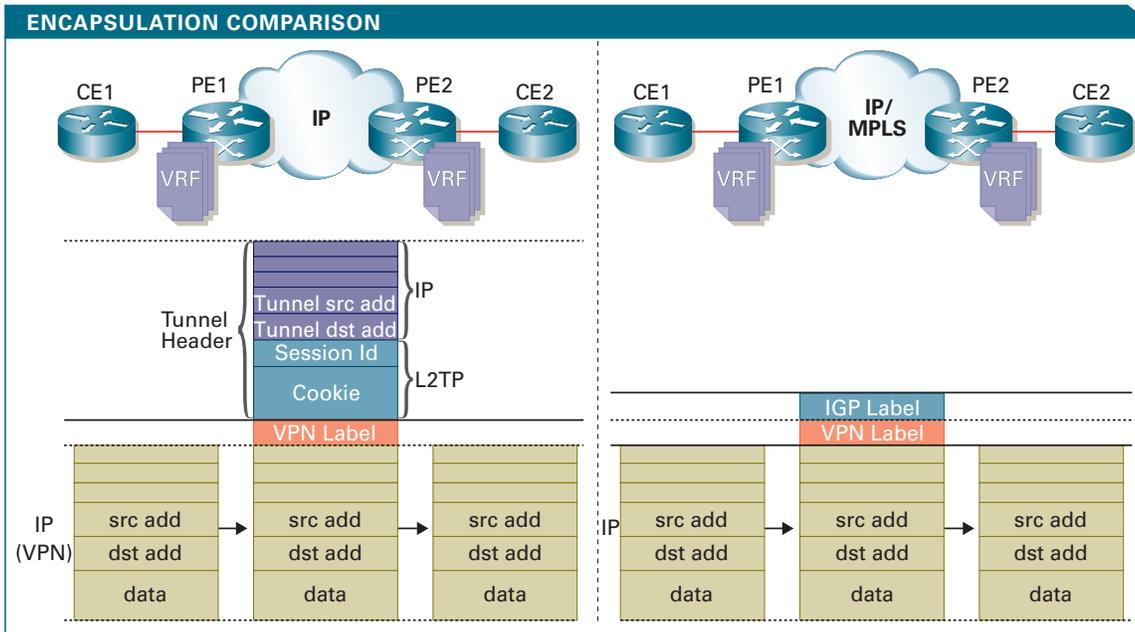


FIGURE 3 Side-by-side encapsulation comparison of an MPLS VPN service being offered over an IP network and over an MPLS network.

100,000 pps). As such, GRE would add little value in this application and bring unnecessary overhead with myriad optional fields to check and verify.

VPN Route Distribution, Tunnel Endpoint Discovery

MPLS VPN uses the same mechanisms for VPN route distribution regardless of the backbone transport (IP or MPLS). However, VPN route resolution operates somewhat differently in Cisco MPLS VPN over IP and MPLS VPN over MPLS. MPLS VPN, in general, requires that PEs perform a recursive route lookup on BGP next hops when an incoming VPNv4 BGP update is processed. When an MPLS backbone is present, the PE is expected to match the next hop with an existing LSP. When an IP backbone is present, the PE is expected to match the next hop with an existing tunnel endpoint.

A successful match selects the multipoint tunnel as the output interface for the packet. This process guarantees that the packet will be appropriately forwarded through the tunnel and with the correct encapsulation. To achieve proper resolution, the BGP next hop is resolved in the separate address space associated with the tunnel. Otherwise, resolution would be attempted against the global routing space in search of an LSP that would not exist.

Cisco MPLS VPN over IP provides automatic tunnel endpoint discovery and signaling of tunnel parameters. Each PE needs to know which other PEs (i.e., endpoints) are reachable via the multipoint tunnel before proper VPNv4 BGP next hop resolution can take place. In addition, each PE needs to know the L2TPv3 Session ID and Cookie that other PEs expect, so VPN packets can be encapsulated appropriately. Manual configuration of this information is not scalable; as the number of PEs increases, the simplicity

of the tunnel's multipoint nature is destroyed.

PEs take advantage of the existing Multiprotocol BGP (MP-BGP) infrastructure to distribute tunnel endpoint information. Cisco MPLS VPN over IP defines a new tunnel address family extension in MP-BGP. This address family is used to signal L2TPv3 tunnel address, Session ID, and Cookie. L2TPv3 is used exclusively as an encapsulation mechanism. The native L2TPv3 control plane is not employed. The operational and processing impact of this extension is marginal given that MP-BGP is already required to distribute VPNv4 route information. In contrast, when MPLS is used as transport, endpoint discovery is tied to VPNv4 advertisements and there is no signaling of encapsulation type (MPLS) and its (LSP) parameters. For more on the tunnel endpoint information learned via MP-BGP, visit cisco.com/packet/171_5c1.

◆ ◆ ◆

Cisco MPLS VPN over IP uses L2TPv3 as the IP tunneling technology, offering inherent anti-spoofing protection that IP alone and GRE both lack. Control plane operation is extended to support tunnel endpoint discovery and VPNv4 next hop resolution through a tunnel. With Cisco MPLS VPN over IP, MPLS VPNs can be deployed over any IP network in a scalable, secure manner today. ■

FURTHER READING

- Cisco MPLS VPN over IP Documentation cisco.com/packet/171_5c2
- BGP/MPLS IP VPNs over L2TPv3 IETF Draft cisco.com/packet/171_5c3
- Encapsulation of MPLS over L2TPv3 IETF Draft cisco.com/packet/171_5c4

HOW RESILIENT IS YOUR BUSINESS?

NETWORK STRATEGIES THAT ENHANCE OPERATIONS AS THEY PROTECT YOUR BUSINESS

By GAIL
MEREDITH
OTTESON

IF YOUR ENTERPRISE lacks a comprehensive strategy for using its IT systems to optimize business resilience, you're not alone. Business resilience includes strategies for business continuance—maintaining operations during and after a disruption—but also improves the organization's overall ability to do its business. It refers to the *operational and technological readiness* that prepares organizations to make daily operations efficient and cost-effective, respond quickly to opportunities with the potential to increase competitive advantage, and react appropriately to unplanned events.

Enterprises need agility to roll out new applications, react to market changes and competitive threats, support business processes, and communicate with employees, partners, suppliers, and customers. The IT resources many enterprises rely on today include not only the data center, but also the company's LAN infrastructure, storage area network (SAN), and the WAN that interconnects all locations, applications, and users. Successful business operation depends on the continuity of *all* of these systems. Because everything upon which a business depends is part of an interconnected system, the entire system must be resilient. Therefore, a business resilience strategy takes into account how IT systems interact with each other. "You can't just look at point products or single systems," says Glen Fisher, manager of the Enterprise Market Management Group at Cisco. "Business resilience is holistic. You have to consider how all IT systems interoperate to achieve your goals for business agility and continuity."

Cisco identifies six components in a business resilience strategy: *network resilience*, *applications resilience*, *communications resilience*, *workforce resilience*, *security*, and *network management*.

Network Resilience

As a strategic business asset, the network is the foundation for business activities and communications that translate to revenue. Network resilience is the result of deliberate design, implementation, and operational practices using an integrated architecture supported by lifecycle services to attain a flexible, secure infrastructure that maintains connectivity, optimizes network performance, and delivers intelligent services during ordinary and atypical circumstances.

Network resilience begins with a high-availability network, which integrates network domains and technologies into an interoperable system that automatically reroutes around failures and ensures consistent conformity to security policies.

Some enterprises deploy lowest-cost, point-product solutions from several vendors without realizing how that approach diminishes their ability to deploy future products and services, and without counting the cost of network downtime. Robbing the organization of the benefits of system-wide network intelligence can increase the complexity and expense of network operations.



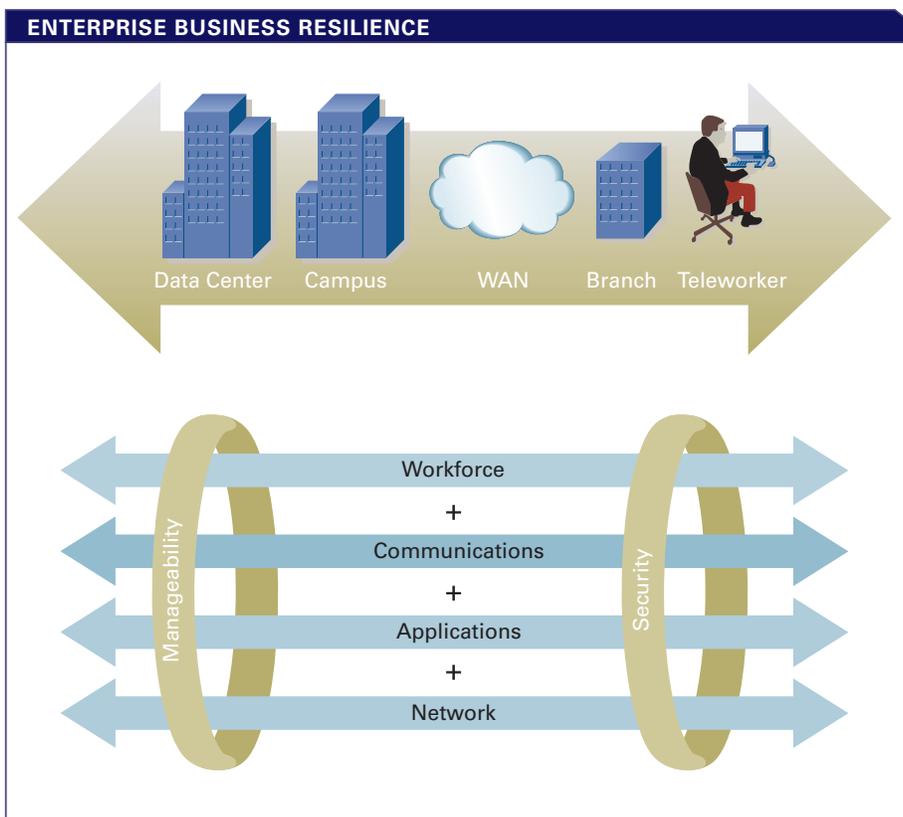
Today many more enterprises are deploying a resilient network, which provides flexibility to adapt an infrastructure to future services and applications with minimal disruption. One emerging technology is radio-frequency identification (RFID), which facilitates inventory management in the shipping and warehousing industries. RFID technology may be adapted to future applications in other industries. Enterprises that already have a resilient network can easily incorporate RFID systems when they need them. A resilient network has end-to-end intelligence that segments (through virtual LANs, virtual SANs, or WANs), prioritizes (through quality of service), and protects (through encryption) RFID traffic without requiring major upgrades, just configuration of existing features. To facilitate adaptation to emerging technologies, enterprises also need a strong relationship with a networking vendor that offers worldwide, world-class service and support. This relationship helps enterprises track and solve problems quickly with access to online training, documentation, and software upgrades, a 24-hour help desk, rapid sparring, and onsite assistance as required.

Applications Resilience

“Most companies might not fully understand the value of the network to their business,” explains Bobby Guhasarkar, senior solutions marketing manager in Enterprise Marketing at Cisco. “But they do understand the value of applications to their business.”

Vital to a business resilience strategy, business continuance systems maintain operations during and after unplanned events. While disruptions certainly result from major events such as earthquakes, flooding, or hurricanes, even more often they are brought on by a power outage or a car that won’t start.

While some applications might require data center mirroring to preserve transactions, others might have higher tolerances for data loss, application downtime, and user accessibility. “The level of business continuance you need depends on the function. Many companies try to implement a uniform business continuance plan. As a result, they overspend on some solutions and underspend on others,” says Zeus Karravala, vice president of Enterprise Infrastructure at The Yankee Group research firm.



BE PREPARED Business resilience integrates systems throughout the enterprise to achieve both market agility and rapid disaster recovery.

A resilient network increases applications resilience. Network intelligence complements server and storage technologies to maintain application availability. Offloading processor-intensive tasks such as encryption, compression, and load balancing into the network increases application resilience and scalability by freeing server and storage processors to perform their core duties.

The trend toward data center consolidation redefines how enterprises attain optimal applications resilience. It begins with redundant network components and server clusters. Cisco partnerships with leading server vendors prove interoperability between servers and network components with tested configurations. Server virtualization technologies protect applications, preventing faults that occur in one application from affecting the performance of others. For more information on data center resilience, see “Ready for Anything,” page 40.

Intelligent storage networking improves both availability and utilization of storage resources. “Acquisition-prone companies can end up with storage resources from

several vendors. Cisco applies networking concepts, such as virtual SANs to storage, allowing IT managers to consolidate multiple SAN islands from different vendors into one integrated fabric with enterprise-wide access,” says Jacqueline Ross, vice president for storage in Product and Technology Marketing at Cisco. “Virtualization technologies make it easier to match application service-level requirements with the appropriate class of storage.”

Regulatory compliance often requires enterprises to build geographically dispersed, redundant data centers. SAN extension technologies, such as Fibre Channel over IP (FCIP) and optical networks support synchronous and asynchronous transaction and data mirroring. Should a catastrophe take all or part of a data center off line, mirrored resources in a backup or hot-standby data center can take over business-critical activities without session loss.

“Regulatory compliance helps people implement things they should have done years ago,” says Karravala. “There isn’t

Continued on page 34

Business Resilience, Continued from page 32
a company I talk to that doesn't say business continuity is important, but when they actually have to spend the money, they always find something else to do instead, and they could get to it next year." A Gartner study found that 88 percent of enterprises were ready to deal with a power outage, but only 38 percent can adapt to a loss of transportation infrastructure, and a scant 13 percent can accommodate a major worker disruption.

Communications Resilience

"Many enterprises have yet to take full advantage of IP communications," says Rick Moran, vice president for IP communications in Product and Technology Marketing at Cisco. "They install an IP PBX [private branch exchange] but stop there. However, IP voice applications can be part of restructuring the way they increase agility. It's easy and inexpensive to do over an IP network."

Converged networks dramatically enhance communications flexibility for both daily operations and disaster recovery. Features such as extension mobility allow employees to use their own phone numbers from any IP phone in the global enterprise network. Unified communications simplify message retrieval by combining voice mail and e-mail into one service that employees can access through a computer or telephone. IP videoconferencing and IP video telephony reduce the need for travel.

Telephone service is the IT function most essential to conducting business; therefore, the network that supports it must be highly available. Branch-office routers must include features that maintain local telephone service and PSTN access should the WAN link to the central management service fail. IP call centers should allow remote agent access from home when weather conditions make commuting difficult or dangerous.

Workforce Resilience

Closely aligned to applications and communications resilience is workforce resilience, which strives for anytime employee accessibility to applications and services from any location. In the campus, conference rooms with wireless LAN access eliminate the traditional "battle for ports." Road warriors can carry a preconfigured broadband router and IP phone kit

that allows them to connect to the corporate network from a hotel room and enjoy the same services they would have access to if they were directly connected at the campus. Teleworkers use a similar means to set up home offices with secure, always-on VPN access to the corporate network. Such flexibility increases employment options, worker satisfaction, and productivity.

Multilayer Security

The most highly available and intelligent network isn't resilient without adequate protection. "An integrated approach to resilience makes it easier to apply and enforce consistent security policies throughout the enterprise," says Fisher.

Attacks such as distributed denial-of-service (DDoS), information theft, and worms and viruses can cripple an organization's ability to do business. For example, after surfing the Internet, an employee can innocently introduce a worm or virus to the corporate network through a remote-access connection (see "Looming Security Threats," page 19). Endpoint security such as Cisco Security Agent, coupled with network-based policy enforcement such as Cisco Network Access Control (NAC), stop the infection before it spreads by prescreening user PCs before allowing them to log into the network. This cooperation between computing and network elements creates a synergy that provides stronger protection than either system can accomplish alone. This principle is the basis for multilayer, modular security blueprints such as Cisco SAFE (cisco.com/go/safe).

While worm and virus attacks often make the news, an IT operator can introduce an exploitable vulnerability through a simple misconfiguration of a router, switch, or firewall. Configuration templates such as Cisco Smartports can help operators avoid common configuration mistakes, increasing network availability and implementing security policies at the same time.

Network Operations

Both elusive and critical to a successful business resilience strategy is network operations. Many enterprises purchase multimillion-dollar IT infrastructures, then manage them manually, which is one reason why so many organizations struggle to control operational expenditures. According to Sage Research, 39 percent of

network outages are caused by configuration errors, 27 percent by upgrade errors, and 10 percent by data entry errors. A Cisco poll found that 62 percent of its seminar attendees preferred using a manual command-line-interface (CLI) for configuring and managing their networks.

Investing in and using integrated management systems, network operators can eliminate configuration errors and speed up routine processes through automation. "You need to pry the intelligence out of your devices using tools," says Brian Junnila, senior manager for network management in Product and Technology Marketing at Cisco. "One process that benefits from automation and tools is change management. If you do the same task differently every time, you'll never progress beyond basic management." For example, it takes 93 hours per year for an operator to manually change passwords in 800 devices each quarter, with a 5 percent error rate. Automating that process using CiscoWorks Resource Manager Essentials drops the error rate to zero and takes less than one hour per year.

Assess Your Resilience

As with any journey, enterprises can plan a cost-managed path toward resilience goals when they know how resilient their IT systems are today. Assessments offer valuable insights about what an enterprise is already doing well and identify areas where it can improve both infrastructure and operations. The Cisco Advanced Services team offers extensive lifecycle-based services to help enterprises with assessment, design, implementation, and operations for high availability networking, security, network operations, and more. ■

FURTHER READING

- Cisco Business Ready Solutions
cisco.com/go/businessready
- CiscoWorks Network Management
cisco.com/go/ciscoworks
- Cisco Advanced Services
cisco.com/packet/171_6a1
- Cisco Smartports Solution
cisco.com/go/smartports
- Cisco High Availability Networking Services
cisco.com/packet/171_6a2
- Cisco Business Resilience Planning Services
cisco.com/packet/171_6a3



By GAIL MEREDITH OTTESON

THERE IS a big difference between concept and delivery. Creating useful business resilience strategies and building business-resilient IT systems requires substantial investment and deliberate focus. Cisco assists enterprises with a comprehensive approach to building resilient IT systems that increase business agility and withstand disruption. This approach traverses the entire organization, focusing on network, applications, communications, and workforce resilience.

First, Measuring Resilience

Resilience itself is difficult to measure, because it entails both quantifiable statistics such as network uptime and less tangible factors such as customer satisfaction.

“The appropriate IT perspective for measuring resilience should be service availability. To understand how resilient a service is, the IT organization as a whole has to marry its systems to business processes,” says Bobby Guhasarkar, senior solutions marketing manager in Enterprise Marketing at Cisco. “Many organizations do not measure service availability, but components of it. Business resilience is measured as the sum of application, server, and network availability combined with business processes.”

RECIPE FOR RESILIENCE

**BUILDING RESILIENCE INTO CAMPUS,
BRANCH OFFICE, AND TELEWORKER IT SYSTEMS**

Enterprises can begin to understand the resilience of their IT systems by measuring uptime. Service providers have long relied upon the “five nines” (99.999 percent) concept for planned downtime. But this metric can play a numbers game with availability without considering the true business impact of an outage. An IT group can justify meeting service-level agreements (SLAs) through statistical interpretation. For example, a 99.5 percent availability target allows a 50-minute weekly service outage. If an enterprise experiences one business-critical application outage of 100 minutes in a given month, the IT group can say it is exceeding SLAs while the business itself might suffer short-term revenue loss, customer dissatisfaction, or possible penalties resulting from regulatory noncompliance and litigation.

It’s more useful to measure the resilience of IT systems from the end-user perspective. At the service level, metrics for measuring availability to users include:

- Mean time between failure (MTBF)—how long a service is operational before it might fail. The maximum MTBF is limited by the MTBF of the least resilient service component.
- Mean time to repair (MTTR)—how long it takes to restore a failed service. The minimum MTTR is impacted by the MTTR of the least resilient service component.

Network Resilience

A high-availability network is the foundation of service resilience. According to Guhasarkar, enterprises that take a short-term, low-cost view toward network resilience are flirting with disaster. Not only do they lose the agility benefits of advanced services delivered through end-to-end network intelligence, they can scramble for hours trying to locate and troubleshoot failures, and then perhaps endure dealing with multiple vendors to solve the problem.

As IT research and analysis firm Gartner, Inc. discusses in its report on “The Real-Time Enterprise,” a high-availability networking strategy represents substantial investment in both capital and operational expenditures, but pays off with significantly higher uptime, greater customer satisfaction, increased revenues, and reduced exposure to regulatory penalties.

Cisco’s strategy for high-availability networking includes the following components:

- Reinforced network infrastructure
- Real-world network design
- Realigned network operations
- Real-time network management
- Relentless network support

Reinforced Network Infrastructure

There are four pillars to building highly available network infrastructures: *device-level resilience*, *network-level resilience*, *manageability*, and *self-protection*. For optimal device-level resilience, enterprises should select hardware platforms with high MTBF and software technologies that lower MTTR. Hardware is more resilient when it includes redundant components such as power supplies, supervisor engines, and routing engines.

DESIGNING CISCO IOS SOFTWARE FOR RESILIENCE

The Cisco IOS Software development team has a systematic, end-to-end approach to developing features that reduce network downtime. The three-pronged strategy includes:

System-level redundancy—combines redundant hardware components and resilience protocols that reduce MTTR of system failures and downtime and protect remote devices. These features include Nonstop Forwarding with Stateful Switchover (NSF with SSO), Hot Standby Router Protocol (HSRP), Gateway Load Balancing Protocol (GLBP), Stateful Failover for IP Security (IPSec), Stateful Network Address Translation (NAT), Warm Reload, and Warm Upgrade.

Network-level redundancy—provides faster network convergence, protection, and restoration in case of a network outage. Embedded Cisco NSF Awareness creates an intelligent protocol fabric, while convergence and self-healing routing protocols such as Intermediate System-to-Intermediate System (IS-IS), Border Gateway Protocol (BGP), and Open Shortest Path First (OSPF) dynamically reroute traffic around link failures. Some other features are IP Event Dampening and multicast subsecond convergence.

Embedded management—delivers proactive fault or event management, configuration management, and availability measurement that interact with automated features in CiscoWorks and tested third-party management tools. Example capabilities are Cisco IOS Embedded Event Manager and Cisco IOS in-service software upgrades.

Network-level resilience provides multiple links among critical devices, especially in the campus core. Branch offices might need secondary or backup WAN connectivity to a headquarters data center. Cisco IOS Software features reduce MTTR in case of device- or network-level failure by detecting imminent or certain outages and automatically failing over to a hot-standby component or converging around a failed link.

“Network convergence for a routing table with over 125,000 unique routes is now less than one second,” says Guhasarkar. “That’s a remarkable accomplishment from a mathematical point of view.”

Cisco IOS Software also includes embedded manageability and self-protection features (see sidebar, “Designing Cisco IOS Software for Resilience”).

Real-World Network Design

Of equal importance to resilient hardware and software, network designs must be proven in real-world scenarios—then tested, tested, and tested. A dedicated team of Cisco engineers develops and tests high availability design best practices for various applications and vertical industries, and publishes reference design

guides for enterprise network domains such as campuses, branch offices, data centers, and teleworkers (see cisco.com/packet/171_6b1 and cisco.com/packet/171_6b2). Cisco also collaborates with third-party laboratories to validate its design recommendations and provide feedback to the product design groups.

Realigned Network Operations

Enterprises must not underestimate how humans can unintentionally contribute to downtime. According to Gartner, one-third of unplanned network downtime incidents are caused by human error. Sage Research puts that number closer to two-thirds of the total. To achieve the highest service levels from high-quality products and excellent network design, enterprises need well-trained operations teams that develop and stick to disciplined operations processes. One of these processes is a thoroughly written, consistently enforced security policy for device access and management.

Cisco Advanced Services publishes operational best practices and assessment guidelines for the software lifecycle, change management, and service-level management (see sidebar, "Resources for Resilience," page 38). Cisco also offers expertise with its portfolio of High Availability Networking services spanning the network lifecycle: planning, design, implementation, operation, and optimization.

Real-Time Network Management

Enabling highly efficient network operations is a comprehensive suite of network management tools, designed to communicate with Cisco network systems to improve the speed and accuracy of fault, configuration, accounting, provisioning, and security (FCAPS) management. Automation features in CiscoWorks tools substantially improve how network managers use their time.

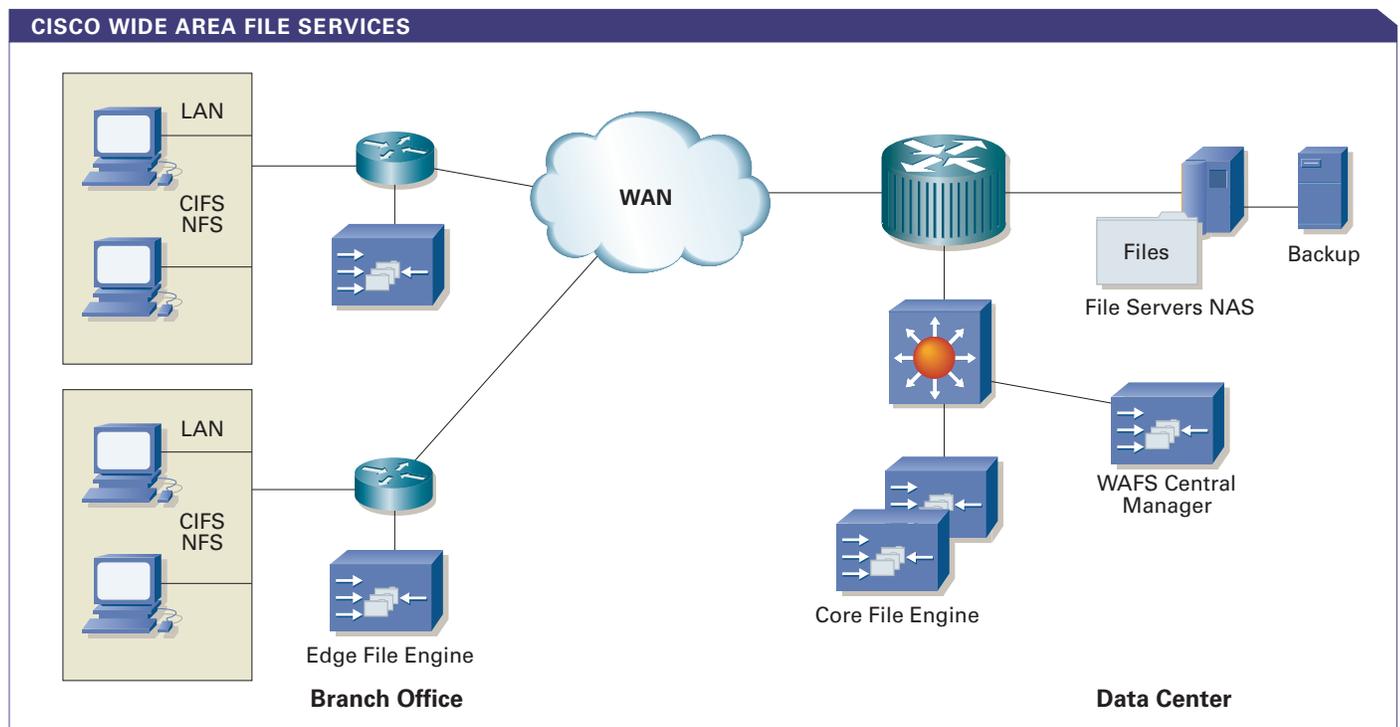
They can spend less time doing routine tasks and devote more effort to projects that increase business resilience.

Consistent configuration enhances network availability. Cisco Smartports is a solution for Cisco Catalyst switches that simplifies the configuration of critical features for Ethernet networks. Requiring minimal effort and expertise, Smartports features pre-tested switch port configurations based on Cisco best practices. Macros enable consistent, reliable configuration of essential security, availability, quality of service (QoS), and manageability features recommended for Cisco Business Ready Campus solutions.

Cisco also collaborates with best-of-breed management companies to offer enterprises a portfolio of problem management, root-cause analysis, and change management systems appropriate for a variety of budgets and resilience goals. Cisco and third-party partner tools help enterprises reduce network downtime with improved operational response times. For example, OPNET's NetDoctor and IT Sentinel products validate configuration changes and automate configuration auditing.

Relentless Network Support

In addition to an ongoing focus on providing support documentation and resources available online, Cisco provides specialized network certifications and training, sparing strategies, SMARTNet coverage, and a worldwide, always-there Technical Assistance Center (TAC) with more than 1600 support engineers including 400 with CCIE certification.



CLOSE TO HOME Using WAFS technology, the Cisco File Engine gives enterprises LAN-like performance for accessing central files from remote locations.

RESOURCES FOR RESILIENCE

To learn more about high availability networking, pervasive security, and building resilience into your network, applications, communications, and workforce, Cisco recommends the following online resources.

Cisco Business Ready Architectures

cisco.com/go/businessready

High Availability Networking

Cisco Best Practices for High Availability Networking

cisco.com/packet/171_6b4

IP Journal: “High Availability in IP Routing”

cisco.com/packet/171_6b5

Packet: “Around-the-Clock Uptime”

cisco.com/packet/171_6b6

Packet: “High Availability Networking”

cisco.com/packet/171_6b7

Packet: “High Availability for Campus Networks”

cisco.com/packet/171_6b8

Packet: “Calculating New Routes Faster”

cisco.com/packet/171_6b9

Cisco Smartports

cisco.com/go/smartports

Pervasive Security

Cisco SAFE

cisco.com/go/safe

Cisco AutoSecure

cisco.com/packet/171_6b10

Cisco Guard DDoS Mitigation Appliances

cisco.com/packet/171_6b11

Planning Services for Network Security

cisco.com/packet/171_6b12

Communications Resilience

Cisco AutoQoS

cisco.com/packet/171_6b19

Cisco 3800 Series ISR

cisco.com/packet/171_6b20

Cisco 2800 Series ISR

cisco.com/packet/171_6b21

Cisco CallManager Express

cisco.com/packet/171_6b22

Cisco Survivable Remote Site Telephony

cisco.com/packet/171_6b23

Cisco Unity Express

cisco.com/packet/171_6b24

Enhanced Security for IP Communications

cisco.com/packet/171_6b25

Applications Resilience

Data Center: SAN Extension for Business

Continuance Overview

cisco.com/packet/171_6b13

Understanding the Alternatives for Extending SANs

cisco.com/packet/171_6b14

Cisco WAFS and Cisco File Engines

cisco.com/packet/171_6b15

Branch office router-based content caching

cisco.com/packet/171_6b16

Branch office router-based compression

cisco.com/packet/171_6b17

cisco.com/packet/171_6b18

Optical Networking Solutions

for Business Continuance

cisco.com/packet/171_6b27

Workforce Resilience

Cisco Business Ready Teleworker

cisco.com/packet/171_6b26

Pervasive Security

Enterprises building a resilient IT infrastructure must invest in end-to-end security. A single unchecked virus or worm can cripple basic business activities (such as e-mail) for days. A denial-of-service (DoS) attack can render business-critical applications unavailable. Information theft can expose organizations to loss of customer confidence and regulatory violations.

Current trends that can create potential vulnerabilities include wireless networking and using the Internet in lieu of private-line WAN services. Enterprises must balance the advantages of workforce mobility and cost containment with security policies and systems that address these vulnerabilities. The Cisco SAFE blueprints provide extensive information about the threats and vulnerabilities for each network domain, along with reference architectures for mitigation with multilayer security solutions (cisco.com/go/safe).

In increasingly complex IT environments, a robust security policy is more important than ever when a business wants to remain resilient during a “day zero” attack. Cisco offers expert security assessment services that assist enterprises with identifying the current security posture of IT systems and developing security solutions and operational best practices to address identified issues. For more on these assessment services, visit cisco.com/packet/171_6b3.

Throughout campus and branch office networks, Cisco Network Admission Control (NAC) is a system-level defense solution incorporating routers, access servers, endpoints, and virus protection software. Cisco NAC allows the network to identify and restrict user and device access, logically isolate unauthorized or corrupted components, and rapidly respond to threats according to policy guidelines.

Firewalls and intrusion detection and prevention devices throughout the campus and branch offices help identify and contain DoS attacks and thwart deliberate hacking or phreaking activities. The CiscoWorks Security Information Management System (SIMS) gathers, correlates, and analyzes the massive amount of security data generated by hundreds of these devices, scaling it into meaningful guidance for identifying and responding to attacks in progress. In concert with Cisco Traffic Anomaly Detectors, Cisco Guard DDoS mitigation appliances let normal operations continue during a DoS attack by filtering DoS traffic from flows and allowing legitimate traffic to continue. In campus wireless LANs, Cisco Aironet solutions provide IEEE 802.1x authentication, wireless encryption, and new wireless LAN intrusion detection.

Branch offices expose campus networks and data centers to back-door attacks. Security operators can use Cisco VPN, IPSec, and encryption technologies to secure branch office connections. Operators also need the ability to quickly lock down branch office routers. A feature of Cisco IOS Software, Cisco *AutoSecure*, offers a “one touch,” remote device lockdown process. It scans the router, closes offending ports, finds vulnerabilities, and suggests configuration changes to prevent future attacks.

Resilient Applications

When it comes to ensuring that business applications are always available and accessible, the data center is the focus of attention. Cisco provides a full range of optical and IP storage area networking (SAN) extension solutions over a variety of media using Cisco ONS 15000 optical networking solutions and the Cisco MDS 9000 Family of multilayer directors and fabric switches. Among the technologies that reside in the data center to address both synchronous and asynchronous replication are local and global server load balancing, intelligent storage networks such as virtual SANs (VSANs), security (e.g., intrusion detection and firewall security), and business continuance. These and other solutions critical to providing applications resilience within the data center are covered in greater detail in the article “Ready for Anything,” page 40.

In addition to providing data center solutions for enterprises and service providers globally, Cisco has been working on providing enterprises with an end-to-end storage consolidation solution for branch-office data, which involves network-attached storage (NAS). To date, this work has largely centered on Cisco Wide Area File Services (WAFS) technology, which gives enterprises LAN-like performance for accessing central files from remote locations. Cisco WAFS uses sophisticated protocol-level caching, compression, and network-optimization techniques to minimize latency penalties associated with file-server access over the WAN. Branch-office users find application performance remains acceptable. Using this technology, the Cisco WAFS Edge File Engine located in the branch office corresponds with a WAFS Core File Engine collocated to the departmental servers in the central data center. The WAFS Central Manager is a management module that provides Web-based facilities for centralized management, configuration, monitoring, and maintenance of all file engines distributed throughout the enterprise (see figure, page 37).

Client file server requests in the branch office are directed to the local edge file engine, which determines whether to manage the request locally using its file system cache or forward the request to the remote file server. In the latter case, the edge file engine encapsulates the Common Internet File Server (CIFS) or Network File System (NFS) request and sends the request over the WAN to the core file engine using the Cisco WAFS transport protocol. This protocol restricts the chattiness of CIFS and NFS protocols over the WAN, protecting bandwidth usage for end-user data retrieval. The core file engine decodes the Cisco WAFS request into CIFS or NFS, issues the request to the file server, and encapsulates the response into the Cisco WAFS protocol for transmission to the branch office edge file server, then to the client. The edge file engine behaves as a pure caching and acceleration system for remote file servers, as opposed to an actual file server. This characteristic simplifies WAFS operations and maintenance, because it eliminates the need to define

and manage users and permissions or manage cache capacity. It requires no client software and no local storage.

In step with providing enterprises with solutions suited to storage consolidation over wide-area networks, Cisco recently entered into an agreement with EMC to sell and support EMC NS500 and NS700 Series/integrated NAS solutions together with the Cisco File Engine. With this integrated offering, IT administrators will be able to take advantage of centralized resources for backup and disaster recovery of their branch-office data. The integrated EMC/Cisco solution will be sold and supported by Cisco directly and through its worldwide partner channel.

Cisco also offers branch office acceleration and WAN compression functions in network modules for Cisco 2600 and 3700 Series routers and the new Cisco 2800 and 3800 Series Integrated Services Routers. The Content Engine Network Module intelligently delivers a number of different file types, including Real Windows Media Technologies (WMT), Darwin, HTTP, PDF, Flash, Shockwave, FTP, and Multimedia Messaging Service (MMS) content, using the demand-pull caching and pre-positioning technologies provided by Cisco Application and Content Networking (ACNS) software. This solution helps enterprises accelerate the delivery of software files, security patches, business video, HTML content, and more to improve application and WAN bandwidth performance. The WAN Compression Module, available for Cisco multi-service and integrated services routers, optimizes WAN bandwidth utilization without sacrificing throughput rates by reducing frame size and thereby allowing more data to be transmitted over a link. With this module, network managers can increase application performance and service availability for end users without expensive infrastructure upgrades.

Resilient Communications

The most essential component of a business resilience strategy is resilient communications, which alone justifies the need for a high availability network. While organizations can withstand limited interruption in many IT applications, near-continuous communication systems are vital to employee safety and business flow, especially during and immediately after a disruption.

The high availability network carrying IP telephony, video, and unified messaging traffic uses QoS mechanisms to prioritize and expedite time-sensitive voice and video traffic. In Cisco Catalyst switches, Cisco AutoQoS incorporates value-added intelligence in Cisco IOS Software and Catalyst OS (CatOS) Software to provision and manage large-scale QoS deployments. It encapsulates Cisco expertise for best practices based on actual deployments. Compared to manual methods, Cisco AutoQoS can reduce deployment cost and time by as much as two-thirds. The first phase of Cisco AutoQoS automates voice-over-IP (VoIP) deployments for enterprises who want to deploy IP telephony but lack the expertise or staff time to plan and deploy IP QoS.

Like its public telephony counterpart, Cisco designed its IP communications solutions for resilience from the ground up. It has long offered clustering of Cisco CallManagers, redundant clustering in secondary data centers, and features such as Power over Ethernet (PoE) to IP phones, extension mobility, 911 locator support, and remote agent capabilities for call center personnel.

Continued on page 89



READY FOR ANYTHING

BUILDING RESILIENT DATA CENTERS

RESILIENCE is most important in the data center: the heart of the IT infrastructure. The trend of consolidating applications, servers, and storage into the data center presents a unique opportunity for organizations to increase their business resilience. IT organizations embarking on consolidation projects must ensure that they build a secure and resilient foundation for current and future applications and technologies. IT managers who adopt resilient data center architectures can also increase business agility as they realize the benefits of service-oriented infrastructure that can rapidly respond to new application requirements.

Network Resilience

The Cisco *Business Ready Data Center* provides a resilient network architecture. Properly deployed, a data center network with IP, storage, and optical components creates a foundation for optimal application service levels. This approach can substantially reduce data center operating expenditures, which can chew up the IT budget.

Essential to network resilience is a high availability design poised for future capacity and service growth. The network investment is a small portion of data center total cost of ownership, yet the META Group (April 2004) states that “the impact of a suboptimal network design can have grave consequences and negate all other investments the business may have made in the design and implementation of a high availability data center.” For a look at the basics of building highly available networks, see the article “Recipe for Resilience,” page 35.

The first tenet for data center resilience is resilient network elements that offer high mean time between failure (MTBF). In large data centers, the aggregation layer of the Cisco IP network is typically based upon the Cisco Catalyst 6500 Series Switch, a highly resilient, scalable, intelligent platform that optimizes and secures access to and between business-critical applications. Its redundant supervisor option offers stateful failover, while specialty service modules allow integration of intelligent network services—such as firewalling, encryption, load balancing, and network analysis—directly into the network fabric. Cisco IOS Software resilience features such as Nonstop Forwarding with Stateful Switchover, Stateful Failover for IP Security (IPSec), and Hot Standby Router Protocol (HSRP) ensure rapid recovery and session protection in case of component or device failure.

Storage Innovation

The Cisco data center storage network is based upon the Cisco MDS 9000 intelligent SAN switch product family. High availability features such as nondisruptive software upgrades and stateful failover are hallmarks of the MDS 9500 Series. In addition, throughout the MDS SAN switch family, Cisco has introduced several technology innovations in storage networking that address network resilience. One such innovation, virtual storage area networks (VSANs), was recently adopted as a standard by the ANSI T.11 Fibre Channel Committee. VSANs allow consolidation of isolated storage “islands” into a single, scalable,

centrally managed and administered physical network without impact on the availability or security of logically separate SANs. Cisco has also added network resilience with inter-VSAN routing, an innovation that combines the flexibility of routing with the stability and scalability of VSANs. With inter-VSAN routing, connectivity can be established between devices across VSAN boundaries while still maintaining the fabric isolation of VSANs. This innovation facilitates common resource sharing (such as tape backup systems) among several VSANs and is essential for scalable data center consolidation.

Resilient network elements within the data center are important, but increasingly customers expect SAN extension solutions to address business continuance requirements. These solutions must offer the same degree of availability and security as the SANs within the data center. Cisco's resilient SAN extension options provide the ability to match application requirements to the most cost-effective connectivity option. These options include Fibre Channel over IP (FCIP), which can leverage a customer's existing WAN infrastructure, and optical solutions such as dense wavelength-division multiplexing (DWDM) and Fibre Channel over SONET or SDH for high-performance, low-latency application requirements.

Supporting hundreds of gigabits-per-second throughput and 50-ms recovery times, optical networking is an attractive solution for business continuance demands, and has become more affordable (see "Adventures in Resilience," page 45). By supporting voice, video, data, and storage, multiservice optical solutions such as the Cisco ONS 15454 consolidate costly redundant WAN networks onto a single highly efficient solution.

All of these network elements must be incorporated into a highly resilient network design that ensures fast, seamless recovery from any foreseeable disruption to the network or attached servers and storage devices. Cisco provides best-practice reference design guides based on tested and validated configurations to help achieve this level of resilience. High availability technical reference guides intended specifically for data center networks are available at cisco.com/go/datacenter.

The final component of network resilience is a disciplined operations team that consistently follows high availability networking best practices, especially for change management. The continually changing data center environment requires collaboration between the groups that are responsible for servers, applications, storage, optical transport/security, business continuance, and the network. Operational procedures need to standardize processes, enforce accountability, and use management tools that automate routine processes. Sage Research traced two-thirds of unscheduled network outages to management errors, most of which can be eliminated through a strict

change management system based on tools such as CiscoWorks Resource Manager Essentials.

Applications Resilience

A major concern during data center consolidation is how to maintain application security, availability, and service levels on a consolidated infrastructure. Unconsolidated, siloed applications offer physical separation, leading to the false belief that they reduce the risk of downtime or security breaches.

"Data center managers must consolidate their resources, because siloed applications and infrastructures are rigid and inefficient," says Jonathan Gilad, solutions manager of Data Center Networking at Cisco. "A well-designed consolidation doesn't sacrifice security or application service levels; it actually improves application resilience."

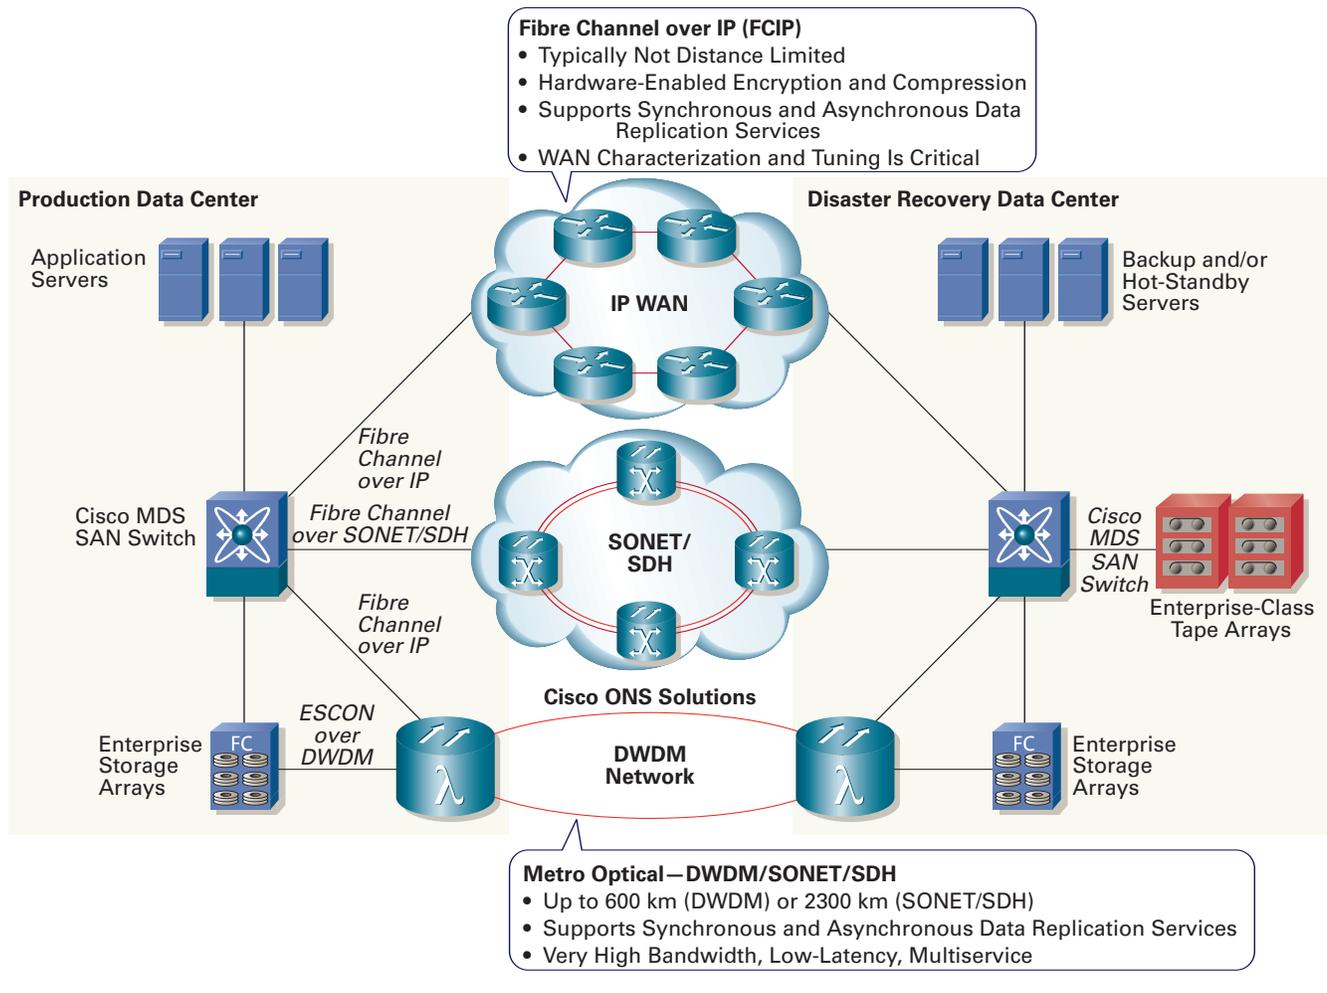
One technology that accelerates data center consolidation is Cisco Wide Area File Services, which provides high-capacity file services in branch offices, encouraging branch office personnel to relocate all files into the highly available and secure data center without suffering performance degradation.

The trend toward building applications resilience in a consolidated data center focuses on virtualization, which uses networks to present available server and storage resources as logical entities for security and management purposes.

The Cisco Business Ready Data Center offers both IP and storage network virtualization solutions. In the IP network, Cisco IOS Software enables application isolation using virtual LANs (VLANs). Traffic remains logically separate so that application misbehavior such as broadcast storms cannot spread to other applications. Firewalls isolate resources to minimize the spread and impact of security breaches or virus and worm infections. The integrated Cisco Firewall Services Module (FWSM) can be configured as multiple virtual firewalls, one for each VLAN, reducing both complexity and expense. Within each security zone, server health monitoring and load balancing services in the Cisco Content Switching Module optimize higher-layer application performance across server clusters.

Cisco enables intelligent SAN virtualization services through its Cisco MDS 9000 family of SAN switches. "Moving storage virtualization into the network increases applications resilience," says Jason Warner, senior technical marketing engineer in the Storage Virtualization Group at Cisco. "Applications don't need to know where data is physically located, only that they can access it. One vendor-independent interface allows storage resources to remain online

By GAIL
MEREDITH
OTTESON



LONG-RANGE SAN Cisco SAN extension enables transparent SAN extension between data centers across town, distances as far away as 20,000 km, or halfway around the globe.

during migrations associated with data center consolidation projects or mergers and acquisitions.”

The Cisco MDS 9000 family provides an open platform for hosting and centrally managing intelligent storage services from multiple partners. These intelligent fabric applications include network-hosted virtualization, data migration, and backup. With this open platform approach, storage managers have the ability to choose storage virtualization and management applications that best meet their needs. Cisco is also driving industry standards in this area as the technical editor for the Fabric Application Interface Specification (FAIS) draft standard.

Business Continuation

In this 24-by-seven world, organizations need IT infrastructure that supports

their business processes, no matter what happens. Business continuation is the art and science of protecting and restoring IT assets in case of disruption. A business impact analysis identifies those applications that need nonstop service and those that can withstand limited downtime without adversely affecting the business. These parameters dictate the scope of business continuation systems that an organization needs to protect itself and control costs.

Cisco supports several IP and Fibre Channel SAN extension options for business continuation between data centers. For short to medium (200–600 km) distances, Fibre Channel directly over DWDM or SONET/SDH is typically optimal due to the high bandwidth and low latencies supported. The Cisco ONS 15000 product

line is qualified for the broadest range of replication and mirroring applications from leading storage vendors and managed service providers. Fibre Channel over IP (FCIP) allows enterprises to leverage their current WAN infrastructure to control costs. Unlike Fibre Channel, which uses buffer credits that limit practical transport distance, FCIP relies upon TCP for transport control, enabling transoceanic or intercontinental transport up to 20,000 km. This storage-optimized version of TCP supports a 32-MB maximum window size.

PACKET ONLINE EXCLUSIVE

For a design guide on data center disaster recovery and load distribution using Interior Gateway Protocol (IGP) and Border Gateway Protocol (BGP), visit cisco.com/packet/171_6c1.

SAN Extension

The introduction of SAN-OS 2.0 for the Cisco MDS family enhances SAN extension services with improved performance and security, as well as support for hardware-based compression. Cisco MDS SAN switches also support tape acceleration and disk write acceleration to support replication over longer distances between sites. Both acceleration features push exchanges to a remote tape backup or disk array without waiting for acknowledgments to each exchange, increasing transfer rates and reducing the impact on application performance by up to 100 percent.

Replication mode—synchronous or asynchronous—determines the SAN extension solution architecture. A high availability solution uses dual switches at each end, with one or more VSANs across each link. Client protection features within arrays perform error recovery and data validation. The Cisco MDS 9000 augments this protection with PortChannel technology. Analogous to EtherChannel in Ethernet switching, PortChannel logically groups up to 16 Fibre Channel or FCIP links for load balancing and capacity reasons. Each PortChannel can trunk multiple VSANs, and boosts the transfer speed of a Fibre Channel connection from 2 Gbit/s to 32 Gbit/s.

In combination with write acceleration, FCIP can service synchronous and asynchronous replication with performance similar to Fibre Channel over optical. The new SAN Extension Tuner feature in SAN-OS version 2 allows storage managers to set minimum and maximum availability and other parameters according to available bandwidth and round-trip time.

Data Center Security

The data center is an attractive target for malicious activity. Hackers, worms, and distributed denial-of-service (DDoS) attacks can cause considerable havoc and costly damage when a center is not properly secured. As data center managers consolidate resources such as servers, storage, networks, and applications, they must consider how these changes affect the security posture and application resilience. Security and network managers must collaborate to understand the particular vulnerabilities and threats to data center resources so

that they can develop a network security architecture that protects against threats that often pervade the enterprise. Cisco security solutions provide comprehensive protection through a multilayer deployment of secure connectivity, threat defense, and trust and identity solutions.

The Cisco Business Ready Data Center achieves optimal end-to-end security, performance and manageability by integrating security directly into the network infrastructure. It takes advantage of the advanced integrated security capabilities of Cisco Catalyst switching and Cisco MDS 9000 intelligent SAN switches. Integrated security software and service modules for the Cisco Catalyst 6500 platform offer firewall, intrusion detection system (IDS), Secure Sockets Layer (SSL), DDoS, and IP Security (IPSec) virtual private network (VPN) services at the higher performance levels required for bandwidth-intensive data center environments.

In storage networks, the Cisco MDS 9000 switches employ Secure Shell, Secure FTP, RADIUS, SNMPv3, and role-based access control (RBAC) against unauthorized management access. In addition, Fibre Channel Security Protocol (FC-SP) delivers IPSec-like functionality to Fibre Channel fabrics and ensures secure switch-to-switch communication. FC-SP, an ANSI T11 standard co-authored by Cisco, provides confidentiality, data origin authentication, and connectionless integrity across the fabric. VSANs also provide security by dividing a common physical SAN infrastructure into separate virtual SAN islands.

Among the many options available for security data centers, enterprises should consider the following:

- Infrastructure security with tools such as hardware rate limiters and control plane protection
- Stateful session inspection with integrated Cisco firewalls
- Endpoint protection with Cisco Security Agent in publicly accessible servers
- DDoS mitigation with Cisco Detectors or intrusion detection devices deployed throughout all data center domains and a Cisco Guard to filter traffic, allowing legitimate sessions to continue

- Trust and identity protection through role-based access control

- Perimeter and domain protection with firewalls and intrusion detection and prevention devices

Communications Resilience

Data center consolidation often includes Cisco IP communications services and application servers. For resilient communications, every data center location with IP communications services should have a direct gateway to the PSTN. A Cisco multiservice or integrated services voice gateway router provides a wide range of packet telephony-based voice interfaces and signaling protocols. Cisco CallManagers can be clustered for high availability. Backup or hot standby clusters can reside in secondary data centers for rapid IP telephony recovery in case of major disruption. Likewise, backup Cisco Unity Unified Messaging and IP Contact Center servers in the secondary data center can preserve message integrity and ensure continuous contact center service.



A thoughtfully considered data center consolidation project helps enterprises achieve greater business agility through resilience. Cisco stands ready to help enterprises plan, implement, and operate secure, consolidated data centers with solutions, operational best practices, and reference designs that address the technical complexities and operational realities of the data center. ■

FURTHER READING

- Cisco Business Ready Data Center cisco.com/go/datacenter
- White paper on high availability data center networking cisco.com/packet/171_6c2
- Cisco storage networking cisco.com/packet/171_6c3
- White paper on extending SANs cisco.com/packet/171_6c4
- Cisco SAN extension products cisco.com/packet/171_6c5
- “Five Steps to Building an Intelligent Networking Infrastructure for Business Continuity” cisco.com/packet/171_6c6
- Cisco High Availability Networking Services cisco.com/packet/171_6c7

RESILIENCE

ADVENTURES IN

By RHONDA
RAIDER



**HOW FOUR
ENTERPRISES
MAINTAINED
BUSINESS
CONTINUITY
IN THE FACE
OF FLOOD, FIRE,
AND POWER
OUTAGES**

FOR A COMPANY without a resilience strategy, an event as trivial as tripping over a cable or a missed fiber installation deadline can threaten business continuity. But for a company with effective resilience measures for the network, applications, communications, and workforce, even a catastrophe can be weathered without business disruption. In fact, catastrophes, either natural or man-made, often provide the impetus for resilience projects. That certainly was the case when Hurricane Floyd struck in 1999, flooding a building at East Carolina University (ECU) in North Carolina, and shutting down the network—and the university—for two weeks. “To restore even limited phone service we had to jump through hoops, including hard-wiring connections to our phone switch, which took a day and a half,” recalls Rob Hudson, director of network services at ECU. “That’s a long time when family members are worried about each other’s safety.”

Determined not to repeat this experience, the ECU IT group upgraded its data network the next year to bolster network and communications resilience, with risk mitigation as the chief objective. “An important component of our disaster recovery strategy was IP telephony,” says Martin Jackson, manager for network engineering. “That would give us more control and the flexibility to restore phone service when unexpected events occurred.”

Today the university can instantly establish an Emergency Operation Center (EOC) to provide phone service, using only the contents of a strongbox containing a Cisco Catalyst 4500 Series Switch with 25 Cisco IP phones. The IT group prepared three potential on-campus locations for the EOC by connecting armored fiber directly to the Campus PSTN Gateway. A site generator powers the phones. "We have instant phone service that's not dependent on any other network links being up," says Hudson. "The center can be up and running within one hour."

Rapid Restoration for Campus Network

As part of its overall network and communications resilience strategies, the university also has redundant Cisco CallManager servers and Cisco 3640 multiservice access routers in two separate locations on campus. To date, 10 percent of the university's 8500 extensions have been converted to IP telephony. Most buildings on campus have dual paths to the carrier: "Even if we lose a remote shelf for voice, we can still keep critical numbers operational," says Jackson. And in the event of an outage in the copper or fiber, IT can take advantage of wireless—a resilience tactic that has paid dividends twice since being deployed in 2001.

The first occasion was when a cut fiber took down a segment of the on-campus medical center network, preventing staff and physicians from accessing patients' electronic medical records (EMRs) and images. Demonstrating workforce resilience as well as network resilience, IT averted the potentially life-threatening outage by restoring the network in just 45 minutes. Jackson recalls, "We just installed a couple of spare Cisco Aironet 1200 Series wireless access points, one in a window of the medical school building and another in the window of a building with connectivity across the street. That extended the network from one building to the other."

In 2004, the innate resilience of IP communications was validated again for ECU—this time when a vendor missed the deadline for installing fiber in a new building scheduled to open its doors to 200 employees the next week. "Without a phone system, the building cutover date would have been delayed, affecting the personnel whose major job function was fundraising," says Hudson. This time IT installed a Cisco Aironet 1400 54-MB bridge to provide connectivity to a nearby building's network, enabling workforce resilience. "The employees, who were used to having a 100-MB connection from their desktops, didn't even notice the difference in application performance or phone service, so no productivity was lost," says Hudson.

Borrowing TDM Connectivity: Kaye Scholer LLP

Paired Cisco Aironet receivers also saved the day for Kaye Scholer LLP, a global law firm headquartered in New York City. On a Saturday in October 2004, a massive fire swept through the law offices, forcing the company to scramble to find space for 1000 employees for the one month needed for repairs. Even more challenging than finding instant office space in New York City was acquiring phone service. "We were moving sight unseen into four separate buildings, two without existing phone service," says John Palefsky, director of technology services. "In one building we had two contiguous floors and one more floor 30 levels below, and in the building across the street we had two discontinuous floors. Getting phones wired in this situation would ordinarily take weeks." If ever there was a predicament requiring communications and workforce resilience, this was it.



Kaye Scholer's Frankfurt, Germany offices had recently deployed a Cisco CallManager server and Palefsky had a hunch its flexibility might serve the current situation, as well. He called the Cisco partner that had deployed the solution, Dimension Data. "On Thursday I asked for 600 phones, and the system was up and running on the following Wednesday. That included ordering, purchasing, shipping, and configuring. Amazing."

Dimension Data installed new Cisco Catalyst 4506 switches in two of the buildings, interconnecting them within each building with fiber that Kaye Scholer had quickly installed after taking over the space. Kaye Scholer's technology staff moved existing Cisco Catalyst switches from its damaged building to the remaining two buildings to complete the network. The building with the Cisco CallManager had three dishes. A Cisco Aironet dish linked with its counterpart in the window of a building across the street to extend the Ethernet network there. Another Cisco Aironet dish faced its counterpart in a building that already had phone service one-half mile away to share connectivity with the PSTN through a Cisco 3745 Multiservice Access Router. A third dish, acquired from an Internet-over-microwave vendor, provided Internet connectivity to the complex. "Acquiring phone service in the traditional way from the phone company would have taken one to two weeks—not an option," says Palefsky. "And yet with wireless connectivity, one Cisco CallManager server and Cisco Unity server supplied phone and voice-mail service to five floors in two buildings. It's even more remarkable that the whole project was completed in less than one week. With any other solution it would have taken five times as long."

Shearman and Sterling LLP

Network and applications resilience is top of mind for New York-based businesses, whose experiences include both the September 11, 2001 terrorist attacks and the August 2003 power grid failure. International law firm Shearman and Sterling LLP, headquartered in New York, deploys a remote, lights-out data center. "Our main objective was to mitigate the risk, which we assessed to be high, of having a data center within our world headquarters," says Tony Cordeiro, chief information officer. Of primary concern is the availability of mission-critical messaging applications, used globally. "Our practitioners simply cannot afford to be out of touch

with their clients,” says Cordeiro. “We wanted to be able to sustain a site-level failure at our New York location without interrupting operations in Germany, the UK, San Francisco, or Asia, for example. These recent experiences, plus the recommendation of our auditors, afforded us the opportunity to reposition our global data center away from the city, off the grid and with its own redundant power supply, HVAC, and fire suppression systems.”

To achieve network and applications resilience, Shearman and Sterling designed two data center locations, which are physically removed from headquarters, in New Jersey and the UK. A primary design objective was continuing to provide the highest level of service for the 1500 employees in the New York office, who were accustomed to LAN-speed performance. “As we planned to move our data center and applications, we asked ourselves how we could replicate LAN performance to our global data center,” says Dan McLoughlin, manager of networks, infrastructure, and voice at Shearman and Sterling. “After analyzing various types of connectivity, we decided on an optical networking solution using Cisco optical equipment. From an operations and technology standpoint, the Cisco ONS equipment provides carrier-class reliability and resilience, as well as the needed throughput. This ensured maximized network uptime and LAN-like performance.”

Multiple Points of Entry for Network Resilience

The Shearman and Sterling IT team designed the optical network with multiple points of entry and exit for each building, creating a resilient infrastructure with no single point of failure. “In Manhattan there is always construction. If a backhoe digs through the building’s primary point of entry, that one is out of service for an unknown period of time. To ensure business continuity we needed another that is diverse,” says James Thomas, senior network engineer.

Based on these considerations, Shearman and Sterling chose to deploy an optical networking infrastructure using the Cisco ONS 15454 multiservice optical networking platform. This platform provides a flexible solution that enables Shearman and Sterling to start with a SONET infrastructure over a single dense wavelength-division multiplexer (DWDM) wavelength and add DWDM infrastructure as throughput demands require. “This highly flexible platform best meets our needs because we can scale up to 32 wavelengths of 10-GB throughput as needed, over time, while immediately supporting Fibre Channel connectivity for tape backups over the metro area network,” says Thomas.

Four Cisco ONS 15454 SONET multiservice provisioning platforms form the network core. Two reside in the New York offices, one at the global data center in New Jersey, and one at a carrier collocation facility in New York City. “By placing one Cisco ONS 15454 in the carrier collocation facility, we can connect to any other major carrier,” says Thomas. “With this design decision, we’ve ensured resilience for our infrastructure by allowing for multiple points of entry through multiple separate paths.”

The network currently runs at OC-192 speeds (10 GB) and is used today to support Fibre Channel communications between the New York office and data center. “We can physically locate some servers and backup devices locally and others in the offsite data center, and connect to them just as if they all resided locally,” explains Cordeiro.

In the event of disaster at any site, the optical network enables the law firm to restore data from the offsite data center remotely. In the highly unlikely event that connectivity to the remote data center is lost, Plan B is to send an employee to the remote data center to retrieve the backup. “Having the site running and humming is one thing; connectivity is equally important,” says McLoughlin. “That’s why dual entries into facilities are so important.”

More Resilient, Less Costly

Shearman and Sterling turned up its New York metro area network in December 2004, and brought its enterprise systems’ development environment live in January 2005. “The investment we’ve made in a resilient network will pay for itself within one year,” says Cordeiro. “We’ve effectively built out a carrier-class infrastructure, providing the needed bandwidth and flexibility for growth.”

This flexibility is not just limited to throughput, because the new network will support not only storage communications but all communication needs including IP, traditional voice communications, IP telephony, and other latency-sensitive applications such as video, increasing the efficiency and cost savings over time.

Communications Resilience for Long Island University

Long Island University (LIU) bolstered its communications resilience strategy after September 11, 2001, when it lost phone service. Within one hour the university regained its ability to make outbound calls by transferring the direct outward dial (DOD) lines to a Cisco CallManager server from the Brooklyn campus to the C.W. Post campus. “Staff and students were very appreciative that they could reach their families,” says George Baroudi, chief information officer. “When people realized how much resilience IP communications provided, we decided to extend it on a much larger scale.”

LIU is migrating from Primary Rate Interface (PRI) and ISDN to IP-based trunks, which will provide resilience for incoming as well as outgoing calls. Each of two campuses has a Cisco CallManager server and Cisco Catalyst 6500 Series Switch and can take over service from the other campus in less than five minutes—something that would take days without IP communications, according to Baroudi. “Our telephone service provider supports IP telephony, so if one campus location is out, we can simply provide a different IP address for DOD or DID [direct inward dial] calls.” Baroudi notes that this concept—decoupling a phone number from a physical location—has long been available for individual phone numbers, with follow-me services. “If it can be done for a single user, why not for a slew of DIDs, like the entire telephone exchange?” The university’s four smaller campuses and six satellite offices connect to the Cisco CallManager server across the WAN. “Should the WAN link become available, the campus router uses its built-in SRST [Survivable Remote Site Telephony] feature to automatically begin routing calls over not only the PSTN, but also the Public Internet Network using a second CallManager over a secure Internet connection.”

Baroudi concludes, “Our network is in use 24 hours a day because that’s the nature of the world. I don’t see any difference between voice, data, or video traffic: If it’s IP-based, we can make it resilient.” ■

Broadband Explosion

Broadband services for consumers = broad revenue growth for providers.

By Janet Kreiling



In just one week last December, *The Wall Street Journal* published a swath of news about broadband technology and use: Verizon offering broadband wireless access through a variant of CDMA2000 called evolution-data optimized (EV-DO); Comcast, Time Warner Cable, and Cox Communications using fiber to deliver video on demand (in addition to offering triple play services); companies in the US lobbying Washington for more wireless spectrum for broadband; British Telecom adding one million new broadband customers in the previous four months. . . .

Broadband in the news isn't surprising. But in these instances, *consumer broadband* is making the headlines. "The battle for the home is no longer over voice, but over broadband," says Christopher Dobrec, director of business development in the Linksys Division at Cisco. "Broadband is the ideal platform on which to converge consumer applications for voice, video, and data."

These days, a typical home equipped with cable or DSL Internet access might see adults and children alike using the Internet to do research, get news, and be entertained: one of them downloading slide presentations; a group of teenagers on their Xboxes matching wits over the Internet; yet another person listening to streaming audio or watching streaming video from a station across the country or abroad. There's tremendous demand for consumer broadband access—and, as this example demonstrates, there is

no killer application. Consumers use broadband for a variety of applications, and the more bandwidth they have the more they use it.

Yet, the broadband penetration rates in the US and Europe are low: just above 20 percent and 15 percent, respectively. The rates are much higher in some countries in Asia and Europe. In South Korea, for example, about 75 percent of households have access to high-speed connections—at speeds up to 40 Mbit/s. Japan, too, has high broadband penetration and high speeds available to consumers. The US and Europe are expected to mostly catch up. By 2010, more than 60 percent of US households should have broadband access, with a similar penetration rate in Europe. Bit rates should also increase, to the 10-Mbit/s range and beyond.

Worldwide, broadband deployments grew from 33 million at year-end 2002 to 97 million by year-end 2003, and were expected to reach 140 million at the end of 2004, according to Ovum Access Forecasts. Ovum projects that revenues from broadband are expected to increase at a compound annual growth rate (CAGR) of nearly 21 percent from 2003 to 2008. That's a *business opportunity*—one many service providers are already banking on. In January 2004, for instance, Verizon stated that it would commit some US\$3 billion of capital over the next two years to bring broadband to the mass market.

"Nevertheless, some providers see broadband more as a threat than an opportunity," says Fernando Gil de Bernabé, managing director of Cisco's Internet Business Solutions Group for Service Providers in Europe, the Middle East, and Africa (EMEA). Broadband can carry voice over IP (VoIP), which, although still a small part of broadband usage, cuts into traditional telephone services revenues. Moreover, broadband allows consumers and small and mid-sized businesses (SMBs) to consolidate their voice and data lines into one connection, further cutting down on fixed-line revenues, says Gil de Bernabé. Factor in the threat that wireless imposes on residential and SMB land lines, and the impact becomes even greater. According to data analysis conducted by Cisco's Internet Business Solutions Group, revenues from traditional fixed-line voice services continue on a downward trend, with CAGR in

the negative single digits, depending on the market. And the total number of fixed lines is now declining in both the business and consumer markets at a rate of between -1 percent and -3 percent in Western Europe.

“There are already many service providers in the world that have more than one million broadband customers each,” says Pankaj Gupta, senior manager for broadband marketing in Cisco’s Product and Technology Marketing Group. Moreover, he says, “VoIP is going to come. It will be the first service to become widespread.”

Points out Dobrec, “What’s happening now is the ‘perfect storm.’ Service providers are prepared to make the necessary capital expenditures in their access networks to increase bandwidth out to the residence or SMB. Wall Street also understands the necessity. And regulators in the US and elsewhere are pushing broadband agendas.”

What Do Consumers Want?

One connection for all. Triple play—voice, video, and data, preferably with wireless, too. “The consumer wants simplicity and flexibility,” Gupta says. “A single pipe, one service provider, and the ability to order up services and even change the width of the pipe as needed—services on demand.”

There are four general areas of consumer needs where broadband applications will flourish: communications services, information services, monitoring and management services, and entertainment services (see Figure 1). Dobrec outlines a scenario that traverses all four of these service areas: A subscriber orders up a video-on-demand (VOD) program, and shortly thereafter his phone rings. Caller ID information running across the bottom of the TV screen indicates that the call is from a parent or child, so he puts the video on hold to take the call. Afterward, the subscriber restarts the video and asks the phone to hold all other calls until it’s over. Oh, and before ordering the video, he requested higher bandwidth while it downloaded.

Beyond commingled services such as those described in Figure 1, it’s likely that customers will eventually use a combination of network- and customer premises equipment (CPE)-based services such as firewalls, which could reside in either place; storage, which would be a network service; and home-based gaming devices that include an Ethernet port and a wireless router.

Japan and Korea, Dobrec says, “see broadband as an economic stimulus.” Around the world, he adds, “as voice services move away from fixed to mobile network delivery, there’s a big drive to turn the copper serving homes back into a valuable, revenue-generating

SERVICES IN THE BROADBAND HOME	
COMMUNICATIONS SERVICES	MONITORING AND MANAGEMENT SERVICES
VoIP IEEE 802.11 phones Presence In-home key systems	Home surveillance “Nanny Cams” Security Fire, utilities, lighting (remote management)
INFORMATION SERVICES	ENTERTAINMENT SERVICES
VPN Home Networking Parental Controls Storage	Video Music Gaming

FIGURE 1 Opportunities for broadband services abound within the residence in four key areas: communications, information services, entertainment, and monitoring and management.

vehicle. And a couple of big questions are being asked: How to scale broadband service and get it to more homes and SMBs? And how can the service provider derive more revenue?”

Making Money in Broadband

Here’s one indicator of the value customers place on broadband: Last December, *The Wall Street Journal* reported that of the many US consumers shopping over the Internet, those with broadband connections did more holiday shopping online than their narrow-band counterparts—and *spent 50 percent more*. The willingness of broadband customers to spend more online might well be a measure of their willingness to spend more for broadband services overall.

Incumbent telecom and cable providers have two potentially lucrative paths to choose from in increasing their revenues using broadband over the next three to five years, Gil de Bernabé points out. One strategy is to offer more communications services and features; the other is to become involved with content. Both paths begin from the same foundation, which is what most providers now have: basic broadband connectivity, flat-rate pricing, hands-off content, basic CPE, direct sales, and mass marketing. Inefficiencies in marketing and provisioning have made it difficult for some providers to break even, although the learning curve has brought the costs of signing up and provisioning a subscriber much lower.

The Communications Services Route

After the foundation is built, the second phase of the communications services path is to improve the obvious—to add or enhance services based on the existing network. For example, Deutsche Telecom has optimized its DSL connections for multiplayer gaming by minimizing delay and jitter. Given that gaming is a real-time experience, subscribers may well want to

What's Going Through Your Network?

"To boost the profitability of broadband networks and expand their service offerings, service providers need more control of their network, especially the ability to identify and classify applications. Cisco's Service Control Application Suite enables both wireline and wireless broadband providers to guarantee performance and charge for high-usage services," says Kevin Mitchell, directing analyst, Service Provider Voice and Data at Infonetics Research.

The new Cisco SCE 1000 and 2000 Series Service Control Engines are purpose-built hardware and software systems that give service providers the information and control to monitor and manage network activity precisely; to price according to QoS, SLAs, and even applications; and to optimize bandwidth at the applications level. The systems reside at the network edge, behind the access aggregation system, the CMTS, or the B-RAS. One system can process up to 2 Gbit/s of traffic with carrier-grade performance.

The Cisco 1000 and 2000 Series Service Control Engines analyze the payload of individual packets using stateful deep packet inspection at Layers 3 through 7. They can detect virtually any network application, including Web browsing, multimedia streaming, and P2P file sharing. As a result, they can fully reconstruct individual traffic flows and the Layer 7 state of each one. Because they maintain state information, the service control engines can identify applications that employ dynamically assigned port numbers and track applications such as VoIP or multimedia streaming that involve multiple interrelated flows.

This level of traffic detail enables service providers to impose rules reflecting admissions policies or data session characteristics. Providers can also shape bandwidth at a very granular level, redirect traffic using specific protocols, and employ quotas.

pay more for better quality of service (QoS). Other subscribers may want bandwidth for peer-to-peer (P2P) file sharing and pay for it. P2P usage now accounts for some 70 percent of all bandwidth used by residences around the globe.

A provider might also offer services such as storage for digital photos, movies, audio recordings, and other large files. Network management is another offering: "Most early adopters have already set up their home networks, and the mass market is less comfortable with doing it themselves. Providers can set up and maintain a subscriber's home network, bringing in a home gateway or router and linking up several PCs along with peripherals such as surveillance cameras or 'nanny cams,'" Dobrec says. "They can also deliver firewall and other security provisions from the network. Many mass-market customers don't have the confidence or know-how to manage these configurations themselves."

Much can be done with marketing at this stage. Providers can bundle services such as local, long distance, and Internet access to maintain customer loyalty and reduce churn, and link up with video providers. "Customer churn is one of the biggest problems facing service providers in competitive markets today. One of the best ways to reduce churn is to bundle services," says Gil de Bernabé. As an example, he notes, in a 2003 analyst conference, management at BellSouth stated that when a local or long-distance customer adds just one additional service—DSL, wireless, or dialup—churn decreases by about 45 percent. "Bundling," continues Gil de Bernabé, "also allows service providers to lower the monthly price of broadband access without sacrificing subscriber profitability, because of higher ARPU [average revenues per user] and lower churn. The more services the provider offers, the longer customers tend to stay, and the more money they spend. When just one service is offered, the average life of a subscriber is 32 months. But when three services are bundled together, the average life of a subscriber doubles."

Providers can also offer tiered pricing with higher prices for higher QoS. Advanced market segmentation can target specific customer groups with different bundles of services. Bundling has been shown to increase customer retention.

Value-Added Services

The third phase of a communications services path includes value-added services. Some service providers already at this stage are reporting annual ARPU of US\$75 to US\$100 per month, and high earnings before interest, depreciation, taxes, and amortization

(EBIDTA) margins. Services offered might include the triple play—voice, video, and data; advanced home networking with features such as home management and security; and healthcare monitoring. Providers might bundle wireline and wireless services to further increase customer loyalty.

In addition, many customers have demonstrated that they will pay more for guaranteed levels of service quality; thus, service-level agreements (SLAs) will probably enter the consumer arena. Phase three may also see the provider offering more sophisticated CPE: gateways and home controllers that make it easy to provision new services such as firewalls, URL filtering, and virus checking, either from the CPE or from network-based systems. Intelligent networks will be able to recognize what application a subscriber is using and adapt QoS to it, providing high-quality performance at lower operating costs.

The Content Route

Starting from the same foundation described earlier, phase two for a provider focusing on content would include digital rights management of the content, which cable companies already do; billing systems capable of handling multiple types of content; and content delivery networks with distributed servers and storage to bring content close to the customer for better quality. Some providers are already offering pay-per-view, personal TiVo-type video recording, and VOD. Providers can deliver a menu of such offerings, individually priced, on top of flat-rate Internet access; they can also deliver different QoS levels for different content. However, this phase requires an intelligent network that can distinguish between different types of content.

In the third phase, providers begin to own some of the content they distribute, although they are probably not creating it. They can also offer Internet portals, content aggregation, subscriptions, and other features. ARPUs might not go up significantly, but loyalty will; as churn decreases, EBIDTA margins could increase.

Which Path Is Best?

As Gil de Bernabé says, “Build on your strengths.” Telecom companies might be more comfortable following the communications services route, while cable companies may find the straight content route the right model. Other types of providers—e.g., alternative or greenfield—might follow a path between communications and content. “There’s no single model that fits all,” adds Gil de Bernabé, “but all service providers will try to enhance their offerings in four areas: coverage, bandwidth, content, and services.”

What’s most important is to move quickly. The key

factor driving broadband growth is competition. In Belgium, for example, there is a strong incumbent telecom provider and two strong cable operators that have spurred one another to offer 3-Mbit/s broadband service for about 39 Euros, or US\$52. Alternative providers such as Yahoo!BB in Japan, FastWeb in Italy, and B2 in Sweden have built new networks that enable them to provide high-quality broadband at very competitive prices.

What Does the Network Require?

“Network infrastructure will undergo phenomenal changes to support broadband value-added services,” says Gupta. “It will need bigger, fatter pipes. It will also need intelligence, QoS, reliability, services on demand, and security. And it will need to be service-driven, not infrastructure-driven.”

Intelligence, he emphasizes, will be crucial. Both cable and telecom operators have limited mechanisms to identify the usage of individual users or applications and to regulate usage by bandwidth hogs. Generally, they now simply pile on bandwidth to solve shortages. This is a costly and not particularly effective solution, because P2P users, especially, use as much bandwidth as they can get. The ability to discern what application and protocol are being used will give providers the ability to deploy, identify, and bill individual customers for services such as VoIP, VOD, interactive gaming, videoconferencing, P2P applications, and IP Security (IPSec).

For example, when the next BitTorrent or eDonkey comes into use, the provider needs to be aware of its presence and identify its protocol so it can be detected when subscribers use it. File sharing is a very good example, Gupta points out, because currently about 15 percent of subscribers using it are hogging about 70 percent of capacity on broadband networks.

Part of the Cisco Service Control Application Suite, the new Cisco Service Control Engine products have already proven able to detect and map a signature for service providers (see sidebar, page 51).

Broadband-Ready, from Edge to Core

Service providers will need to deliver residential gateways as a means to connect to their broadband services as well as provide all the appropriate LAN interfaces to connect to devices inside the home, says Dobrec. “The gateway to the network really needs to be agnostic regarding the protocol the subscriber uses,” he explains. “It should accept DSL, cable, Ethernet, Wi-Fi, or any other broadband signal.”

Cisco’s broadband-remote access server (B-RAS) is available on both the Cisco 10000 Series Router and

Guiding Broadband

Responsible for setting industry guidelines and developing standards for DSL service is the DSL Forum, of which Cisco is an active and principal member. The DSL Forum numbers more than 200 leading service providers, equipment manufacturers, and other interested parties around the world. Its reports, such as TR-059 released in 2003, in particular, are guiding the future of DSL in all its flavors, which specify an IP-centric network architecture and general requirements for key components such as B-RAS, and TR-092, which defines B-RAS requirements much more precisely.

TR-059 recommendations for network architecture spell out new infrastructure capabilities that will be needed as well as interoperability requirements for these networks. A key provision is that QoS be delivered via IP, rather than ATM, which enables providers to deliver, for example, VoIP and gaming within certain latency and jitter parameters.

Focusing on the B-RAS, TR-092 specifies details such as interfaces, protocols, traffic management, policy management, operations, and many others; in addition to the commonly thought-of services such as triple play, a compliant B-RAS system will also support services such as multicasting and VPNs.

Cisco has actively contributed to both documents, as well as other works of the DSL Forum. "Cisco's participation enables it to remain a leader in IP technology and offer expertise to its customers," Gupta says. "Cisco has, in fact, more background and expertise in IP QoS than any other vendor." And by participating in IETF, ITU, Metro Ethernet Forum working groups as well as the DSL Forum, "Cisco helps bridge between these groups, thus helping to ensure compatibility of standards and visions."

the Cisco 7000 Series Router. Fully compliant with the newest DSL Forum standards (see "Guiding Broadband"), the Cisco 10000 Series Router significantly reduces the cost of delivering broadband services with a capacity of 60,000 simultaneous sessions

tightly coupled with QoS and other features. The Cisco 7600 Series Router, especially appropriate for metro Ethernet networks, handles up to 32,000 sessions per module while delivering Ethernet services, and the Cisco 7200 and 7300 Series support up to 16,000 subscriber sessions per chassis. For cable providers, Cisco offers the industry's most complete line of DOCSIS-compliant cable modem termination system (CMTS) solutions that can support differentiated services such as VoIP, gaming, video, and bandwidth-on-demand for large numbers of subscribers.

Cisco continues to work with telecom and cable providers to create service solutions that enable these providers to offer the most appealing and most economically delivered suites of services. Examples include its Broadband Local Integrated Services Solution (BLISS), the Cisco Gigabit Ethernet Optimized VOD Solution, and its PacketCable Multimedia capability.

Providers can also offer home-based systems such as the new Linksys Wireless-G ADSL Gateway and Wireless-G Cable Gateway, together with peripherals such as the Linksys Analog Telephony Adapter, Wireless-G Gaming Adapter, and Wireless-B Internet Video Camera.

In addition to offering a range of end-to-end products, Gupta says, Cisco works with individual providers to develop broadband offerings that suit their own business plans and markets, to help them offer broad, creative, and reliable services. "We're with the provider throughout the service cycle, from planning the infrastructure, designing the service offering, implementing it, and optimizing it," he says. "We can help improve service delivery, which increases customer satisfaction, which, in turn, increases revenues and profit growth." ■

FURTHER READING

- Cisco Broadband Aggregation Solution
cisco.com/go/broadband
- Cisco Cable Solutions for Service Providers
cisco.com/packet/171_16a1
- "Driving Revenues in Consumer Broadband"
cisco.com/packet/171_16a2
- "Consumer Broadband: The Path to Growth and Profitability"
cisco.com/packet/171_16a3
- "Connected Homes: Essays from Innovators in Consumer Broadband"
cisco.com/packet/171_16a4
- Cisco 1000 and 2000 Series Service Control Engines
cisco.com/packet/171_16a5
cisco.com/packet/171_16a6

Wireless Patrol

Cisco Metropolitan Mobile Networks enhance public safety and law enforcement.



STAYING CONNECTED Law enforcement agencies and city governments can access rich applications such as video surveillance from virtually any location, allowing the entire network to stay connected at all times.

By David Baum

In an era of growing concern about law enforcement and security, policy makers and government administrators face constant pressure to improve their ability to respond quickly to criminal activities and security threats. In some cases, new technology is a catalyst that enables them to improve operations.

Consider the exciting advancements taking place in mobile computing, as city governments and law enforcement agencies deploy Cisco Metropolitan Mobile Network (MMN) solutions to more effectively manage services in local communities. MMN technology from Cisco is extending the edge of the IP network and enabling new types of applications for the defense, public safety, and commercial transportation markets. These secure, standards-based broadband mobility solutions integrate wired and wireless IP infrastructure across a city or regional area so authorized users can access crime databases, fingerprint files, photo images, and other pertinent information from any location.

“Real-time access to law enforcement information empowers organizations to make faster, better, more informed decisions, ultimately increasing their productivity and effectiveness in the field,” says David Yuan, a mobility manager in Product and Technology Marketing at Cisco. “Rich applications such as IP video surveillance and government agency databases can now be accessed in real time from virtually any location, allowing an entire network—not just a single client—to stay connected at all times.”

For example, in the state of Washington, the City of Everett Police Department (everettpolice.org) has deployed Cisco MMN solutions to supply network-roaming access to its patrol cars. The in-vehicle network supports rich applications and services to make offenders’ photos, communications tools, scheduling tools, and management tools available to officers—without having to return to the station. “It’s great to have access to records and be able to network to other parts of the country to look for outstanding warrants or similar unsolved crimes,” says Sgt. Boyd Bryant, police sergeant and public information officer for the City of Everett Police Department, and supervisor of the department’s technology projects.

According to Bryant, the average police officer spends about four hours per 12-hour shift in a police station—partly because that’s the only way they can access police records and other computer-based information. “Officers are out of touch with the community during that process,” he adds. “Their eyes are no longer engaged in what’s happening on the street.”

The City of Everett wanted to devise a mobile wireless strategy and associated networking infrastructure that could support rich applications and services so it could keep officers fully connected while in the field. After evaluating technology and speaking with other police departments, Everett chose Cisco and partner Northrop Grumman to configure and deploy the department’s new mobile network solution. The solution is built on Cisco 3200 Series Mobile Access Routers, Panasonic laptops in the squad cars, and Cisco Aironet bridges.

Broadband Wireless Solutions

Thanks to its compact size and rugged design, the Cisco 3200 Series is easy to deploy in public safety vehicles. It can withstand the harsh demands of a mobile environment, and it uses the Mobile IP standard to allow network nodes to roam across multiple wired or wireless networks while maintaining live connections. “The Cisco 3200 extends the edge of the IP network into the field, allowing users to maintain secure data, voice, and video connections while their vehicles are in motion,” explains Marc Bresniker, product manager for the Cisco 3200 Series.

Both the Cisco 3200 Series routers and the Cisco Aironet bridges use IEEE 802.11b/g broadband wireless technology to supply much higher bandwidth than the City of Everett’s legacy system could deliver—up to 54 Mbit/s. This enables Bryant and his team to deploy rich applications and bandwidth-hungry voice and video communications. “The flexibility of the Cisco Wireless and Mobile Access Router solution is the key,” Bryant says. “Nothing else available on the market allows you to incorporate a variety of standards-based, network-connected applications with field printers, cameras, presentation tools, and an almost endless variety of other devices into a network on wheels.”

With the ability to access photos of offenders and crime scenes, instantly updated operational documents, and outstanding warrants—all delivered to the community through a networked multimedia theater on wheels, Everett police officers are not only better equipped to apprehend criminals, they are more connected to the community than ever before. Additionally, the department’s mobile command vehicles can serve as core networks in a disaster zone—supporting emergency personnel on the scene, transmitting observations back to a command center, and controlling mobile video cameras in areas where officers can’t go.

“Ultimately, the goal is to use 802.11 to provide overlapping coverage areas across the majority of the city,” says Bresniker. “Future use of the network could include other city agencies such as the fire department and department of public works.”

Flexible Deployment Options

As the City of Everett is demonstrating, Cisco Metropolitan Mobile Networks use the 802.11 standard to provide overlapping coverage areas. These networks use Cisco Aironet 802.11 access points and bridges and the Cisco 3200 Series as an outdoor wireless router to create the coverage areas. Cisco Aironet 1400 Series wireless bridges are used mainly for point-to-point and point-to-multipoint backhaul links. (Backhaul links are useful where fiber or wired lines are not available.) Cisco Aironet 1300 Series access points and bridges can be used either as bridges or access points. The Cisco 3200 Series, used in rugged outdoor enclosures, combines multiple 802.11 radios with Layer 3 routing to provide added flexibility in where and how coverage can be deployed without the need for wired backhaul.

In addition to in-vehicle deployments, wireless access points, bridges, and wireless routers can be placed on city buildings, fire stations, communications towers, or—as will soon be the case in London’s Westminster borough—on top of light poles along with closed-circuit TV (CCTV) cameras.

The Westminster City Council (westminster.gov.uk) is using Cisco MMN technology not only to fight crime, but also to provide real-time information to city workers so they can better manage street services such as parking, premises licensing, and environmental waste.

The Westminster City Council’s initial motivation was to create a flexible platform for Wi-Fi-based monitoring to supplement existing CCTV systems. As part of an extensive pilot project, the city is deploying the Cisco 3200 Series with 802.11b/g wireless capability to enable these wireless connections. The equipment is being mounted on lampposts and buildings to extend the metropolitan-area network. Flexibility is a chief advantage of this mobile infrastructure. Seventy Wi-Fi access points and 40 CCTV cameras will be deployed around Westminster. Westminster expects that the deployment will pay for itself within two years and that the productivity of street-based services will improve by around 20 percent.

Because of the planning, coordination, and regulation needed to move traditional CCTV cameras—which involves moving street poles, digging trenches, and running cables—it used to take from three to six months to set up a new monitoring station. With the Cisco mobile infrastructure, the Westminster City

Council can relocate its monitoring systems within three hours—simply by provisioning a cherry picker and moving the box that contains the router and related wireless gear. For example, a camera was quickly moved into position to record the activities of drug dealers, who were later apprehended. As a result, the system helped eradicate drug dealing in the coverage area.

The Cisco 3200 Series can use multiple 802.11 interface cards configured as bridges or access points, as well as directly connect application hardware such as cameras and sound sensors. City Guardians and other Westminster personnel can use personal digital assistants (PDAs) and laptop computers to wirelessly tap into any of the cameras, telemetry devices, and noise monitoring devices that are integrated into the network. IP is the fundamental enabling technology in the project, combined with MPEG2 encoders, a PC-based management system, and sufficient storage to hold images and other data from all the cameras for a five-day period.

The Westminster network can also be extended to include applications such as automated number plate recognition (ANPR) software, which works with the camera footage to recognize license plates and compare them with a central database. Mobile fingerprint recognition systems can be enabled so that law enforcement officials can make instant arrests with positive ID.

Andrew Snellgrove, network manager at Westminster, says the city in effect extended its corporate data network from Council buildings (a LAN/WAN environment) into the street (a wireless metropolitan network). “Running multiple applications over the same IP infrastructure provides the tools to take the Council to our residents. The wireless CCTV implementation provides mobile workers access to back-office systems at broadband speeds and the ability to manage on-street assets. The radio coverage in the wireless domain has exceeded the plan established by the RF survey and has given us greater flexibility to deploy devices. Historically, our CCTV cameras have been at fixed locations due to limitations with the technology. The Cisco 3200 Series Mobile Access Router allows mobile cameras in vehicles to be connected online, together with viewing of video streams from mobile devices from the street.”

“A particular strength of the Cisco 3200 is the capability to support traditional and wireless technologies. Multiple radios and the use of telemetry are essential for our implementation, and this is also provided by the Cisco 3200,” he adds.

Using Wi-Fi technology raises security issues. At Westminster City Council, following testing, Snellgrove is confident that security is more than adequate to protect the content the Council delivers. “We use standards-based encryption and 802.11x, combined with firewall and Cisco VPN technology. For an enterprise

Wi-Fi network, a robust security strategy must be adopted and this needs to have a multitiered approach while balancing the network design between security and performance. Network security has to be fit to a purpose; unnecessary layers have a detrimental effect on network usability. Wi-Fi in itself is not secure, it is how you deploy it and what you deploy.”

Broad Technology Horizons

Mobile routers enable a shift in the way organizations communicate and share information. Solutions such as Cisco IP communications, Cisco MMN, and Cisco integrated security provide the infrastructure on which regional and local governments can base new ways of doing business. These technologies can be used not only to bolster law enforcement agencies, but also to create intelligent information networks and connected communities. They can be applied to many public sector agencies, such as transportation and public works, as well as to other industries altogether, from telemedicine to factory automation.

“MMN technology is a good fit anywhere in which secure, ruggedized wireless equipment is useful,” points out Yuan. “For example, ruggedized wireless equipment might be required at a port or factory. Similarly, fleets of vehicles in a campus environment could be equipped with mobile infrastructure to extend the learning or work environment.”

In short, Cisco MMN solutions can be applied wherever users can benefit from a secure, scalable, broadband network that integrates wired and wireless IP infrastructures.

“In public safety and government service, you need to understand three things,” says Boyd Bryant. “First, you are becoming more technology-dependent for your efficiency, so your systems must be backed by a corporation you know will be there several years from now. Second, you have to be certain you have the capacity to transition to another service if you need to. Third, you need to be able to support redundant network connections to help ensure that you’re not only covered, you’re resilient. The Cisco solution is the only one we’ve found that gives us all of that.” ■

FURTHER READING

- Cisco 3200 Series Wireless and Mobile Routers
cisco.com/packet/171_7a1
- Cisco MMN Flash Demo
cisco.com/packet/171_7a2
- Cisco Aironet 1300 Series
cisco.com/packet/171_7a3
- Cisco Aironet 1400 Series
cisco.com/packet/171_7a4

When Organizations Converge

Companies and employees that have successfully adopted converged networks share lessons learned.

By Rhonda Raider

When organizations begin sending voice over their data networks, two previously separate responsibilities converge. “Traditionally, the IT department took care of the data network, and a separate telephony department handled phone traffic and faxes,” says Alex Hadden-Boyd, director of marketing for IP communications in the Product and Technology Marketing Organization at Cisco. “When companies adopt IP communications, both types of expertise remain needed and both staffs have the opportunity to learn something new.”

One Strategy: Infrastructure and Services Teams

A textbook example of a successful organizational transition to IP communications is Liz Claiborne (lizclaiborneinc.com), the New Jersey-based clothing retailer. Before deploying Cisco CallManager for IP telephony and call center operations in 2003, the company’s datacom and telecom teams rarely interacted. However, the two teams already reported to the same director, thanks to the foresight of Vice President of Information Technology John Kovac and IT Director Anthony Iadisernia. “We knew there would eventually be convergence, so we organized under one director from the outset, to avoid turf wars,” says Kovac.

Once the project was approved, Iadisernia shared his plan with staff members: the two organizations would merge to become a single IP communications team comprising an infrastructure group and a services group. Each group would include nearly an even number of members from the datacom and telecom groups. The infrastructure group would manage implementation and project planning, while the services group would take charge of day-to-day monitoring, reporting, provisioning, vendor relations, and adds, moves, and changes.

Today the two groups work side by side. “We intentionally seated the teams together so that we could informally leverage each team’s experience and technical knowledge,” says Kovac. All that remains of the original division along datacom and telecom lines is one chief architect for data and another for voice, who also sit together and report directly to Iadisernia.

Before and during the transition, Liz Claiborne cross-trained employees on voice and data. “If we upgraded the network, for example, we’d assign one or two voice specialists to observe and participate,” says Iadisernia.

When IT Inherits IP Telephony

Adapting the IT organization to support IP communications is somewhat simpler for companies that previously outsourced PBX operations to vendors and partners, as was the case for Tequila Herradura (herradura.com) of Mexico. Previously, Tequila Herradura had managed its data network with inhouse resources and outsourced management of its telecommunications network and devices. Therefore, when the company migrated to IP communications, the first task for IT Director Irvin Valencia was to train inhouse IT staff on IP telephony management, using the services of Hewlett-Packard. “The IT staff found it easy to learn to use Cisco CallManager and the Cisco IP Phones, and were excited to learn a new technology,” says Valencia. “Working with IP telephony makes their jobs more interesting and improves their value in the workplace.”

Adding voice traffic to the network did not significantly increase the workload for the Tequila Herradura IT group. “In fact, it’s easier to manage the converged voice and data network than it was to manage data alone because we’ve begun using Web-based tools like Cisco CallManager Administration,” says Valencia. “The voice aspect works by itself almost 100 percent of the time.” Troubleshooting, in particular, is faster. “Before, identifying the source of a failure took time because we had to schedule an appointment with the PBX supplier,” says Valencia. “It was a waste of time and bad for productivity. Now our own staff can identify and resolve problems just as they’ve always done for data.”

Managing Internal Change

Tequila Herradura introduced IP telephony to its employees gradually, starting in February 2002 with just 50 employees. “We emphasized to the participants that they were an important part of the organization, and people were very eager to volunteer,” says Valencia. “Soon, seeing the benefits of IP telephony, such as the directory, four-digit dialing, and unified messaging, more and more people in the

company approached us to participate, including the general director of our company.” By March 2004, every employee in the organization had a Cisco IP Phone. “Productivity is difficult to measure,” says Valencia. “But the benefits of IP communications were very evident to our employees, leading to enthusiastic adoption.”

Like Tequila Herradura, Liz Claiborne took pains to ensure its employees knew what to expect from Cisco IP communications. Members of the services group met with executive assistants before the transition to understand how to use their phones and which new IP telephony features would benefit them. “We provided lots of hand-holding to ensure high adoption,” says Iadisernia.

Cisco’s Own Experience

Liz Claiborne and Tequila Herradura benefited from lessons learned by Cisco during its own migration to IP communications in 1998–2001. At the time, Graham Hosie, now a Cisco IT director, was senior manager for global voice services. When asked to head up an effort to deploy Cisco CallManager in place of legacy TDM switches, Hosie assembled a 60-member global team, whose members contributed expertise in voice, data, and hosting.

“Until that time, Cisco’s voice and networking groups were located on the same floor but had no idea what the other did,” says Hosie. “To deploy a multiservice network, we needed both groups to work together for planning, process, and support.” If a call didn’t go through, for example, Cisco needed a process and organizational model to determine the cause and assign the appropriate people.

“If the PBX Is Going, What About Me?”

When companies decide to adopt IP communications, IT staff—especially telephony specialists—can wonder about the effect of the transition on their jobs. In fact, telephony specialists possess a collection of skills other than PBX administration that remain essential for companies that adopt IP communications.

- *Deep understanding of end users’ business needs*, which can be as simple as whether a manager’s and administrative assistant’s phones ring simultaneously or one after the other, and as complex as casual contact centers—where employees interact directly with customers in addition to their other job responsibilities. “These real-world aspects of telephony don’t disappear just because you change the infrastructure and devices,” says Cisco’s Hadden-Boyd.
- *Expertise in working with carriers*, a skill that remains indispensable.
- *Interpersonal communication skills*. “Network engineers usually don’t need to concern themselves with the effect of network change on users, because the change is generally invisible,” says Hadden-Boyd. “But the moment you change a password or require users to press a different key sequence, communication with

users becomes very important, and telephony specialists have honed this skill.”

“When we discussed the transition, we always emphasized the continuing importance of the telecom staff’s role in the organization,” says Liz Claiborne’s Kovac. “We also explained that they would have the opportunity to learn a new technology that would help them in their careers.”

Cisco followed a similar approach with its own IP communications deployment. “Through training and informal brown-bag lunch sessions, the staff on the voice side started to understand that their careers were not over,” says Hosie at Cisco. “On the contrary, they’d get a boost by advancing to the leading edge of new IP telephony technology, and their knowledge was critical to the effort. They understood that their long-term career success hinged on their also learning a networking skill set.”

“Ultimately, IP communications is an opportunity for both sides to learn something new,” says Doug McQueen, a solution implementation manager at Cisco. “Network staff want to learn about voice, and voice staff want to get their teeth into the network.”

Hosie enlisted the support of a senior engineer who was known for his technical knowledge to break down the barriers between the voice and networking organizations. Almost immediately he physically relocated the groups so that they worked side by side. "Co-locating the groups fostered idea sharing and an understanding of the other group's concerns and expertise," says Hosie. "Networking and telephony people began building personal relationships, eating lunch together, getting to know each other, and realizing that the other team was not a threat, but rather provided complementary skills to achieve common goals."

Next, to better address the technology and business needs of IP communications, Cisco reorganized IT into three groups: IP Telephony Operations, Emerging Technologies, and Foundation Technologies, which includes both voice and data. "We've gotten away from networking and voice silos," says Hosie. "The exception is that we've retained a voice services group for client-facing activities." Cisco deliberately used traditional terminology, "voice services," so that employees know who to call when they have a question about their phones or unified messaging.

"The organization that is customer facing is the overall owner," Hosie continues. "If an employee has a question or problem with their IP telephone, they want to talk to the telephony people. They don't care that it's running on the data network."

Blurring Turf Boundaries

At the outset, Cisco voice and data engineers hesitated to provide each other with access to their equipment. "In particular, the networking group had reservations about the voice people touching the routers that provide access to the PSTN, and voice people didn't think data people should have access to the Cisco CallManager," says Hosie. "People worried about someone from the other group inadvertently bringing down the network." Training alleviated these concerns. "We learned to understand, respect, and play well in each other's sandbox," says Dennis Silva, an IT manager for IP telephony at Cisco.

Like management at Liz Claiborne, Cisco provided both formal and informal training so that the networking and voice groups could learn the other technology. The informal training, which Hosie describes as "cross-pollination," included brown-bag lunch sessions to talk about new technologies. "We took people with extensive voice and limited networking experience, or vice versa, and trained them to support converged voice over IP," says Hosie. Employees appreciated the reassurance about the continued importance of their jobs, as well as the opportunity to learn new skills. "We didn't lose anybody because of concerns about job loss or change," Hosie notes.

A New Perspective on Change Management

Among the biggest cultural shifts to emerge from IP communications at Cisco was a different view of change management. Traditionally, networking people had the luxury to make changes to the network during the day because people rarely notice a brief interruption. That's not the case when voice travels over the network because even a brief network outage might interrupt live conversations.

"Voice is the most visible application we'd deployed on the network," says Silva. "In the past, if we had to bring down the network to make a change, we could do it at 5:30 p.m. and almost nobody would notice. But if someone does a failover to a backup router when voice runs on the network, you can potentially kill thousands of phone calls. We had to come to an understanding about when we could do changes." Voice-impacting network changes are done after 9 p.m. local time.

New Job Opportunities

Rather than eliminating jobs, converging the data and voice networks often creates new job opportunities (see sidebar, "If the PBX Is Going Then What About Me?"). "Our staff needs didn't dwindle through convergence," says Cisco's Hosie, citing that convergence brings on new capabilities in call processing, Cisco Unity Unified Messaging, Cisco IP Contact Center (IPCC), Web collaboration, and videoconferencing. Similarly, Liz Claiborne retained the same size staff. Because it achieved some efficiencies by converging its network, the company was able to redeploy existing staff to projects that had languished on the back burner for want of people, such as training, implementation, project planning, maintenance, and support.

Hosie's team began speaking to key individuals about these new job opportunities in the planning stages of the transition. "We coined the term 'IP telephony engineer' to capture the blend of existing and new skills we would need," he says. For example, someone who used to be a PBX technician and worked with a PBX connected to the PSTN could become an IP telephony engineer who worked with Cisco CallManager connected to the LAN and WAN as well as to the PSTN. "Now there's more sophistication and more excitement about new opportunities for learning," says Hosie. IP telephony engineers not only have to understand PBX functions on Cisco CallManager, but also how the network is used to route voice packets to the appropriate end device.

Continued on page 64

Lessons Learned

Make Organizational Changes Early

Liz Claiborne's Ladisernia suggests that converging the telephony and networking groups should occur either at the same time as the network convergence or immediately afterwards. His organization had intended to merge the groups about a year after implementation, but ultimately did it after 10 months.

Retain Subspecialties Within the IP Communications Group

Even with its converged network, Cisco distinguishes among issues that concern PCs, networking, hosting, and telephony applications. "If a client uses Cisco IP Communicator, a softphone application that enables voice communication from a laptop, the IP Telephony Operations team does not manage the laptop. Instead, they partner with the PC support team that manages the laptop," says Holloman.

Communicate Frequently and Openly with Networking and Telephony Staff

Cisco, Liz Claiborne, and Tequila Herradura all credit the success of their transition in large part to a policy of open communication. "Had we not been so open about our

strategy, we might have had a lot more concern about job roles," says Kovac. Adds Hosie, from Cisco, "Don't underestimate your audience or try to mask what you're doing. Explain the business drivers for the transition to IP communications, and clearly communicate the career opportunities. This is the opportunity of the future." The importance of communication applies to successful user adoption as well. "To just say, 'Here's a new phone,' is not as successful a strategy as saying, 'Now you'll have incoming caller ID, a directory, and four-digit dialing,'" says Hadden-Boyd.

Organize Under One Director

"Cross-training on voice and data technologies was very successful, and this might have been a difficult road if both data and voice didn't report to me," Ladisernia adds. "I had to ensure that voice traffic received priority, which was easy to do when both groups reported to me. If we had been organized with voice reporting into facilities whereas data reported to IT, it might have been a difficult battle."

While the voice and data specialists who come together to support IP communications increase their worth by learning something about the other technology, they still retain their main focus. "IP telephony is an application that uses the IP network as its transport," says Silva. "When an organization adopts IP communications, telephony engineers continue to apply their same skills, just at a different layer."

The same applies to network engineers. "Network engineers don't necessarily want to become application specialists," says Marc Holloman, global operations manager for Cisco Intelligent Network Services. "Their specialty is the plumbing underneath. There is a need to learn about the voice application, and some of our engineers initially expressed a little of 'We weren't hired for this sort of thing.' But at the end of the day, network engineers viewed IP communications as an opportunity to learn and become even more valuable employees."

Valuable Diversity

Liz Claiborne's Kovac views his company's experience with IP communications as underscoring the value of diversity. "The reason for diversity is to get different perspectives, ultimately resulting in better decisions," he says. "Having created a single IP communications team, we're getting different technical perspectives, which increases the strength of the IT organization." Valencia, of Tequila Herradura, agrees. "Our IT staff is thrilled that they have the opportunity to learn something new." ■

FURTHER READING

- Cisco IP Communications
Cisco.com/packet/171_7a1
- "Migrating to IP Telephony?" (*Packet*, Second Quarter 2004)
cisco.com/packet/171_7a2
- Cisco IT@Work: IP Communications
cisco.com/packet/171_7a3

Technology Wakeup Call

Hotel operators deploy converged IP networks to vanquish fierce competition and indulge sophisticated guests.

By Joanna Holmes

The hospitality industry took a beating between 2001 and 2003 as a plague of difficulties—the SARS epidemic, the September 11, 2001 terrorist attacks on the US, and the economic downturn—daunted all but the most intrepid travelers. 2004 proved to be a better year, but even in fatter times, hoteliers are challenged to stay competitive amidst a new generation of business demands. Enter the IP network.

With IP technology, hotel operators are finding a quartet of compelling opportunities: creating new revenue streams, improving the guest experience, improving operational efficiency and staff productivity, and reducing the costs associated with real estate.

Hoteliers' Hurdles

To understand what converged, multiservice networks have in store for the hotel industry, take a look at the challenges that impede hotel profitability.

Fostering brand loyalty and differentiation—Hotels today face increased competition among the major brands for a limited number of travelers. Building brand loyalty is pivotal to profitability.

Increasing revenue, maintaining margins—Yesterday's revenue sources—high-speed Internet access (HSIA), for example—are standard (and gratis) fare today. Moreover, guests are abandoning hotel telephones services in favor of personal cell phones.

Increasing guest satisfaction—Guests generally expect the same business or entertainment technology in their hotel rooms as they employ in their homes and offices.

Reducing OpEx, improving efficiency, and boosting productivity—Hotels, like all businesses, must continuously seek new ways to lower overhead and improve processes.

Improving guest safety—Recent, deadly attacks on hotels have made guest safety a new priority. For more information on this growing concern, visit cisco.com/packet/171_7c1.

Problem-Solving with Technology

Many well-respected hotel chains worldwide, such as the Sheraton, Holiday Inn, Mandarin Oriental, and Crowne Plaza, have already invested in converged IP networks, retiring their legacy PBX systems and



FIVE STAR NETWORK In Poland, the Sheraton Krakow's IP network includes 600 Cisco IP Phones and wireless connectivity.

moving voice networks onto modern, streamlined IP telephony infrastructures. Products and technologies for wireless networks, as well as Cisco's IP communications solution suite, enable hotels to offer highly personalized services for guests, plus Wi-Fi connectivity throughout hotel premises for guests and staff.

With guest-friendly, in-room phones such as the Cisco IP Phone 7970G, which features a high-resolution color touch screen, hotel IT staff can create a personalized interface that enriches the guest's experience and builds loyalty, while opening limitless revenue opportunities. IP phones provide hotels with flexibility. The soft buttons, which allow the hotel to constantly change the phone features and upgrade services, are a major leap from the old analog hotel phones with hard-coded buttons and paper templates (and which generally support only dialup Internet).

Behind the scenes, new models of efficiency and automation are at work. The same network infrastructure that allows a guest to e-mail a voice message to a colleague over a Cisco IP Phone helps the hotel itself to monitor building access, manage climate controls, and perform myriad tasks that previously have run over disparate systems (and required separate staff members to manage them).

Hong Kong's Langham Place Hotel

The Langham Place Hotel, a 665-room, five-star business and leisure property, opened in July 2004. The hotel invested in an advanced IT system that includes a Cisco voice and data network and provides high-speed wired and wireless Internet access throughout the hotel.

“IP technology gives us a competitive advantage that lets us grow our market share and stay ahead of the curve,” says Brett Butcher, managing director for the Langham. “And as Asia’s most technologically advanced hotel, we believe we’ll continue to win the hearts of business travelers who are technologically savvy. Our guests appreciate the speed and convenience this hotel provides.”

Acknowledging the hospitality industry’s need to provide highly personalized service, Butcher says, “We are able to achieve even more than usual with the hotel’s IP technology.”

Langham guests can use their mobile Cisco IP Phones anywhere in the hotel. “Guests are impressed with our ‘Wi-Fi’ bubble, which offers wireless broadband connectivity in every part of the hotel. We receive excellent feedback on our IP telephony; guests find the features of the phone—including SMS [Short Message Service] and Web connectivity—very helpful,” says Butcher.

The first year’s results are not in yet. However, says, Butcher, “Based on the pure exposure we’ve gained in the market place as an industry leader, we’ve already made a handsome return on investment.”

Mandarin Oriental, New York

Nick Price, CIO and CTO for the Mandarin Oriental Hotel Group, also chairs the forward-looking Hotel Technology Next Generation In-Room Technology workgroup (htng.org). Perhaps that’s why the chain’s recently opened New York property is a 38-floor showcase for today’s most leading-edge hotel applications.

Opened in late 2004 in Manhattan, the Mandarin Oriental is New York’s first hotel to offer plug-and-play live broadcast capabilities, providing guests with services for live television feeds, videoconferencing, and Web streaming. The hotel’s in-room offerings set new industry standards by featuring systems as technologically advanced as they are user friendly. (See “Mandarin Oriental’s IP-Enabled Entertainment System,” at cisco.com/packet/171_7c2.)

Revenue Revisited: Ousting the PBX

According to many in the hotel business, the role of the cell phone has been to make the hotel PBX obsolete. “The revenue opportunities that have existed in hotels have historically been very significant,” says Price. “Hotels routinely made half-million dollar investments in PBX technology in full expectation that they would recoup that money handsomely.” But those returns are rapidly disappearing as guests bypass the hotel phone in favor of personal cell phones.

“If people aren’t making external calls through the PBX because they have cell phones,” Price observes. “You have to question the very presence of a PBX in the hotel,” he continues.

Thus hotels face two fundamental choices, he says: Either reduce PBX costs (and much PBX functionality), or work within the same cost structures to introduce features and functionality that compete with the guests’ cell phones. “And that is the role of IP telephony,” says Price.

“We’re charting uncharted waters,” he continues, asking rhetorically, “How do we bring guests back from a cell phone experience?” The Mandarin’s decision on this point was to optimize its in-room handsets for conference calls. “Conference calling is widely done in hotel bedrooms today—but all that hotels typically see of it is

either a free 800-number call or a 50-cent charge for the 800 call number.” Price describes how the Mandarin Oriental’s IP telephony system replaces the bridge normally used for conference calls, creating a genuine revenue opportunity for the hotel with each conference call.

Also critical to weaning guests away from cell phones, says Price, is the hotel handset itself. “We have a very superior, off-hook hands-free experience with the Cisco handset, and that contributes significantly to the guests’ desire to use it,” Price observes, noting that the phones’ usability experience compares favorably with cell phones, but offers far more features. “The Cisco IP Phone is clearly the handset of choice, particularly if you’re doing conference calls or long calls,” Price says. “We anticipate that the feature function and services we can add onto it will enable us to reverse revenue losses—and do so in a legitimate, meaningful way that adds value for guests.”

Whether PBXs have had their day is a point on which Price feels strongly. “IP telephony is probably the only opportunity we have to stop our revenues going to zero,” he says. “I don’t see any other traditional telephony provider or technology [in hotels] that currently presents any service beyond a dial tone. And if that’s all you’ve got, users are not interested.”



GUEST FRIENDLY Cisco IP Phones with color touch screens are helping hotels stay competitive.

Mandarin Oriental, New York, invested significant research in the installation of live broadcast capabilities. The Time Warner Center, where the hotel is situated, features advanced broadcast cable installations that allow guests and hotel IT staff to bypass satellite trucks or the lengthy process of contracting a camera and sound crew. Guests can book the services to broadcast content such as live newscasts or keynote speakers for live Web streaming.

Cisco IP Phones provide enhanced in-room capabilities, including color touch screens with language options (based on stored user profiles) in some of the top suites. Conference call capabilities include hosting calls for as many as six parties at once. One of the most effective ways the Mandarin Group gains differentiation for its New York property is by offering guests a much-needed business service: “We’re enabling self-provisioned conference calling, which is extremely difficult on typical hotel phone systems,” says Mandarin Oriental’s Nick Price.

With its Cisco platform, the Mandarin Oriental’s converged guest services network allows the hotel to add services as needs arise, making expansion capabilities limitless. The interlinking of the various hotel systems allows operations to run more efficiently, and that lowers operating expenses and makes for happier guests.

Sheraton Krakow

June 2004 saw the opening of the Starwood Group’s Sheraton Krakow, the first international five-star hotel in this Polish city. The Sheraton Krakow features an integrated Cisco IP communications system that leverages voice over IP (VoIP) and prominently features Cisco color IP Phones in every guest room. The network comprises Internet, data, and voice access for all guests and hotel staff, and includes more than 600 IP telephone handsets and wireless connectivity.

“More and more often, guests require Internet access, and companies themselves now expect their employees to stay in touch with their colleagues left in the offices,” says Warwick Gunning, general manager of the Sheraton Krakow.

Addressing these needs, the Starwood Group deployed a converged IP network leveraging Cisco products and technologies. The resulting infrastructure is an integrated Fast Ethernet network that can carry data, voice, and video, while enabling future services such as TV over IP.

Wireless LAN access is available throughout all hotel common areas. The Sheraton Krakow operates its own private wireless network for staff voice communication through wireless IP phones. The hotel also uses the wireless network for public access, creating a separate virtual LAN for high-speed wireless networks for groups of guests.

All 232 of the Sheraton Krakow’s rooms are equipped with the Cisco IP Phone 7970G with color touch screen. As well as providing Internet access, these phones allow guests to send e-mail, check hotel bills, book meeting rooms, and access a raft of useful information services.

While making a positive impression on guests, the Sheraton Krakow’s Cisco infrastructure also improves the hotel management tasks. “The IP network allows our staff to work more efficiently, resulting in cost savings and further helping our profitability,” says Gunning.

The IP infrastructure is based on Cisco CallManager call-processing servers and uses software and the V/IP Suite Server from Cisco partner Nevotek. The key function of the whole solution is its integration with the hotel’s property management system through Nevotek’s V/IP Suite. This software enables hotel staff to use a phone as a simple terminal on which they can communicate room status or mini-bar restocking needs.

The Nevotek V/IP Suite on the large, colorful LCD displays of Cisco’s IP phones gives guests touch-screen access to a wealth of services. The screen works as an Internet browser, using popular industry standards such as XML and HTML, so the whole system is easy to enhance in the future, with virtually no limits.

Making the Move to Converged IP Networks

To make a case for implementing IP telephony, says the Langham’s Brett Butcher, “First, you have to want to provide your customers with cutting-edge technology and understand the inherent value of this proposition in attracting and maintaining your clients. And second, you have to get close to what the technology can now offer and envisage what the future could hold.” By migrating hotel services and legacy voice systems into a converged Cisco IP infrastructure, that future can hold unlimited possibilities. ■

Metro Ethernet Coming Your Way

Metro Ethernet gives providers the flexibility and QoS to deliver *real* broadband services—and customers are buying.



By Janet Kreiling

Ethernet is slowly but surely making a major dent in metro networks. Between 2003 and 2007, Ethernet will greatly impact metro telecom equipment spending, accounting for an estimated US\$24.9 billion over the five-year period, according to Infonetics Research. Each year during this period, metro Ethernet will account for a larger portion of metro capital expenditures (CapEx), driving a projected compound annual growth rate (CAGR) of 27 percent. Technologies such as Resilient Packet Ring (RPR), Multiprotocol Label Switching (MPLS), and very-high-data-rate DSL (VDSL) are paving the way for Ethernet to take its place as a respected, telecom-grade option for metro networks.

“One of the most important drivers of the metro Ethernet evolution is that both enterprises and residential customers are increasingly seeking customizable services from their providers. They not only expect higher bandwidth connectivity services, but also want the bandwidth delivered at greater levels of granularity,” says Wei Wang, product marketing manager for metro Ethernet in Cisco’s Product and Technology Marketing Group. Ethernet as a technology offers many benefits. Chief among them, notes Wang, is that service providers can deliver bandwidth up to 10 Gbit/s to their customers, support demanding applications, and tailor the bandwidth

to deliver performance that meets the needs of specific business applications. Metro Ethernet also enables providers to seamlessly offer new services, such as videoconferencing, managed storage, and online interactive gaming, to enterprise and residential customers in one network. Additionally, notes Wang, “New and improved intelligent Ethernet equipment delivers advanced network security, and rich quality of service [QoS] functionalities allow service providers to differentiate their offerings from competitors, improve their profit margins, and generate more revenue over the long run.” For the past few years, service providers around the world have been jumping into metro Ethernet with both feet. Asia continues a massive buildout of backbone, core, and edge infrastructures. In November 2004, Videsh Sanchar Nigam Limited (VSNL), India’s leading telecommunications and Internet service provider, announced the deployment of India’s largest broadband metro Ethernet solution for Tata Indicom Broadband Services. Based on Cisco gear, the solution will provide VSNL’s enterprise and residential customers with high-quality broadband services of 10/100-Mbit/s connectivity.

Europe is another growth spot for metro Ethernet. In the past three years, a slew of service providers, municipalities, as well as governments across continental Europe and the UK, have been

building extensive fiber infrastructures and rolling out metro Ethernet services. Just some of the service providers are FastWeb in Italy, Lyse in Norway, and InTechnology in the UK.

This article looks at three service providers—UNET in The Netherlands, Hong Kong Broadband Network Ltd. in Asia, and Time Warner Cable in the US—who are benefiting from metro Ethernet deployments.

Almere: Fastest-Growing City in Europe Lays Fiber

In 2002, business and municipal leaders in Almere, a city of 175,000 in The Netherlands built on land reclaimed from the sea, decided that attracting new business required a foundation of widely available “real broadband”—with speeds of 100 Mbit/s and up to every home and business. Dutch service provider UNET was chosen to help design, build, and monitor this new metropolitan network.

“Real broadband is a way of attracting companies who want to do business, work, and live in a modern city,” says Mayor Annemarie Jorritsma. Working with Cisco’s Internet Business Solutions Group, Almere has thus far installed fiber to 1700 homes and 500 businesses as a pilot. About 15,000 homes and businesses should be fibered by the end of 2005. The network uses Cisco Catalyst 4500 Series switches in the core, Catalyst 4500 and 3500 Series switches in the access network, and Cisco gateways on customers’ premises.

The city owns the network, which operates as a virtual LAN (VLAN); broadband services and network management are handled, respectively, by UNET and its subsidiary, First Mile Ventures. Development of services is also open to other providers. The first houses came on-stream in March 2004; by mid-year residents received triple play service—symmetric 10-Mbit/s Internet access, digital radio with 20 channels and television with 75 channels, and IP telephony through Cisco CallManagers. Businesses receive a symmetric 100-Mbit/s or 1-Gbit/s link.

“We had a 10 percent sign-up rate before we were even ready to deliver service,” says Ger Bakker, chief technology officer at UNET. “We’re now at about 20 percent, and our business plan calls for 40 percent three years from the start.” Prices are very competitive, especially for serious Internet users, he adds: 80 Euros per month (about US\$104) for the lowest triple play service levels.

Besides basic high-bandwidth services, UNET is beginning to roll out additional services—video on demand (VOD) is one of the first, as is remote backup for both homes and businesses. “The service automatically backs up the user’s computer daily, so the customer doesn’t need to worry about hard disk or CD backups,” Bakker says. Others to come are storage on demand and security; security companies can locate IP cameras on their customers’ premises for monitoring

after an alarm or around the clock, and the alarm itself would go instantly over IP, rather than a dialup connection. Healthcare monitoring is also a possibility. “Cameras in the home—turned on only on demand, of course—would enable healthcare workers to check on the elderly and the elderly to remain longer in their homes,” Bakker points out. Schools could also provide home access to classes or distance learning.

UNET is now selecting a network operator for all real broadband service in Amsterdam, which should begin later this year, and is planning to offer service in the next four largest Dutch cities. “By 2010,” Bakker says, “we expect that 90 percent of homes and businesses in these cities will be served by fiber.” He also expects many homes to be using not just 100-Mbit/s service but rather 1 Gbit/s. “Five digital TV sets in a home, and you already need more than 100 Mbit/s, and don’t forget the upcoming HDTV over IP.”

Hong Kong Homes Will Enjoy 1 Gbit/s

Hong Kong Broadband Network Ltd. (HKBN) announced last November it would offer customers on its all-residential network symmetric 100-Mbit/s then and 1000-Mbit/s (1-Gbit/s) service by the second calendar quarter of this year—it believes the latter a worldwide first. The company already passes 1.2 million homes—60 percent of all those in the city—and has an aggregate base for voice, broadband, and pay TV of more than 500,000 subscribers.

Competition in Hong Kong is tougher than in most cities, because the building density means that most homes are within three kilometers of their local telephone exchange so ADSL can easily deliver 6 to 8 Mbit/s. But, says HKBN’s chairman Ricky Wong, “We knew that pay TV would be vital for our success, and the sooner we got to market the better. We were faced with a difficult decision: We had to either make do with leased circuits or upgrade our network and buy an optical core.” Further impetus, according to Wong, came from the desire to offer additional broadband services in the future.

Over the past four years, HKBN, a unit of City Telecom (HK) Ltd., has lowered its initial reliance on wireless local multipoint distribution service connections by acquiring or installing 100,000 core km of fiber, covering most of Hong Kong. It expects its own fiber network will pass 90 percent of homes by the end of June.

HKBN’s backbone consists of 2x200 optical core fibers that can support up to 64,000-Gbit/s transmissions. Developed with Cisco’s help, it uses the Cisco ONS 15454 Multiservice Transport Platform, which employs dense wavelength-division multiplexing (DWDM), Cisco Catalyst 6500 Series switches, Cisco RPR architecture in the core, and Cisco Catalyst 4507R and 3550 Series switches to distribute service to apartment buildings—more than 2500 so far.

The network includes more than 10,000 Cisco switches and more than 800 Cisco routers. HKBN owns the building gateways and Cisco Catalyst 2950 Series switches. This system enables upgrades from 10 to 100 Mbit/s to be provisioned remotely. Then, Category 5e copper wiring already in the buildings takes signals to the individual units.

“Cisco’s ONS DWDM platform is the only one in the market that can offer HKBN all the optical legacy services that will support our future growth, including Layer 2 and Layer 3 services,” says Sam Leung, technical director at HKBN. In addition, the Cisco Catalyst switches offer wire speed, multilayer switching with granular QoS, advanced security, and predictable performance.

The cost, HKBN reports, was around US\$130 per home pass, a fraction of the cost for networks of comparable bandwidth elsewhere. It charges US\$35 per month for unlimited local 100-Mbit/s service and 20-Mbit/s international access. And come HKBN’s BB1000 service, 1 Gbit/s will be carried over Category 5e copper. HKBN dedicates four pairs to each customer, two for IP access and two for video. With the separate ports, explains Lim Wong, consulting services engineer in the Asia Pacific Consulting Group at Cisco, the company has even better control over QoS and to whom the TV signal is delivered.

Cable Company Delivers Real Broadband

Time Warner Cable, through its offering called Road Runner Business Class, provides services within 31 regional divisions located in major urban and suburban areas in the US. In the last few years, Time Warner Cable has invested heavily in fiber networks to connect these locations—and now has a well-groomed infrastructure for providing metro Ethernet. It simply extends fiber from its existing network to an enterprise’s single or multiple locations, linking them together via metro Ethernet in urban areas and via the rings across regions. The service delivers from 5 Mbit/s up to 1 Gbit/s.

“Customers are asking for an alternative to incumbent carriers via a completely separate network, which we have, as well as reliability, SLAs, flexibility in provisioning, and competitive pricing,” says Kurt Fennell, vice president of technology, operations, and field integration in Time Warner Cable’s Commercial Services Organization. “Our fiber network is very robust, and we can compete with any local or competitive carrier on reliability.”

Over its fiber to the enterprise, Road Runner Business Class provides basic high-bandwidth Internet access, and private line and point-to-point or multipoint services. “We offer three basic SLAs, with active monitoring and reporting on intelligent CPE and core

network. And we can provide just about any bandwidth the customer wants.”

Road Runner Business Class’ dedication to customer service is evident in its corporate structure. For example, each Time Warner Cable division is staffed with sales consultants, engineers, and technical support agents as well as administrative personnel to serve customers in the market. The “distributed architecture” makes the company very nimble in responding to customer needs. “Our sales consultants will work with the enterprise organization to develop the right solution to fit into their complex environment,” says Fennell.

The Road Runner Business Class commercial network is built with Cisco ONS optical products, Cisco Catalyst 3550 Series switches as CPE and core access devices, and Catalyst 6500 Series switches or Cisco 10720 routers to manage the urban transport networks.

“Larger pipes enable our customers to better deliver services in a more efficient and customer-effective manner. Metro Ethernet opens paths for the delivery of just about any service,” Fennell says. “Bigger pipes enable customers to do their business better, and our costs are lower. We have a tremendous advantage in being positioned to serve customers over a converged network. I can only see customer demand for these services growing.”

Metro Ethernet—a Future-Proof Technology

Metro Ethernet’s advantages have been often enumerated: high speed, ease of delivering any amount of bandwidth, QoS, customer desire, and many others. Bakker of UNET, an experienced provider, cites a crucial advantage: “If you’re going to go to that level of service and broad spectrum, you want to make your network as simple as possible. A basic engineering principle is that complex is bad. . . . IP Ethernet is fast, simple, and future-proof.” ■

FURTHER READING

- Cisco Metro Ethernet Solution
cisco.com/go/metroe
- Metro Ethernet for Service Providers
cisco.com/packet/171_8a1
- “The Service-Driven Network”
cisco.com/packet/171_8a2
- Cisco Metro Ethernet Products
cisco.com/packet/171_8a3
- Metro Ethernet Network Blueprint
cisco.com/packet/171_8a4

The IP NGN Journey

Cisco innovation and technology advancements are helping service providers on the journey toward IP-based next-generation networks.

Gone are the days when providing connectivity was the name of the game. Today, service providers of all stripes must look toward offering new, value-added services for revenue growth, greater competitive differentiation, and increased customer loyalty. Carriers have adopted a strict laser focus on achieving efficiencies in operating expenses (OpEx) and capital expenditures (CapEx) to boost profitability. And in this intensely competitive environment, it's increasingly important for providers to gain control of their networks and the services that run on them and, in the process, to regain greater control over their business from the ever-changing market.

Service providers also need flexible solutions that help them cost effectively address the unique requirements and tap the opportunities of their various customer segments—consumers, small and midsize businesses, large enterprises, and wholesale customers. For example, in the consumer space, gaming, network-based personal video recorders, video on demand (VoD), Wi-Fi networks, and mobility are growth areas. Small and midsize businesses are likely to increase their interest in and use of a range of managed services such as hosting and security. Meanwhile, enterprises will experience increased demand for Layer 2 and Layer 3 virtual private networks (VPNs), remote access, storage, security, and Ethernet. For their part, carriers will seek revenue from wholesaling access, local and long-distance voice and services including collocation, peering, transport, and content delivery.

To address these diverse markets, service providers need a single infrastructure capable of evolving to provide a wide range of new services that will increase revenues and customer loyalty, as well as yield efficiencies in OpEx and CapEx. The industry generally calls this forward-looking infrastructure a *next-generation network (NGN)* and has near-unanimous consensus that IP will be the foundation technology to make it a reality.

“Many in the industry have narrowly defined the term NGN to address only a small piece of the very significant transition required by service providers,” says Jeff Spagnola, Cisco's vice president of service provider marketing. “Cisco takes a more comprehensive view

of an IP-based NGN that addresses a wide range of issues that service providers must resolve. We believe that IP NGNs bring about a broad network transformation that encompasses not just the service provider's network but its *entire business*.”

Nor does this network transformation end at a single point. Like providers' business and service plans, the IP NGN is a continuum. It will constantly evolve to adapt to customer demand and new technology opportunities. “IP NGNs refer to the idea of one network that can not only cost effectively deliver and manage all the voice, video, and data communications options available today, but one that can also adapt and grow to handle any new communications options that will inevitably evolve,” says Mike Volpi, senior vice president of Cisco's Routing Technology Group.

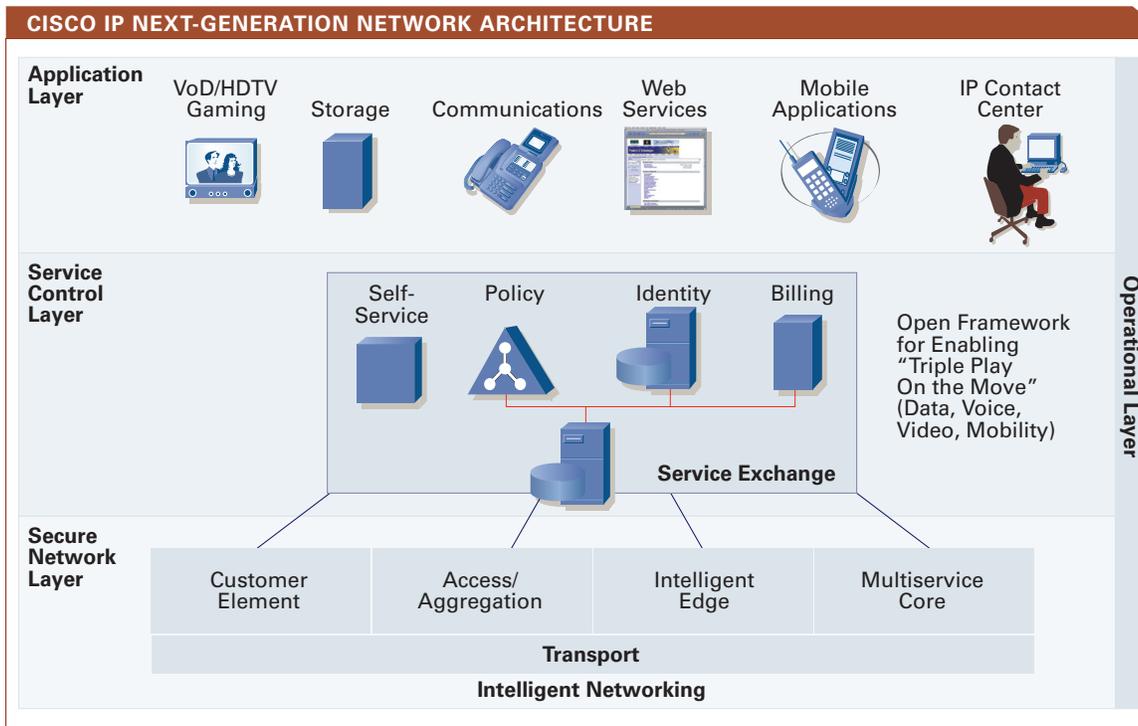
Many service providers are already moving toward IP NGNs. Though they might use different terms for NGN, broadly speaking, they share many of the same basic concepts in their visions for tomorrow's carrier infrastructure. AT&T, for example, is pursuing an NGN through its “Concept of One, Concept of Zero” initiative, and British Telecom characterizes NGN as the “21st Century Network.” Individual service providers will migrate to an IP NGN at their own pace based on their business and regulatory requirements.

The phased development of the IP NGN, emphasizes Volpi, involves creating an *intelligent* infrastructure from which application-aware services are delivered by service-aware networks. This type of intelligent IP NGN will open new opportunities for service providers to offer end customers advanced, value-added, and personalized all-media services securely and seamlessly over wireline and wireless connections.

Convergence Is at the Heart of IP NGN

Central to an IP NGN are three fundamental areas of convergence already being enabled by service providers today:

- *Application convergence*—integrating new, innovative IP data, voice, and video services over a single broadband infrastructure.
- *Service convergence*—Providers are migrating toward delivering “triple play on the move,” which combines voice, video, data, and mobility services. Service convergence includes network access and control that is technology-agnostic and seamlessly



FOUNDATION FOR SUCCESS The goal of the Cisco IP NGN architecture is to provide rich, personalized, value-add multimedia services. To do this, carriers need a service control framework that supports the key business transition from a basic "highway" to value-added "tollway" service structure.

compatible with any networking medium: mobile, wireless, cable, DSL, or Ethernet.

- **Network convergence**—Providers are migrating from deploying, managing, and maintaining multiple service-specific networks to delivering all services on a single network, most often an IP Multiprotocol Label Switching (IP MPLS)-based network.

Of course, service providers prioritize these areas of convergence in different ways, depending on their business. Many mobile operators, for example, might focus most of their efforts on service convergence, whereas cable operators target their efforts at application convergence.

The Cisco IP NGN vision and architecture address these three primary areas of convergence (see figure).

"Providers worldwide are building networks to create revenues, not just to move bits," says Tom Nolle, president of CIMI Corporation, an industry analysis and consulting firm. "Cisco's IP NGN architecture and vision offer them a compelling model for generating revenue from new services that focus on delivering a network experience based not just on transport and connection but on linking applications and networks in a seamless way to achieve carrier goals."

Recent advancements by Cisco, largely in the areas of service control and the secure network layer, underscore its commitment to building, acquiring,

and partnering to develop technology and solutions that help service providers transform their networks to profitable IP NGNs.

Cisco IP NGN: Service Control Layer

To achieve true service convergence, companies must be able to operate, bill, and manage services over a range of access media. To this end, Cisco and its technology partners have developed and are continuing to advance an open Service Exchange Framework, which allows providers to facilitate and control customer access and use wireline and mobile IP services with no limits on the types of applications that can be deployed.

While this framework contains a range of different products and solutions from Cisco and its partners, one of the most recent additions comes from Cisco's acquisition of P-Cube, a developer of IP service control platforms. The *Cisco Service Control* solution overlays intelligence and application-level control on existing IP transport networks, allowing service providers to analyze, control, and meter and charge for multiple application- and content-based services—all on a common network infrastructure. The hardware components of the solution, the Cisco SCE 1000 and 2000 Series Service Control Engines, are programmable network elements that reside behind an aggregation device such as a Cisco 10000 Series Router, broadband remote access server (B-RAS), or cable modem termination system (CMTS). The Cisco SCE interoperates with subscriber authentication and management components, as well as billing, data

collection, and policy provisioning systems, to deliver transparent, application-differentiated broadband services to subscribers.

Running on the service control engines, the Cisco Service Control Application Suite is composed of three software applications: Service Control Application for subscriber service monitoring, Cisco Collection Manager for capturing and reporting service data, and Cisco Subscriber Manager for individualized traffic accounting and control.

The Service Exchange Framework is further enhanced by Cisco's recent acquisition of dynamicsoft, a maker of carrier voice-over-IP (VoIP) software based on Session Initiation Protocol (SIP). The integration of dynamicsoft's technology with Cisco's carrier VoIP products, such as the Cisco BTS 10200 Softswitch, will help service providers offer SIP-based integrated communications services (telephone, mobile phone, e-mail, and instant messaging) that enable users to be contacted via a single device.

These new Service Exchange Framework components complement the Cisco Mobile Exchange (CMX) portfolio, which addresses the interface between the evolving radio access network and an array of Internet services offered by IP networks. CMX gives mobile operators, application providers, and system integrators flexible solutions that enable them to offer value-added data services to mobile subscribers.

Cisco IP NGN: Secure Network Layer

At the foundation of an IP NGN is the secure network layer, composed of a customer element, access/aggregation, intelligent IP MPLS edge, and multiservice core components with transport and interconnect elements layered below and above. The secure network layer is undergoing fundamental change compared to just a few years ago. For example, IP MPLS is being integrated throughout each section of the network, and edge and core areas are converging, with each adopting capabilities of the other and providing greater efficiencies to service providers.

Cisco has played a major role in developing IP MPLS communications infrastructures, the foundation for large-scale, converged, next-generation IP networks. "For several years, IP MPLS has been recognized as a foundation enabler of network convergence," says Spagnola. "Cisco has more than 250 service provider customers worldwide who deploy IP MPLS. By virtue of the fact that these customers have chosen Cisco and its intelligent IP MPLS solutions, they are already joining us on our mutual IP NGN journey."

Cisco is leading the industry in delivering innovative technology to drive network convergence and help service providers lower infrastructure costs. This is

most evident with the Cisco CRS-1 Carrier Routing System and the recently launched CRS-1 8-Slot Single-Shelf System. The world's most advanced routing system, the CRS-1 has a system capacity of up to 92 Terabits per second (Tbit/s) and is designed to provide continuous system operation, service flexibility, and extended system longevity to telecommunications providers and research organizations. Designed to fit into half of a standard 19-inch rack and with 640 Gbit/s of total switching capacity, the Cisco CRS-1 8-Slot System extends the reach of CRS-1, providing a foundation for network and service convergence.

Global carriers and research organizations worldwide are adopting the Cisco CRS-1 for building out their IP network infrastructures and to deliver advanced multimedia services. A few examples are:

- Broadband content and services provider SOFTBANK BB Corp. in Japan (provider of Yahoo!BB) has chosen the Cisco CRS-1 for its IP NGN super backbone core router. SOFTBANK focuses on services such as broadband Internet access, video-on-demand, and online gaming.
- SuperSINET, the largest national academic research network in Japan, will deploy the Cisco CRS-1 as the core routing system to enable research of grid, supercomputing, and other scientific applications.
- The Pittsburgh Supercomputing Center, a leading scientific research organization, has been measuring IP NGN performance using the Cisco CRS-1, to gauge performance levels required for advancing next-generation scientific research.
- Telecom Italia is in trials with the CRS-1, which serves as the network foundation for delivering advanced multimedia applications to its customers. So far, the CRS-1 is meeting the carrier's top requirements for availability and service flexibility and is slated to be a key component of Telecom Italia's IP NGN.

◆ ◆ ◆

Cisco's strategy in the service provider arena is to innovate and to provide the technology, solutions, and expertise carriers need as they transform their networks and move along the IP NGN journey. Deploying solutions that deliver greater network intelligence, integration, and overall flexibility will not only provide carriers with short-term relief but, in the end, enable them to combat competitive pressures, address new market opportunities, and increase profitability. ■

FURTHER READING

- Routing Solutions for Service Providers
cisco.com/packet/171_8b2
- Cisco Service Control
cisco.com/packet/171_8b3
- Cisco CRS-1 Carrier Routing System
cisco.com/packet/171_8b4

Read the full white paper, *Cisco and the Service Provider IP Next-Generation Network Journey*, at cisco.com/packet/171_8b1.

Spinning IOS into Gold

How to Turn IOS Technologies into Profitable Services

The intelligent functionality in Cisco IOS Software has long led technology innovation throughout the networking industry. Service providers already take advantage of many IOS capabilities to build and manage their network backbones more efficiently. Cisco has targeted an integral set of IOS features that service providers can use to add revenue-generating security, management reporting, and route-management service offerings to their portfolios. With this functionality, called *IOS Technologies for Managed Services*, Cisco is helping service providers identify specific IOS technologies that they can productize and market to business customers as profitable, add-on complements to their managed IP VPN services.

The following IOS technologies can each be turned into a commercial service: Cisco IOS Firewall; Cisco Intrusion Prevention System (IPS); Cisco IOS IP Service-Level Agreements (IP SLAs); and Cisco Enhanced Interior Gateway Routing Protocol (EIGRP), when supported on the provider edge (PE) router.

Selling Edge Security Services

Service providers can easily sell managed security services at the perimeter of their enterprise customers' WANs as part of a managed service using the firewall and intrusion prevention features in Cisco IOS Software. "Managed WAN routers are the business customer's first line of defense," observes Lily Lu, marketing manager in Cisco's Products and Technology Marketing Organization. "The service provider that offers managed VPN services would do well to become the enterprise's first-tier security partner with firewall and intrusion prevention services at the WAN edge."

To do so, the provider can simply activate the Cisco IOS Firewall security-specific option in the customer edge (CE) routers that they install and manage on behalf of business customers. The high-performance Cisco IOS Firewall is embedded in Cisco IOS Software on a broad range of router platforms, from the very low to very high end. It statefully filters TCP, UDP, and Internet Control Message Protocol (ICMP) traffic in accordance with the business customer's policy—blocking some traffic while permitting access to other traffic—based on source IP address (including IPv6), protocol type, MAC address, user ID, password authentication, and other criteria.

For business customers running voice over IP (VoIP) across their IP VPN services, the Cisco IOS firewall software supports voice traversal using deep-packet inspection. The firewall recognizes the application

protocol and can follow the call on a per-flow basis as it "hops" among ports and new channels are opened. The firewall currently recognizes and supports H.323v2, Session Initiation Protocol (SIP), and Cisco's own Skinny Call Control Protocol VoIP signaling protocols.

When selling managed router services based on Cisco's new line of 1800, 2800, and 3800 Series Integrated Services Routers for branch offices, service providers can layer on dynamic intrusion prevention. The routers can dynamically load the 740 signatures supported by the Cisco IDS Sensor appliance platforms in real time. Providers who have control of the customer's WAN access router can use the *Cisco IOS Intrusion Prevention System (IOS IPS)* feature to modify an existing signature or create a new signature to address newly discovered threats.

Differentiated Services and Management Reports

Service providers, of course, require ways to measure their customer network service levels just to make sure their offerings are competitive and of high quality. In addition, sophisticated and granular performance monitoring and measurement tools can be put to work to offer money-making differentiated service classes for supporting application traffic with differing performance requirements and to offer the SLA performance information and reports that business customers are beginning to demand. These metrics can more or less be difficult for the service provider to measure on a customer-by-customer basis. For example, one customer contract might require that a set of metrics be met on a monthly basis across all sites on its network (fairly easily enforced), while another might delineate specific metrics on a per-site, per-day basis (much more difficult). The more granular and specific the metrics are, the more meaningful they are to an individual enterprise site/user's experience, but the more difficult they are for the service provider to deliver.

"IP and SLAs are converging and IP performance monitoring needs to be application-aware," says Tom Zingale, Cisco IOS product manager for IP SLAs and NetFlow. This transformation is critical for new IP network applications such as VoIP, audio and video, enterprise resource planning (ERP), customer relationship management (CRM), and material requirements planning (MRP). SLA measurement needs to be end to end for today's VPN services, adds Zingale.

Service providers can charge more as agreed-upon SLAs become more stringent. With *Cisco IOS IP*

For more on Cisco IOS Technologies for Managed Services, visit cisco.com/packet_171_8c1

SLAs, providers can deliver performance management reports and self-monitoring capabilities—required add-ons to managed services for which providers can charge. Cisco IOS IP SLAs use unique service-level assurance metrics and methodology to provide highly accurate, precise service-level assurance measurements and are embedded in the network for flexibility and cost-effectiveness.

Cisco IOS IP SLAs enable the service provider to measure delay, jitter, packet loss and VoIP quality in real time on a hop-by-hop basis between any two routers under its control, including CE to CE, CE to PE, PE to PE, as well as from point of presence (POP) to POP. When providing a managed IP VPN service, for example, the provider can measure these metrics between any two customer routers and generate billable reports for that customer on the various metrics.

“Service providers are competitively driven to offer SLAs, and the enterprise needs to verify that the SLA is indeed being provided,” says Zingale. “The IP SLAs functionality is embedded in the network and provides information for end-to-end SLA reporting and verification. If both the service provider and its customer have the technology, they can monitor their QoS [quality of service] so the service provider is sure it is providing ‘gold’ service to the customer, and the customer can validate it.”

The most recent additions to the IP SLAs feature have related to VoIP metrics. With embedded simulated codec capabilities, IP SLAs allows service providers to easily create test calls, for which they can continually measure performance and also provide a mean opinion score (MOS) for measuring user perceptions of voice quality. In February 2005, the IP SLAs functionality was enhanced to measure *post-dial delay*—the amount of time a call rings, rings busy, or takes to connect using SIP or H.323 protocols; and *gatekeeper delay*, the time it takes for a device to register a number in the IP address-to-phone number database

Global service provider Equant has implemented IP SLAs with its MPLS-based IP VPN customer DuPont in South America, though the scientific-products maker expects to go global with the implementation, which includes VoIP, in 2005. Equant supplies DuPont with end-to-end reporting for multiple classes of service on delay, packet loss, and jitter. “When we contemplated Layer 3 VPN service from Equant, we realized we needed a way to measure SLAs,” says Mike Dowler, global services integration manager at DuPont Telecommunications. “The ability to create and verify different classes of services with different priorities [including VoIP] would be compromised without a reliable means of measurement.”

Easing the Enterprise’s Routing Burden

Providers of MPLS-based IP VPN services can leverage Cisco routing software to offer a special route

distribution and management service that eliminates the enterprise customer’s requirement to learn and implement the Border Gateway Protocol (BGP) in its network between the PE and CE. In effect, the provider manages the customer’s internal EIGRP routes across the VPN WAN in addition to (or in lieu of) managing the customer’s CE router. To do so, the service provider supports EIGRP on the PE to which the customer CE attaches. More than 60 percent of enterprises currently run EIGRP as their interior gateway protocol of choice because of its fast convergence, ease of configuration, and network efficiency benefits. PE support of EIGRP prevents enterprises from having to convert their EIGRP networks to BGP to utilize BGP-based MPLS VPN services, says Jim Crockett, a systems engineering manager at Cisco whose SE team helped Verizon Communications roll out EIGRP-based MPLS VPN services in April 2004.

Native EIGRP instead becomes the routing protocol running between the CE and PE. Without EIGRP PE-CE support, normal redistribution of EIGRP into BGP at the PE would result in inter-site EIGRP routes appearing as external routes in the target customer network, Crockett explains. “If there is a ‘back-door’ EIGRP route between sites that doesn’t use the MPLS VPN—and there often is, such as an ISDN backup link or a connection from a merger or acquisition—traffic will always take the EIGRP route instead of using the MPLS VPN unless the enterprise converts its entire network to BGP.”

By supporting EIGRP on the PE, service providers preserve the enterprise customer’s EIGRP metrics across the MPLS VPN backbone using Multiprotocol-BGP (MP-BGP) extended community attributes. The enterprise’s internal EIGRP routes are redistributed into BGP with extended community information that is appended to the provider’s BGP route. BGP then carries this route over the provider’s MPLS VPN backbone, with the EIGRP route information appearing as any other MPLS label-encapsulated data. Once the peering site receives the route, BGP redistributes the route into EIGRP, which extracts the BGP extended community information and reconstructs the original route.

Another technology to run from the managed CE across WAN and ISP links is *Cisco IOS Optimized Edge Routing*, when there are two or more physical or logical paths across the WAN. The best path is selected based on latency, packet loss, reachability, throughput, link balancing, and/or monetary cost reduction.

The ability to offer these route distribution and management outsourcing services leverages the fact that Cisco IOS technologies run from low- to high-end platforms, enabling service providers to take over management of the primary routing protocol in use in the majority of enterprise environments today as they traverse the WAN backbone. ■

Passing the Test

Early Field Trials participants discuss the challenges and benefits of working with Cisco to evaluate pre-release products.



Noel Hedrickson

STAFF THE BATTLESTATIONS Gordon C. Hawkins, network and systems engineer for Vancouver Film Studios, tests Cisco products in a unique environment.

By Fred Sandmark

We all know people who can be considered “early adopters.” They’re those individuals who have a burning desire to use new technologies before most other people do, and who are eager to share their knowledge and enthusiasm with others. (Many *Packet* magazine readers can probably get a good look at an early adopter by looking in the mirror.)

The man who coined the phrase *early adopter* in 1962—a sociologist named Everett Rogers—also had a term for the 2.5 percent of the population that uses technologies even before the early adopters do. He called them *innovators*, and described them as “venturesome,” that is to say, risk-takers.

When Cisco Systems develops a new product, it wants to thoroughly test it in real-world situations before officially releasing it. In a sense, the company looks for that small segment of venturesome innovators who are willing to take risks with brand-new technology in order to stay ahead of the pack. To do this, Cisco

has developed Early Field Trials (EFT), a program through which it selects customers and partners to test equipment before it becomes commercially available.

An EFT is the very first test of new hardware or software products in a customer setting. EFTs follow alpha testing, which occurs internally at Cisco, and precede beta testing, which occurs in a larger number of customer sites than EFTs. (Not all Cisco business units run both EFTs and beta testing.)

EFTs are a vital part of a product’s path to technical and commercial success. “Cisco has a reputation that, when a product comes to market, it has been well-developed,” says David Hope, director of sales and marketing for DSi, an EFT participant and network infrastructure service provider. “And I think the testing is an important step in that.”

SOPHISTICATED, BUT NOT NECESSARILY LARGE

DSi performed three months of EFT testing for the Cisco Catalyst 4500 Series Supervisor Engine II-Plus-TS in what Senior Network Engineer Mike Cotrone describes as a “semi-production” environment in DSi’s laboratory.

“We didn’t switch over anything to it in terms of full production, but I had traffic flowing through it and stress-tested it with packet generation,” he explains. “I used it for a lot of IP telephony demonstrations and studies for the months I had it installed.”

This made sense for DSi, because the company is not only a customer, but is a Cisco partner serving small and medium-sized businesses (SMBs). Forty of its 60 employees are in technical positions, and 22 of its engineers carry Cisco Career Certifications.

“I just saw [the new product] fitting into a lot of our accounts, in terms of small and medium businesses,” says Hope. “We wanted to test it, get our hands on it, and make sure we understood its capabilities.”

The EFT program isn’t limited to companies like DSi that have deep engineering benches, but participation does require a high level of technical expertise. Cisco engineers interview prospective testers and review their technological skills and network environments. The company also requires EFT testers to sign a nondisclosure agreement (NDA), agree to invest time and energy in the test process, carefully document their tests and results, and communicate regularly with the product’s development team. And, yes, Cisco expects to get the equipment back at the end of the test period.

In spite of these rigorous requirements, the EFT program attracts a wide variety of companies. Take, for example, Watt Commercial Properties, a 150-person real estate firm with ten offices and an IT staff of five.

“Our network environment is not a Coca-Cola or a General Motors,” says Dan Campbell, the company’s chief information officer. “Ours is a fairly sophisticated environment, but not a very large one.”

Campbell and his staff performed EFT testing on the Cisco Catalyst 4500 Series Supervisor Engine II-Plus-TS and Cisco Network Assistant software. The new products proved a good fit for his company.

“It’s a small and very new IT organization here, so we were looking to derive some efficiency,” he says. “We ran the products for a little while in a parallel environment.” Campbell and his team used Cisco Network Assistant to replicate standard router configurations and do remote setup—an important real-world test,

because the growing company opened three branch offices during the test period.

“Being part of the program gave us a glimpse of what these products could do without actually having to purchase them.”

—Gordon C. Hawkins, Vancouver Film Studios

APPLYING EMERGING TECHNOLOGIES

Watt agreed to participate in the EFT program because it was in the market for the very product it was asked to test.

“We had already identified a need to have a product like that, and the test seemed like a good opportunity to give our comments on what we wanted to see,” Campbell says.

And, for Watt, participation had a big payoff: “There were a couple of pieces of equipment that we used that weren’t supported [in the EFT release] that Cisco moved up the [priority] list because of our input,” Campbell says.

Participating in the EFT program also allowed Watt to directly communicate with Cisco—something that might not otherwise happen in a company of its size.

“For want of a better term, we got a little bit of free consulting as a part of the implementation,” Campbell says. “We could ask Cisco questions at a level that we probably wouldn’t be able to access by just buying a product from a third party.”

Seeing a new product in exchange for an investment of time is an incentive for some EFT participants.

“Being part of the program gave us a glimpse of what these products could do without actually having to purchase them,” says Gordon C. Hawkins, network and systems engineer of Vancouver Film Studios (VFS) in Vancouver, British Columbia, Canada. “Definitely that was an advantage to us, because we wanted to test-drive the stuff and have an extended trial. Being a small company it’s sometimes difficult to get more IT dollars.”

Hawkins participated in EFT on the new Supervisor II-Plus-TS for the Catalyst 4503 in an IP communications environment. VFS has 17 buildings and 30 employees (including an IT staff of three), and is part of the larger parent company, The McLean Group, a real-estate development company in British Columbia.

But those numbers don't tell the whole story: At any time, hundreds of film and television professionals might be working on VFS's site, often for months at a time. These people need telephone and data services, so the company's network environment is constantly changing to accommodate their needs.

"The test period was a month and a half, maybe two months," Hawkins says. "We tested a bunch of IP communications gear and Cisco 7900 Series IP Phones, and we connected it to our network. We played

"We had already identified a need to have a product like this, and the test seemed like a good opportunity to give our comments on what we wanted to see. There were a couple of pieces of equipment that we used that weren't supported that Cisco moved up the list because of our input."

—Dan Campbell, CIO, Watt Commercial Properties

around with the configuration, looked at the documentation, and got a feel for some of the capabilities of the supervisor blade. After that we moved to a satellite building and configured some of the virtual LANs and the Quality of Service features. We created a routed interface where all the IP phones in this particular building would connect into that router for the IP and routing functions. However, when they needed to access an external network, like for IP telephony, it went across a routed boundary and handed it over to the Cisco Catalyst 6500 in our core network . . . and it worked fantastic."

In his high-pressure environment, Hawkins says that participation in EFT testing provides an important element of professional growth for himself and his staff.

"Work is so hectic and fast-paced these days," he says. "You really need to spend some time in the lab, and value that time, to get your IOS knowledge up to par and make sure you're aware of emerging technologies. It's all about taking those emerging technologies and applying them to our business to make it better.

You'll never get your head around that unless you have time in the lab."

COSTS AND BENEFITS

Indeed, for some IT shops, the opportunity to take time out from everyday tasks to experiment with pre-release equipment is considered a perk. "It gave my staff the ability to push the envelope and do something different from their day-to-day operations," Campbell says of the EFT testing. "They enjoyed it.

That's not to say that participating in the EFT program is easy. It requires a time commitment—and there are costs, albeit soft costs.

"I had to come up with a structured test plan and procedure," explains Cotrone. "I had to pull myself out of the field, away from my work and [my role as] a team lead. And there was a very large documentation [requirement] that took quite a bit of time on the back end. It's just a large time investment, which equals a certain revenue loss."

Still, many EFT participants say the benefits outweigh any costs. They also like the fact that the EFT program opens up communication channels to people within Cisco.

"We're talking to account managers, project leads, and development engineers," Hawkins says. "Now, they're experienced in with working with us, so the next time that we call them or open a support case, they'll know who we are. It has an impact not just on the Catalyst platform for which we did the EFT testing, but any kind of Cisco technology."

Campbell, Hawkins, and Hope are all enthusiastic about their EFT experiences, and all three would participate in the program again.

"I think Cisco is developing an attractive product set for the SMB market, and we believe that's important for us," Hope says. "We hope to continue to play a role where Cisco is developing products in this space." ■

Cisco business units recruit participants in Early Field Trials in different ways, depending upon the type of product that needs testing.

Customers interested in learning more about the EFT program should contact their account managers or channel partners to get more information. Go to cisco.com/packet_171_9a1

FURTHER READING

- New York University EFT Experience cisco.com/packet_171_9a2
- Cisco IT's EFT of Wireless LAN Services Module cisco.com/packet_171_9a3
- Bang Networks EFT Experience cisco.com/packet_171_9a4

Who's at Your Service?

Setting up support contracts for your technology requires flexibility on both sides of the negotiating table.

By Howard Baldwin

Knowing what's on the network and its relative priority within the scope of the business is just one element of keeping your operations up and getting the optimal value from your technology investments. Setting up a support contract is a simple but important way to offload that worry to a third-party provider. You want to worry about deliveries, not downtime.

"SMBs have come to rely on their IT infrastructures not only to reach business goals but also to differentiate themselves in the marketplace," says ML Krakauer, a vice president in HP's services division focusing on SMBs. "When there's a problem, they face the same issues as enterprises in terms of potential loss of revenue, lower customer satisfaction, and negative publicity, but they often don't have the same level of IT support resources."

As a result, it's vital to make sure your company has the right support package. You need to understand what level of support you require, what program provides that support, who will deliver the support, how you will deal with multiple providers, and the parameters of your contract. It's important to manage everyone's expectations, including those of your provider.

Vendors understand the increasing importance of the SMB market. "We see it as one of the fastest-growing customer segments," she says. "Up until now SMBs haven't had solutions that are designed to address their needs." Cisco is launching a program targeting SMBs called Cisco SMB Support Assistant. You can also look to your vendors' channel partners, value-added resellers, or even third-party independent support firms. The costs vary depending on the level of service and the complexity of your system, but contracts may run for as little as US\$40 per year to more than US\$1,000.

Negotiating the Contract

No matter who will handle your support, the contract helps define your relationship with your provider. If you only start looking at a support contract when you need it, you're already too late. Strive for a contract that both gives you protection and the ability to amend it if your needs change.

"The costs have to be predictable, and the contract has to have flexibility and clarity," warns Michael Lauricella, vice president of telecommunications



The First Quarter issue of *iQ* features articles about the winners of the annual Cisco Growing with Technology Awards.

You'll also find articles to help you with wireless networking, technology adoption in Europe, network security and selecting technology vendors.

Find the articles online at cisco.com/go/iq. Subscribe today to get *iQ Magazine*, a free quarterly publication from Cisco for small and midsize businesses. cisco.com/go/iqmagazine/subscribe/packet.

research at AMI-Partners market-research firm.

That's why it's important to know exactly how important something is to your business. You may want 24x7 protection with a two-hour response time for your e-commerce server, because if it's down, you don't make sales. The priority is the same for your mail server if you rely on e-mail communications. On the other hand, a Web server that contains only product brochures might not require the same urgent response.

But you also need flexibility in the contract, because you may be paying a single fee for a high level of protection and then decide you can safely cut back. "You may want to move to a 'pay as you go' system in the second year of a contract," suggests Lauricella, "with a cap on the amount you pay."

SMBs have an advantage in negotiating the terms of service contracts: There's a lot of competition for this huge customer base, so vendors will want to offer flexibility.

"No two channel partners offer the same service contract," notes Helen Chan, manager for SMB strategies at AMI-Partners. "Every one involves customization." Think about how your business may evolve over the life of the contract.

"SMBs are entrepreneurs—they can smell a bad deal a mile away," says Lauricella. "But they also have to realize that there are limitations as to how far a support provider can go in terms of what it offers. They have to understand when they have unrealistic expectations."

Cisco SMB Support Assistant

Customer Advocacy has been a key organization and philosophy at Cisco ever since the 1980s. Today, with SMBs facing increased complexity in their networks and computer installations, Cisco is focusing on improving the way it delivers support offerings.

“We’re investing more than US\$2 billion in new products and services specifically designed for SMB customers, so we need to pay attention to their needs when it comes to service and support,” says Wim Elfrink, senior vice president of Customer Advocacy.

Cisco offers a variety of services designed to give SMBs options, flexibility, simplicity, and predictability. A new program, Cisco SMB Support Assistant, is designed to simplify network operation, while also keeping the cost predictable and creating a more flexible portfolio of services to meet the unique needs of SMB customers.

“Service options need to be granular enough so that the customer gets the right service at the right price,” says Elfrink. “But you have to balance this with simplicity and manageability.”

Cisco provides a range of annual support offerings that enable SMBs to continue focusing on resources to run

their businesses, while still being able to afford industry-leading support.

For example, if a customer or a partner needs around-the-clock support and requires deep technical expertise, the SMARTnet service features a range of predictable service levels and prices. SMB Support Assistant is more appropriate for customers that need business-hours support and want to improve their productivity, and need simple, easy-to-use support.

For an SMB with limited IT bandwidth, minor events can create major disruptions. Having the appropriate level of support to address its unique needs provides cost-effective assurance upon which a company can build its network strategy.

SMBs need to think carefully about vendors that gain financially when a customer has a problem and calls for support. Sometimes “pay as you go” service agreements can become extremely expensive.

Cisco SMB Support Assistant can address network configuration and connectivity issues, hardware failures or software bugs. The offering combines basic diagnostic and troubleshooting assistance with software bug fixes and product replacement.

Group Dynamics

Frequently, there are multiple relationships to manage. It may be unlikely that you’ll find a company that is equally expert at supporting networks, computers, phones, Web-based e-commerce, and application software. Roll in the issues of branch offices and home-based workers in different regions and the complexity increases.

You may want to work with a partner that has created its own network of partners, so that you have a single chief support partner who is the “master” of a group of subcontractors.

Unfortunately, she notes, some providers don’t want to take on that liability. “You won’t find a partner to help you in everything.”

Thus, you need to account for these “group dynamics” as well when devising contracts. Chan recommends including a clause that requires your vendor to help you configure each new employee’s computer and network-access needs.

More important, recommends Krakauer, is the idea of looking at your business from a different perspective than you normally do when you’re thinking about support. What’s your most crucial system? Is it the one that handles sales? The one that handles accounting? Is it e-mail?

“Rather than starting with the technology and looking out at customers, start from the outside and look in,” she says. Then you’ll understand the customer experience you are trying to provide, and what you need to proactively protect to make it happen. ■

FURTHER READING

- Cisco SMB Support Assistant
cisco.com/packet_171_9b1
- iQ Article: Service Providers Are Helping Out
cisco.com/packet_171_9b2

Rate-Based Satellite Control Protocol

Enhancing IP Performance over Satellite Connections

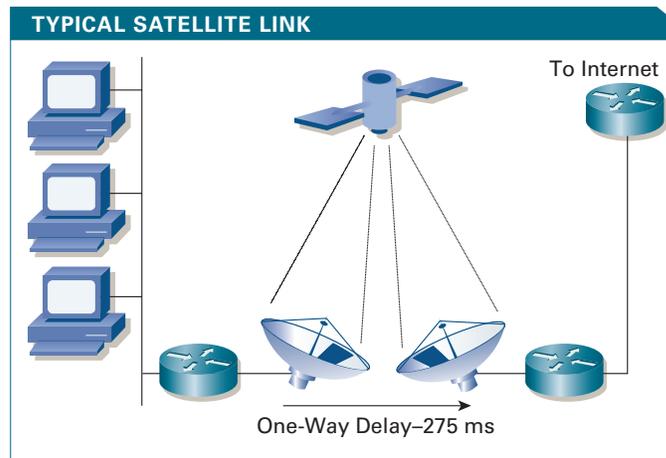
By Peter Lei and Randall Stewart

Rate-Based Satellite Control Protocol (RBSCP), a new feature in Cisco IOS Software Release 12.3(7)T, is designed for wireless or long-distance delay links with high error rates, such as satellite links. Using tunnels, RBSCP can improve the performance of certain IP protocols, such as TCP and IP Security (IPSec), over satellite links without breaking the end-to-end model. Organizations are increasingly deploying satellite technology to reach the “last mile” in remote locations that require high-speed, broadband Internet access (see figure). This technology is typically provided by a geostationary two-way satellite. The result is a *satellite round-trip time*, or *s-RTT*, of 500–600 milliseconds (ms) and error rates that are much higher than typical wired technologies such as DSL or cable. (*s-RTT* is the round-trip time between routers closest to the satellite link, *not* the typical RTT time measured between end hosts. For a satellite link, the end-host RTT would be the *s-RTT* plus any additional Internet and intranet transit times.)

Although poor performance is common with these connection types, mitigation techniques are available for some of the problems.

The two predominant Internet transport protocols, TCP and Stream Control Transmission Protocol (SCTP) use very similar congestion control techniques to prevent congestion collapse. (For details on congestion collapse, refer to RFC 896 at cisco.com/packet/171_11a1.) First, any lost packet is assumed to be a sign of congestion, which causes an exponential backoff in the sending rate. Second, both protocols use slow start and congestion avoidance to slowly ramp up the sending rates, which helps prevent a sender from flooding the network with packets. The traditional term for these techniques is *Additive Increase Multiplicative Decrease (AIMD)*. AIMD works well in the normal wired Internet and has been quite successful at preventing congestion collapse while ensuring users a “fair share” of the available bandwidth. However, when these techniques are applied to a satellite link, the results can be less than optimal.

Satellites tend to have long *s-RTT* times, so it takes longer for a connection to ramp up its sending rate. In addition, packets are often dropped due to errors on the link; each packet drop is interpreted by the AIMD algorithm as an indication of congestion, which results in collapsing the sending window (*cwnd*) to one maximum transfer unit (MTU) (If the drop is detected using a fast retransmit, AIMD will only halve the *cwnd*.) The end result of high error rates combined with long delays is that the sender often



SIGNAL DELAY Performance in satellite links is limited by the delay inherent in geosynchronous systems and the probability of bit errors in any wireless system.

stays collapsed in slow start, sending only one packet per RTT. These AIMD techniques can quickly make a satellite flow of 1.5 Mbit/s appear as slow, or slower, than a 56-kbit/s dialup line.

Another factor, even when there are no errors, is the receiver window (*rwnd*). In both TCP and SCTP an *rwnd* provides end-to-end flow control. Each acknowledgment that a receiver sends contains an *rwnd* that specifies how much available data buffer space the receiver has that the sender can use. Therefore, for each RTT a sender can have no more than *rwnd* bytes outstanding toward the receiver. In a typical wired network this is not a problem, but with the long *s-RTT* a connection with a small *rwnd* becomes quite slow.

For example, consider a traditional cross-town Internet connection with a 20-ms round-trip time and a 32-kilobyte (KB) window. On this connection the maximum transfer rate possible would be 50 windows of 32 KB of data every second, or about 13 Mbit/s. While this is more than adequate for today’s DSL or cable technology, which provide about 3 Mbit/s or less, if you take the same connection and route it using a 550-ms *s-RTT* the result is a 570-ms RTT. This gives you 1.75 windows of 32 KB of data per second, or about 450 kbit/s per second. The satellite link can be rated to send 1.5 Mbit/s, but in reality less than one-third of the rated speed is utilized.

The new tunneling technology provided by RBSCP helps with this problem without breaking the end-to-end model. For more details on RBSCP, the solutions available for handling poor satellite connection performance, and the various tradeoffs, visit Packet Online at cisco.com/packet/171_11a2. ■

PETER LEI is a technical lead in the IP Technologies Engineering Group at Cisco. He can be reached at peterlei@cisco.com.

RANDALL STEWART is a senior software engineer in the IP Technologies Engineering Group at Cisco specializing in IP transport technologies, and the primary author of SCTP. He can be reached at rrs@cisco.com.

SPOTLIGHT ON:

Cisco ONS Family: New and Enhanced Solutions for Edge, Network, and Core Optical Networks

The Cisco ONS family offers several platforms for transporting traffic and delivering high-speed, high-capacity services over optical fiber networks. Several new and enhanced Cisco ONS platforms are now available.



With the new Cisco ONS 15310-CL SONET Multiservice Platform, service providers can deploy private-line and switched Ethernet services that support link capacity adjustment scheme (LCAS), virtual concatenation (VCAT), and generic framing procedure (GFP) without redesign or disruption of the entire network. This compact, one-rack-unit edge platform is interoperable with infrastructure using the Cisco ONS 15454 SONET Multiservice Provisioning Platform (MSPP). The enhanced Cisco ONS 15302 and Cisco ONS 15305 SDH MSPP models also support similar services.

For metro networks, the Cisco ONS 15454 MSPP now offers a Carrier Ethernet (CE)-Series line card and high-density DS3 line cards. The CE-Series line card supports LCAS, VCAT, and GFP and helps carriers migrate to optimized data services from one platform. The high-density DS3 line cards help service providers deliver time-division multiplexing (TDM) services more cost effectively, save termination costs, and free shelf slots for new service offerings.

For core networks, Cisco offers a new any-service, any-port (ASAP) line card that enables software-selectable bandwidth and protocol options as well as faster broadband switching functionality for the Cisco ONS 15600 SONET/SDH Multiservice Switching Platform (MSSP). Through small form-factor pluggable (SFP) optics, this line card provides Gigabit Ethernet bandwidth connectivity based on OC-3, OC-12, and GFP. The new Cisco ONS 15600 MSSP single-slot cross-connect (SSXC) card enables the network to transparently pass SONET/SDH overhead traffic and supports more than 32 two-fiber, bidirectional line-switched rings.

Two new line cards are available for aggregation applications using the Cisco ONS 15530 DWDM Multiservice Aggregation Platform. The four-port 1-Gbit/s or 2-Gbit/s Fibre Channel/FICON Aggregation Card aggregates up to four 2-Gbit/s Fibre Channel or FICON services on a single 10-Gbit/s wavelength to support high-density storage area network (SAN) extension applications. An eight-port Multiservice Muxponder supports consolidation of numerous storage, data, voice, and video services over dense wavelength-division multiplexing (DWDM) using optical and copper client interfaces.

cisco.com/go/optical

Edge Routing, Access, and Aggregation

Cisco Service Control Engine Products

Cisco Service Control Engine (SCE) products help service providers maximize the use of network resources, control service delivery, and accurately bill for value-added broadband services. These devices process broadband subscriber traffic in a point of presence (POP), cable headend, or distribution hub. Cisco SCE 2000 Series engines provide line-speed processing of 4-Gbit/s traffic over 2-Gigabit links, managing up to 2 million concurrent unidirectional application flows. Cisco SCE 1000 Series engines process 1-Gbit/s traffic over 2-Gigabit links, supporting 1 million flows. Used with SCE devices, the Cisco Service Control Application Suite for Broadband encompasses three software applications: the Service Control Application for subscriber service monitoring, the Cisco Collection Manager for capturing and reporting service data, and the Cisco Subscriber Manager for individualized traffic accounting and control.

cisco.com/packet/171_npd1 



Core Routing

Cisco CRS-1 8-Slot Single-Shelf System

Designed for service providers building IP next-generation networks, the Cisco CRS-1 8-Slot Single-Shelf System is the newest member of the Cisco CRS-1 Carrier Routing System family. The CRS-1 8-slot system is deployed in a service provider's point of presence (POP) and supports a total switching capacity of 640 Gbit/s in half of a standard 19-inch rack. The CRS-1 8-slot

system is half the size and capacity of the previously released 16-slot system. In a multishelf configuration, the Cisco CRS-1 can reach a maximum capacity of 92 Tbit/s. The CRS-1 8-slot system's midplane design provides slots for eight modular services cards and four fabric cards in the rear of the chassis, as well as eight interface modules and two route processors in the chassis front.
cisco.com/go/crs

Switching

Cisco Catalyst 6500 Series Switches: New Supervisor Engine and Modules

New hardware for Cisco Catalyst 6500 Series switches offer capabilities for a variety of deployments. The Cisco Catalyst 6500 Supervisor Engine 32, with Policy Feature Card 3B, extends hardware-based security features to the network edge and provides two 10 Gigabit Ethernet XENPAK-based uplinks or eight Gigabit Ethernet small form factor pluggable (SFP)-based uplinks. New Cisco Catalyst 6500 LAN Access Interface modules include a 48-port 10/100/1000 module, and a 48-port 10/100 module with enhanced quality of service (QoS) and cable fault-detection capabilities. A 96-port 10/100 (RJ21) module provides high port densities in a compact form factor. All of these modules support standards-based Power over Ethernet (PoE) with field-upgradeable daughter cards. Additional new hardware options include a 48-port 100BASE-X module to support highly secure 100-Mbit/s fiber deployments and a 6000-Watt power supply for high-density PoE deployments.
cisco.com/go/catalyst6500

Cisco Catalyst 4500 Series Switches: New Supervisor Engine and Power Supply

New hardware choices for the Cisco Catalyst 4500 Series offer greater flexibility for switch deployment. The Supervisor Engine V-10GE includes dual, wire-speed 10 Gigabit Ethernet ports (x2 optics) and four alternatively wired Gigabit Ethernet ports (SFP optics) on the faceplate. The triple-input 1400-Watt DC power supply is optimized for central-office deployments by service providers. Multiple inputs enable technicians to connect the supply to smaller fuses and breakers and customize the output power to meet site-specific application needs.
cisco.com/go/catalyst4500

Cisco Catalyst 3750 and Cisco Catalyst 3560 Series Switches: New 10/100/1000 Ethernet and PoE Models

Four new models each for Cisco Catalyst 3750 and Cisco Catalyst 3560 Series switches give enterprises a broader range of 10/100/1000 Ethernet and PoE deployment options with integrated security, availability, and quality of service. The Cisco Catalyst 3750 and 3560 Series switches provide 24 or 48 ports of 10/100/1000 Ethernet, four SFP-based Gigabit Ethernet ports, and optional support for PoE, all in a single-rack-unit chassis.
cisco.com/go/catalyst3750
cisco.com/go/catalyst3560



Security and VPNs

Cisco IPS 4240 and Cisco IDS 4255 Sensors

New products for intrusion detection and prevention provide inline, real-time traffic analysis to protect against malicious or unauthorized network activity. The Cisco IPS 4240 Sensor provides intrusion protection with 250-Mbit/s performance in switched environments, on multiple T3 subnets, and, by using 10/100/1000 interfaces, on partially utilized Gigabit

Ethernet links. The Cisco IDS 4255 Sensor supports 600-Mbit/s performance to protect gigabit subnets and traffic that traverses aggregation switches connected to numerous subnets. Both platforms initially support four on-board 10/100/1000 monitoring interfaces for copper links.
cisco.com/go/ids



Content Networking

Cisco Content Switching Module with SSL

The Cisco Content Switching Module with SSL (CSM-S) for the Cisco Catalyst 6500 Series combines high-performance Layer 4 to Layer 7 content switching with integrated Secure Sockets Layer (SSL) acceleration. This combination provides scalable performance, connection persistence, and ensured uptime for business-critical applications and offloads the CPU-intensive task of processing SSL transactions in backend servers. The integration of content switching and SSL also enables Layer 7 load balancing while ensuring that data remain encrypted while on the network.
cisco.com/go/csm-s

Wireless

Cisco Aironet 1130AG Series and Cisco Aironet 1230AG Series Access Points

The Cisco Aironet 1130AG Series 802.11a/b/g Access Point simplifies deployment of a wireless LAN in offices and similar RF environments. This model features integrated antennas and dual IEEE 802.11a and 802.11g radios for predictable coverage and a combined traffic capacity of 108 Mbit/s. The Cisco Aironet 1230AG Series Access Point supports dual-band 802.11a/b/g radios with dual antenna connectors for wireless LANs in rugged environments or installations that require specialized antennas. The Cisco Aironet 1230AG Series combines antenna versatility with high transmit power, receive sensitivity, and delay spread for reliable performance and throughput in high multipath and indoor environments.

cisco.com/go/aironet/abg 



Storage Networking

Cisco File Engine Series Appliances

The new Cisco File Engine Series appliances simplify management and increase protection of file-based data located at enterprise branch offices. Based on Cisco Wide Area File Services (WAFS) technology, the Cisco Edge File Engine is deployed at each branch office and replaces local file and print servers. The Cisco Core File Engine is deployed at the data center and connects directly to one or more file servers or network-attached storage (NAS) gateways for processing WAN-optimized file requests on behalf of each Edge File

Engine. The Cisco WAFS Central Manager software provides management and monitoring functions for all file engine devices. The Cisco File Engine Series is covered in greater detail in the article beginning on page 35.

cisco.com/packet/171_npd2 



Voice and Video

Cisco IP Phone 7971G-GE

First in the industry, the new Cisco IP Phone 7971G-GE brings the benefits of a Gigabit Ethernet network to the desktop. The two-port Ethernet switch allows for direct connection to a 10/100/1000BASE-T Ethernet network through an RJ-45 interface with a single LAN connection for both the phone and a collocated PC. Like the Cisco IP Phone 7970G, the Cisco IP Phone 7971G-GE offers a high-resolution, color touch screen for display and control of user features. The IP phone provides access to eight telephone lines and can be powered through Power over Ethernet (PoE) or a local power supply.

cisco.com/packet/171_npd3

Cisco EGW 2200 Enterprise Gateway

The new Cisco EGW 2200 Enterprise Gateway facilitates a phased migration from traditional private branch exchange (PBX) networks to converged Cisco IP communications systems, while offering important network call routing and number analysis capabilities. A software application that runs on select Cisco Media Convergence Servers, the Cisco EGW 2200 supports interworking of voice signaling protocols, such as Digital Private Network Signaling System (DPNSS) and Q.SIG, with Cisco CallManager and Cisco Unity voice mail and unified messaging solutions. Commonly used DPNSS and legacy voice-mail features are supported by the Cisco EGW 2200 Enterprise Gateway.

cisco.com/packet/171_npd4

ABOUT NEW PRODUCT DISPATCHES

Keeping up with Cisco's myriad new products can be a challenge. To help readers stay informed, *Packet* magazine's "New Product Dispatches" provide snapshots of the latest products released by Cisco between October 2004 and January 2005. For real-time announcements of the most recently released products, see "News Release 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/.

Networked Home

Linksys Wireless A/G Media Center Extender

The Linksys Wireless A/G Media Center Extender enables users to wirelessly stream digital entertainment content such as music, videos, or photos that are stored on a Microsoft Windows Media Center PC to connected televisions or stereo systems around the home. Users can also watch and pause live television shows or make digital recordings for later viewing. The Media Center Extender connects to a home stereo or television using standard consumer electronic cables and communicates with the Media Center PC via a home network that uses Wireless-A, Wireless-G, or 10/100 Ethernet cabling.

cisco.com/packet/171_npd5

Linksys Wireless A+G Products

New Linksys products allow home users to share a wireless network using equipment compatible with Wireless-A or Wireless-G standards. The Dual-Band Wireless A+G Broadband Router provides Internet connectivity and a four-port, full-duplex 10/100 Ethernet switch to connect up to four PCs in the home; the router can connect to additional hubs and switches for a larger network. The Dual-Band Wireless A+G Broadband Router also contains two wireless access points, each supporting all three wireless networking specifications. Three new adapters provide dual-mode wireless access by client devices: the Linksys A+G PC Card for notebooks, Linksys A+G PCI adapter for desktop computers, and the Linksys A+G USB Adapter.

cisco.com/packet/171_npd6

Linksys Wireless-G Internet Video Camera

The Linksys Wireless-G Internet Video Camera sends live video with sound through the Internet to a Web browser. The camera contains its own Web server that enables it to connect directly to a network, either over Wireless-G (IEEE 802.11g) or 10/100 Ethernet cable. The advanced MPEG-4 video compression in the camera produces a high-quality, high frame rate, audio/video stream at up to 640x480 resolution. Security Mode sends a message with a short video attached to as many as three e-mail addresses whenever the camera detects motion in its field of view.

cisco.com/packet/171_npd7

Linksys Wireless-G CompactFlash Card

The Wireless-G CompactFlash (CF) Card enables wireless networking on a personal digital assistant (PDA) with the PocketPC 2002 or PocketPC 2003 operating systems. The CF Card installs directly into the PDA using a CompactFlash Type II slot and communicates over wireless networks at speeds up to 54 Mbit/s. The IEEE 802.11g CF Card is also compatible with Wireless-B (802.11b) networks. The card allows the PDA to roam seamlessly among multiple 802.11g/b access points or routers, communicate without an access point to download data from a wireless PC, and share data directly with other wireless PDAs.

cisco.com/packet/171_npd8

Security Challenges, Continued from page 22

itself. For example, an IPS device can catch an application-embedded attack that a firewall might miss. However, the IPS device might not have the appropriate enforcement action that the firewall offers for dealing with the attack. By converging firewall and IPS capabilities, network security administrators have all the mitigation actions and resilience of a firewall with all the inspection intelligence of an IPS.

An additional limitation of IPS devices, however, is that while they have a fine-grained view of network traffic, they are signature-based; that is, they must receive updates that tell them what to look out for. Signature updates can take from 24 to 48 hours, making them ineffective against tomorrow's flash threats. This is where network antivirus software comes in, with its dynamic outbreak prevention updates. Antivirus software can be updated very quickly and can disseminate the information rapidly through an infrastructure to all endpoints. If this infrastructure is merged with IPS and firewalls, companies gain more than just the power of each: they now have a security threat defense system, a way to rapidly update information and deeply analyze packets for identification of worms and viruses, as well as the firewall capability to block those packets from entering the network and a solution that is highly resilient.

This type of systems approach transforms security from operating as separate siloed technologies in a reactive mode—with limited and static detection methods—to functioning as a coordinated, proactive threat defense system that adapts to the threat environment.

According to Pope, these systems will provide numerous benefits: improved detection, greater event classification accuracy, lower operating costs, streamlined administration, and services extensibility that integrates the most advanced security technologies as they are developed. Most importantly, these converged systems will not compromise the quality of security in any given category, but instead combine the strength of each in complementary ways to deliver a tighter, coordinated defense. ■

FURTHER READING

- Cisco Security and VPN Information
cisco.com/go/security
- Cisco Self-Defending Networks
cisco.com/go/sdn
- Cisco Intrusion Prevention Alert Center
cisco.com/go/ipsalert
- SANS Institute Internet Storm Center
isc.sans.org
- eSecurity Planet Online
esecurityplanet.com
- SecurityTracker
securitytracker.com
- "Are hackers using your PC to spew spam and steal?" (*USA Today*, September 2004)
usatoday.com/tech/news/computersecurity/2004-09-08-zombieuser_x.htm
- "Code that steals for its creators" (*NetworkWorld-Fusion.com*, March 2004)
nwfusion.com/weblogs/security/004453.html

Load Balancing with Cisco CSM, CSS, and Their SSL Modules

The Cisco Networking Professionals Connection is an online gathering place to share questions, suggestions, and information about networking solutions, products, and technologies with Cisco experts and networking colleagues. Following are excerpts from a recent Ask the Expert forum, “Advanced Load Balancing with the Content Switching Module (CSM), Content Services Switch (CSS), and Their Secure Sockets Layer (SSL) Modules,” moderated by Cisco’s Gilles Dufour. To view the full discussion, visit cisco.com/packet/171_10a1. To join in on other live online discussions, visit cisco.com/discuss/networking.

Q: *Is asymmetric traffic allowed through the CSM? For example, can a real server issue a request to a backend server (database server) through the CSM (the real server is configured with a default route pointing to the CSM) and receive a reply through a router without having to go through the CSM?*

A: There are scenarios where it is possible. For example, if you are in bridge mode, and the traffic initiated by the real server does not hit a vserver, then you can have the return traffic bypassing the CSM. Another working scenario is when the connection initiated by the real server hits a vserver (bridge or routing mode does not matter). In this case, you can configure the vserver with the command `unidirectional` to tell the CSM to keep the flow alive as long as it sees traffic coming from one side. Finally, in routing mode and if traffic does not hit a vserver, I believe it may work as well as long as the traffic seen by the CSM comes from a real server. I recommend creating a vserver to catch the traffic from the real server going to the backend server, and using the `unidirectional` option.

Q: *Can you schedule a TCL script to run on the CSM at a specific time?*

A: You can’t schedule script on the CSM. The CSS can do this but not the CSM.

Q: *Can you please tell me any command that helps in load balancing EIGRP [Enhanced Interior Gateway Routing Protocol], as load balancing is not taking place between A end and B end? We are using 64 kbit/s. I have tried using `ip load-sharing per-packet` command, but still load balancing is not taking place.*

A: This type of load balancing requires you to have at least two routes to the same destination. These routes should be equal and appear in the routing table. If you have these two routes, `ip load-sharing per-packet` will only work if you have Cisco Express Forwarding or no `ip route-cache`.

If you do not have two equal routes, you must work on getting them. This matter would be more appropriately addressed by the Routing/Switching Discussion Forum on the Networking Professionals site.

Q: *I need to do a dual keepalive (http and https) to provide a successful response. I assume a script could accommodate this. Are there any other methods available? For example, if http and https both succeed, then service is up.*

A: Let me first say we don’t have https keepalive on either the CSS or CSM. All you can do for https is check whether you can open a TCP connection with the server. If you are OK with just checking the port, you will indeed need a script to combine the two keepalives into a single one.

Q: *Besides http and https, can CSS or CSM support other protocols, such as FTP, SSH [Secure Shell], etc.? If we define the Web other than standard port (TCP 80 or 443), will CSS or CSM still treat the traffic as Web traffic?*

A: CSM and CSS do understand other protocols such as FTP and Real-Time Streaming Protocol (RTSP), and this applies whichever port is being used.

Q: *I will have a server farm with identical Web application servers and will be running SSL for data protection. Can the Cisco CSS 11501 without SSL termination load balance SSL Web application servers traffic if the load balancing is not based on higher-layer application data? In such a situation, what are the choices for load balancing apart from simple round robin?*

A: SSL is a TCP protocol, so you can do load balancing based on IP or TCP. I would recommend round robin or leastconn for the balancing method. You should also use the `advanced-balance ssl` command to enable stickiness based on SSL ID.

Do you have a question about load balancing with the Cisco CSM, CSS, and their SSL modules? Ask the NetPro Expert. Send your question to packet-netpro@cisco.com, with the subject line “Advanced Load Balancing.” ■



GILLES DUFOUR, CCIE No. 3878, is a customer support engineer in the Cisco Technical Assistance Center, Europe, the Middle East, and Africa (EMEA). He has been a member of the content networking team since June 2002. He can be reached at gdufour@cisco.com.

Recipe for Resilience, Continued from page 39

An industry first, Cisco Survivable Remote Site Telephony (SRST) for branch office routers maintains basic phone service in case of interruption in WAN service to the headquarters CallManager cluster. In the case of a failure, the local Cisco IOS Software-enabled SRST router provides basic call processing, and also supports secure calls with authentication and encryption for signaling and media transmission for Cisco IP phones using Cisco CallManager 4.1.

For enterprises with frequent WAN outages and fewer than 240 employees per branch office, Cisco CallManager Express delivers a full set of commonly used key system and low-end private branch exchange (PBX) features. For added resilience, enterprises can run Cisco CallManager Express on two routers and use HSRP to provide immediate failover capabilities or, for added transparency, run CallManager Express on one router and SRST on another. Cisco Unity Express provides branch users with local integrated voice mail and automated attendant features.

Resilient Workforce

An enterprise without a workforce resilience strategy squanders its most valuable asset—its people. Workforce resilience empowers employees with anytime, anywhere access to corporate resources, fortifies applications with identical services, and provides contingencies for dealing with disruptions. Anytime, anywhere access gives employees the flexibility to do their jobs and live their lives—increasing productivity and morale. A workforce-resilient organization could include the following:

- Mobility features via Cisco Aironet wireless LAN solutions provide campus-wide network connectivity through wireless access in conference rooms, cafeterias, and patios
- IP Phone Extension Mobility allows users to log into any Cisco IP Phone in any corporate or branch office and receive identical service as at their desk phones, especially voicemail access
- Client VPN services and Cisco SoftPhone provide remote data and telephony services to mobile workers through broadband or wireless LAN services in hotel rooms, conference centers, airline lounges, and other hotspots
- Router-based broadband VPN services and Cisco IP Phone or Cisco SoftPhone enable secure connectivity from a home office

An enterprise business continuance plan must include workforce resilience solutions. If a primary data center is lost to fire or flood, for instance, IT systems automatically reroute active sessions to the secondary data center, where employees continue to access business-critical services without a hiccup. Or during hazardous weather, employees can work from home using the same services they would have while in the office.

In extreme cases, Cisco customers have set up ad-hoc wireless LANs in hotel rooms with a secure broadband connection to the corporate network.

Workforce resilience technologies should deliver powerful attributes that are easy to configure, requiring minimal management and user training. For example, the Cisco Business Ready Teleworker uses a single broadband connection for PC and IP phone access to the corporate network, with voice and data encryption over an IPSec VPN for security. The remotely manageable Cisco 800 Series Router contains preconfigured VPN, security, and QoS features that are transparent to users. Teleworkers receive the same services—such as video-and audioconferencing, business-critical applications access, and IP phone extension services—they would have while sitting in a cubicle at their corporate headquarters.



Careful consideration should be given to the factors that contribute to highly available IT systems—both hardware and software—and the resources and solutions allocated to building resilience throughout the network, applications, communications, and workforce. With this done, enterprises will have taken the most important steps in turning their business resilience strategies from concept to delivery. ■

 **PACKET ADVERTISER INDEX**

ADVERTISER	URL	PAGE
ADC - The Broadband Company	www.adc.com/truenet	D
AdTran	www.adtran.com/info/wanemulation	2
Aladdin Knowledge Systems	www.eAladdin.com/Cisco	IFC
American Power Conversion (APC)	http://promo.apc.com	13
BellSouth Business	www.bellsouth.com/business/answers	OBC
Boson Software	www.boson.com	A
Cisco Press	www.ciscopress.com	B
Cisco Systems	www.cisco.com/poweredby	F/18/33
eiQ Networks	www.eiqnetworks.com	6
Global Knowledge	www.globalknowledge.com/train4free	4
Interstar Technologies	www.faxserver.com	80
NetScout	www.netscout.com/ad/cii	54
Network General	https://networkgeneral.mnl.com/c1	44
New Edge Networks	www.newedgenetworks.com	48
NIKSUN	www.niksun.com/packet	20
OPNET Technologies	www.opnet.com	68
Panduit	www.panduit.com/pp12	IBC
Pulver.com	www.von.com	58
Solsoft	www.solsoft.com/packet	8
SurfControl	www.surfcontrol.com/go/cisco	26
Trend Micro	www.trendmicro.com/cisco	62/63

CACHE FILE

Snippets of Wisdom from Out on the Net

CYBER QUOTE

“The first rule of any technology used in a business is that automation applied to an efficient operation will magnify the efficiency. . . . automation applied to an inefficient operation will magnify the inefficiency.”

—Bill Gates
Chairman, Microsoft Corporation

Content-Based Online Activity Leads in the US

Content ranked as the leading US online activity at the end of 2004, according to a study by the Online Publisher's Association. Conducted in conjunction with Nielsen//NetRatings, the study tracked the online activities of 40,000 Internet users, parsing their activity into four categories: commerce, communications, content, and search. Content—defined as Websites and Internet applications designed primarily to deliver news, information, and/or entertainment—led across the study's metrics. The report attributes increased broadband penetration as a main contributor to the rise in content-based activity. More on this and other broadband trends can be found at clickz.com/stats/.

Where Is Your Blog?

Blogs (short for weblogs) are proliferating on the Internet—doubling every five months over the last year and a half, according to blog analysis firm Technorati (www.technorati.com). The current number of blogs is now more than 8 times bigger than the 500,000 blogs the firm measured in June 2003. Technorati tracked 3 million blogs as of the first week of July 2003, and has added more than 1 million blogs since then. Meanwhile, Pew Internet & American Life (www.pewinternet.org) reports that a new weblog on the Internet is created every 5.8 seconds. These personal journals, frequently updated and intended for general public consumption, typically represent the personality of the author or reflect the purpose of their hosting Website.

Banner Ads Grow on European Websites

The volume of banner ads on European Websites grew 24 percent last year, from 76,375 in November 2003 to 94,939 in November 2004, according to research conducted by Nielsen//NetRatings (nielsen.netratings.com). Sweden, France, and The Netherlands led European growth in banner ads—each registering more than 30 percent increases. Lagging were Germany, Norway, Spain, and Belgium with growth of 10 percent or less.

Net Lingo

Packet monkey—Someone who intentionally inundates a Website or network with data packets, resulting in a denial-of-service situation for users of the attacked site or network. Packet monkeys typically use tools created and made available on the Internet by hackers (whatis.com).

THE 5TH WAVE



“He saw your laptop and wants to know if he can check his Hotmail.”

©The 5th Wave, www.the5thwave.com