

PACKET

CISCO SYSTEMS USERS MAGAZINE

FIRST QUARTER 2006



UNLEASH YOUR NETWORK SERVICES

The National University of Singapore Did 22

How to Achieve Network Nirvana 28

Behind the Scenes of a DMVPN Deployment 48

Navigate the Winding Road to IMS 61



CISCO.COM/PACKET

PACKET

CISCO SYSTEMS USERS MAGAZINE

FIRST QUARTER 2006
VOLUME 18, NO. 1



ON THE COVER

Unleash Your Network Services

22

Cisco Service-Oriented Network Architecture outlines how enterprises like the National University of Singapore can evolve their network to increase efficiencies, lower costs, and strengthen business agility.

Network Nirvana

28

Through the process of simplification, networks achieve new levels of resiliency—and lower total cost of ownership.

Push to Talk Everywhere

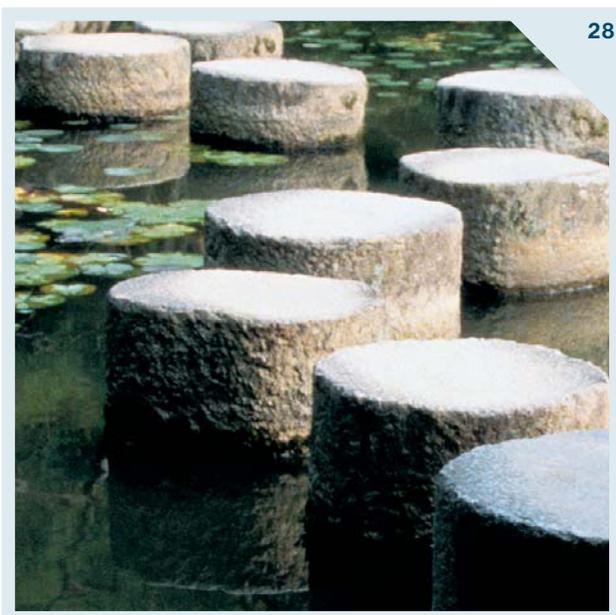
34

Cisco IPICS creates communications interoperability by joining radio systems with IP networks.

Minding the Store

38

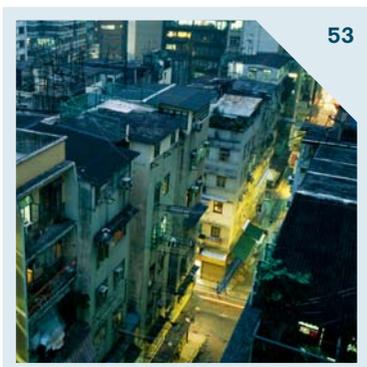
As IT managers build out their organizations' storage-area networks, they face performance, security, and management issues.



COVER PHOTO: Tommy Hor, director of the computer center at the National University of Singapore



43



53



65

IN EVERY ISSUE	
Mail	3
Acquisitions	5
Tech Tips	11
Advertiser Index	73
Cache File	74
The 5th Wave	74

TECHNOLOGY

ROUTING: EIGRP Efficiency	15
Best practices for scaling EIGRP neighbors in hub-and-spoke networks.	

SWITCHING: Dynamic Buffer Limiting	19
Industry's first hardware and flow-based congestion avoidance at wire speed.	

ENTERPRISE SOLUTIONS

Primary Wireless	43
Employees at Intel's Jones Farm campus will soon be using a wireless LAN as the primary access method for data, voice, and video.	

DMVPN Deployment	48
A British broadcaster moves from a mixed-media architecture to MPLS with the help of mGRE technology.	

SERVICE PROVIDER SOLUTIONS

Say It with IPTV Services	53
Tapping the flexibility of a converged network infrastructure is smart business for Hong Kong Broadband Network.	

Gaining the Ethernet Edge	57
New OAM protocols enhance a carrier's service deployments and make managing and monitoring Metro Ethernet networks easier.	

The Winding Road to IMS	61
Non-IP Multimedia Subsystem applications are still key in the service provider evolution toward IP Multimedia Subsystem (IMS).	

SMALL AND MIDSIZED BUSINESSES

Modular to the Core	65
Modular features and advanced services modules make the Catalyst 6500 Series Switch an affordable option for midsized networks like Maryland's nonprofit Columbia Association.	

DEPARTMENTS

From the Editor	1	New Product Dispatches	68
Big Changes for <i>Packet</i>		What's new from Cisco over the past quarter.	
User Connection	5	NetPro Expert	72
Networking Academy Advanced Technology Training • New Edition of Routing TCP/IP		Advice from Cisco's Jazib Frahim on implementing and troubleshooting IPSec redundancy.	
Tech Tips & Training	7		
Enterprising MPLS • Top 5 Freeware Tools • Reader Tips			

PACKET MAGAZINE

David Ball
Publisher and Editor in Chief

Jennifer Redovian
Executive Editor

Susan Borton
Managing Editor

Suzanne Jackson
Joanie Wexler
Contributing Editors

Robert J. Smith
Sunset Custom Publishing
Project Manager

Nicole Collins, Amy Mackey,
Mark Ryan
Sunset Custom Publishing
Production

Jeff Brand
Art Director

Emily Burch
Designer

Ellen Sokoloff
Diagram Illustrator

Bill Littell
Print Production Manager

Valerie Marliac
Promotions Manager

Richard Koh, Eightfish Ltd.
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 address corrections and other correspondence direct to packet@cambeywest.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, iQ, Linksys, *Packet*, and PIX are registered trademarks or trademarks of Cisco Systems, Inc., and/or its affiliates in the USA and certain other countries. All other trademarks mentioned in this publication are the property of their respective owners.

Packet copyright © 2006 by Cisco Systems, Inc. All rights reserved. Printed in the USA.

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

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

This magazine is printed on recycled paper.



10%
TOTAL RECOVERED FIBER

FROM THE EDITOR

Big Changes for *Packet*

Cisco is a company that practices what it preaches: that Internet technology will truly change the way we live, work, play and learn. By IP-enabling many of its own business processes, Cisco generates gains in productivity and business agility that translate to millions of dollars in savings each year. The fact that Cisco uses its own networking technology to gain a competitive edge is clearly demonstrated in the Cisco IT@Work CDs we have distributed with this magazine over the past year.

Whether mobilizing its workforce with secure wireless networking, improving customer service via IP contact centers, or increasing employee uptime and reducing travel expenses with desktop video conferencing, Cisco prides itself on being its best reference customer. Our offices are quickly becoming showcases unto themselves for demonstrating how networking technologies can be turned into real business differentiators.

Cisco migrated away from printed technical documentation in favor of electronic distribution years ago. Every internal transaction, from signing a purchase order, submitting an expense report, or rewarding an employee, is done completely via Web-based applications. Not a piece of paper in sight. So it should come as no surprise that Cisco will be placing significant emphasis on the digital distribution of *Packet* magazine in the coming year.

Cisco launched *Packet Digital Edition* last June. In just six months, we went from zero to nearly 25,000 subscribers worldwide. Clearly, readers are responding positively to the new format. However, not everyone is ready to jump headfirst into the paperless society. Many of you still prefer reading a printed publication, enjoy the mobility of print, or just plain like the feel of a real magazine in your hands.

For this reason, we have decided not to do away with the print edition of *Packet*. Instead, we will be offering readers a choice: You can subscribe to *Packet Digital Edition* for free, or pay a subscription price to continue to receive the print edition of *Packet*.

Starting with the Third Quarter 2006 issue, *Packet* will be a paid subscription publication. You will receive one more free issue after this one you hold in your hands.

But what an issue it will be!

Starting with the Second Quarter 2006 issue, *Packet* is being revamped and redesigned to be the premier publication for increasing the knowledge and expertise of Cisco networking professionals. That means more tech tips, more deployment guides and troubleshooting tricks, more how-to articles and configurations, and more real-world examples than ever before.

We think you'll like what we've done and hope you'll give the new *Packet* magazine a try, in whichever format you choose.

David Ball
Editor in Chief
daball@cisco.com



Rob Brodman

MAIL

More Long-Lasting Routers

The recent reader letters about router uptime that you published in your past two issues sparked me to look at our own routers on the network.



```
BRISDCNR01>show ver
Cisco Internetwork Operating System
Software IOS (tm) 3660 Software (C3660-
JS-M), Version 12.1(2)T, RELEASE SOFT-
WARE
(fc1)
Copyright (c) 1986-2000 by cisco Systems,
Inc.
Compiled Tue 16-May-00 19:43 by ccai
Image text-base: 0x60008900, data-base:
0x613B4000
```

```
ROM: System Bootstrap, Version 12.0(6r)T,
RELEASE SOFTWARE (fc1)
```

```
BRISDCNR01 uptime is 4 years, 40 weeks, 5
days, 17 hours, 1 minute System returned
to ROM by power-on System restarted at
14:26:24 gmt Tue Dec 12 2000 System image
file is "flash:c3660-js-mz.121-2.T.bin"
```

```
cisco c3660 (R527x) processor (revision
C0) with 83968K/14336K bytes of memory.
Processor board ID JAB0429844K
R527x CPU at 225Mhz, Implementation 40,
Rev 10.0, 2048KB L2 Cache Bridging soft-
ware.
X.25 software, Version 3.0.0.
SuperLAT software (copyright 1990 by
Meridian Technology Corp).
TN3270 Emulation software.
```

```
3660 Chassis type: ENTERPRISE
2 FastEthernet/IEEE 802.3 interface(s)
8 Serial network interface(s)
DRAM configuration is 64 bits wide with
parity disabled.
```

125K bytes of non-volatile configuration memory.

16384K bytes of processor board System flash (Read/Write)

Configuration register is 0x2102

No problemo.

—Robert McCallum CCIE No. 875,
THUS, Glasgow, Scotland

Unsafe Driving



Recently, while browsing through your Third Quarter 2005 issue, as ever looking for more excellent material to show my students, I was appalled to find that a video clip on “Internet Benefits” (part of the enclosed Cisco IT@Work CD) showed an extremely dangerous work practice—specifically someone driving a car with one hand while dialing on his mobile phone with the other. By all means promote more efficient work practices, but not at the expense of endangering lives.

—Austin Kinsella, Institute of Technol-
ogy, Carlow, Ireland

First-Time Reader

As I'm writing this e-mail I'm holding in my hand my first-ever *Packet* magazine (Fourth Quarter 2005). I am so excited to be among those reading it as I start my Cisco career as a consultant. I have read almost all of the articles included in this issue and I love every word of it. This valuable tool will not only help me grow my business and knowledge, but will also inform others how Cisco is making a difference in this field. Keep up the good work. I am now officially a *Packet* subscriber and a Cisco super fan.

—Erasmus J. Medina, London,
Ontario, Canada



A Party for a Router

Router uptime seems to be a big deal these days. I guess we're geeks. After reading all the recent letters from your readers, we found that we have a Cisco 2610 Router located at our Hobby Lobby store in Des Moines, Iowa, that has been up for almost five years and is not on a UPS. Here's where the geek part comes in: We had a party to celebrate the five-year uptime. Fun was had by all and the cake was great (see photo). Thanks for making durable and reliable products.

—Thomas R. Tucker, Matt Bowman,
and Sherry Bowman, Hobby Lobby
Network Services, Oklahoma City,
Oklahoma, USA

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.

Cisco Networking Academy Advanced Technology Training

The Cisco Networking Academy has broadened its security and wireless training programs. Course enhancements and exam updates include the following.

Network Security 2.0

- Replaces the “Fundamentals of Networking Security” course.
- Combines Cisco PIX Firewall and Cisco IOS Software integration; includes new course content.
- Covers advanced topics; curriculum updates complement the Security+ exam.
- Offers new equipment bundles, Cisco 1800 Series Integrated Services Router and Cisco PIX 515 Firewall, as the standard classroom package.
- Incorporates newer tools and technologies to align with certification exams.

Wireless LANs 1.2

- Includes lab updates of 802.11g.
- Offers new, dynamic simulations.

For more information about the Cisco Networking Academy Program, visit cisco.com/go/netacad. ■

New Edition of Routing TCP/IP, Volume I

Cisco Press has introduced the second edition of *Routing TCP/IP, Volume I*. The new edition includes information about protocol changes, and describes Cisco features that are used to enhance routing integrity, secure routers from attacks initiated through routing protocols, and provide greater control over the propagation of routing information for all IP interior routing protocols. Information within each section of the new edition is enhanced and modified to include new developments in routing protocols and Cisco implementations.

About Cisco Press

Cisco Press is the Cisco Systems authorized book publisher of Cisco networking technology and Cisco certification self-study materials, and Cisco Networking Academy Program materials for networking students and professionals. Leading authorities from Cisco and other industry innovators write and contribute to the titles and series that make up the Cisco Press product family. Products from Cisco Press are part of a recommended learning path from Cisco.

For more information about *Routing TCP/IP, Volume I, Second Edition*, visit cisco.com/packet/181_3b1. ■

Stay Informed with Beyond Basic IP

Cisco-based networks are moving far beyond basic IP communications to gain the software intelligence required to self-adapt to network conditions. Beyond Basic IP (BBIP), a Cisco monthly electronic publication, keeps networking professionals informed about the latest networking software innovations. Each BBIP issue covers topics such as network security, voice over IP (VoIP), high availability, high-performance design, troubleshooting, and quality of service (QoS).

To read the current issue of BBIP and to register for a free subscription, visit cisco.com/packet/181_3c1.

Recently Announced Cisco Acquisitions

Acquired		Employees	Location
Cybertrust's Intellishield Alert Manager	Web-based security intelligence service that provides daily information about information security threats and IT product vulnerabilities that affect the entire corporate information technology domain. Intellishield Alert manager will become part of the Cisco MySDN intelligence website. The Intellishield team will join the Technical Support Service's group in Cisco's Customer Advocacy organization.	26	Herndon, Virginia, USA
Scientific-Atlanta, Inc.	Provider of set-top boxes, end-to-end video distribution networks, and video system integration. Scientific-Atlanta will become a division of Cisco's Routing and Service Provider Technology Group.	7,500	Lawrenceville, Georgia, USA

Enterprising MPLS

Thinking About Moving to an MPLS MAN or WAN? Here Are Some Things to Consider.

By Janet Kreiling

Many enterprises that have relied on ATM or Frame Relay transport for their metropolitan-area networks (MANs), WANs, or both, are now exploring IP Multiprotocol Label Switching (MPLS) for its flexibility and scalability. In some cases, the impetus comes from the service provider, who invites an enterprise customer to take advantage of the benefits of its core MPLS network. In other cases, the impetus is from the enterprise itself. In either instance, most IT departments report knowing less about MPLS as an enterprise tool than they would like, whether the enterprise is setting up an MPLS network itself or ensuring that an outsourced MPLS WAN meets its needs.

MPLS networks, incidentally, are just one alternative for MAN and WAN deployments. To date, most enterprises using IP across their WAN have chosen IP Security (IPSec), while others opt for multiple virtual routing and forwarding (VRF) segmentations or encrypted private connectivity. Like MPLS, each of these is a viable solution that depends on the needs of the enterprise (see sidebar, page 8).

The primary market for IP MPLS MAN/WAN networks is composed of organizations that use Layer 2 technologies such as time-division multiplexing (TDM) circuit switching, Frame Relay, ATM, or SONET/SDH for transporting data and possibly voice. According to Kevin Loo, solutions manager in the Enterprise Systems Engineering group at Cisco, these organizations are looking for provisioning flexibility, broad geographic availability, little or no distance sensitivity in pricing, the ability to mix and match access speeds and technologies, and to segment multiple departments or operating units, applications, and services securely within a single network. As with any architecture that provides VPNs over shared wide-area or metro-area facilities, MPLS networks can yield significant cost savings while providing fully meshed connectivity among locations, high inter-location bandwidth, and end-to-end quality of service (QoS).

If your enterprise is considering switching to MPLS, should you install and manage the MPLS network in-house or outsource those tasks to a service provider? The largest organizations, with the IT expertise and desire for full control over their internal networks,

might choose to self-deploy MPLS. For the lion's share of enterprises, however, outsourcing will be a more practical choice.

Outsourcing Your MPLS Network

Simplifying operations, lowering costs, and mitigating the risk of changes in technology are among the top reasons companies choose to outsource rather than build their own MPLS MANs or WANs, according to a June 2005 report conducted by Forrester Research. When evaluating different service providers, you should keep these and other business goals in mind. For example, will the proposed network be able to link business applications, demarcate sufficiently and securely among various segments of your business, or deliver the needed bandwidth to different locations?

The link between the enterprise and service provider networks can take place at either Layer 2 or 3. With peering at Layer 3, the provider's network routes the customer's IP packets through its shared network, while guaranteeing secure transport. It does this by installing a VRF table for each customer that isolates that customer's traffic from others. One of the advantages of Layer 3 peering is that the two networks can exchange routing information directly, says Bob Vigil, an engineer in the Service Provider Systems Engineering group at Cisco. In addition, most service providers can provide QoS with greater intelligence in Layer 3 than in Layer 2. MPLS also offers true any-to-any connectivity and thus more efficient routing.

Any Transport over MPLS (AToM), in which Layer 2 packets or cells are carried over an MPLS network, is a good solution for some enterprises, especially those with ATM, Frame Relay, or Ethernet networks that need point-to-point Layer 2 connectivity, says Vigil. The virtual point-to-point circuits characteristic of Layer 2 networks are set up through VPNs.

Whether the managed MPLS service requested is at Layer 2 or Layer 3, many of the key questions enterprises should ask prospective providers are the same.

Availability, Addressing, and Topology

Availability is critical. What is the provider's availability percentage? What backup is offered? Is the fallback Frame Relay, ATM, leased lines, or the public Internet? How will your data be protected?

Enterprise MAN/WAN Deployment Options

In addition to IP MPLS, Cisco offers the following solutions for enterprise MAN/WAN deployments:

Encrypted private connectivity adds Advanced Encryption Standard (AES) as well as Digital Encryption Standard (DES) and Triple DES (3DES) to existing Frame Relay, ATM, or other links. It's ideal for enterprises with moderate growth expectations that require secure, dedicated connectivity.

IPSec VPNs use the same strong encryption standards for transport over public and private networks; most often used to link to branch offices and teleworkers.

Multi-VRF segmentation is an extension of VPN technology that helps keep traffic segregated across the WAN. Used with IP VPNs, IPSec VPNs, and encrypted private connections, multi-VRF segmentation is designed to enhance security between departments, business functions, and user groups.

Then, starting at your enterprise's front door, how will IP addressing be handled on the links to and from the provider? Often the provider determines the addressing scheme, which might use either the provider's or enterprise's address space, private address space, or unnumbered addressing on the link. When the enterprise's private address space is used, whoever does the addressing must ensure that the customer retains some private address space and that standards-compliant addresses transmitted across the link are not confused with either the enterprise's or provider's own addresses.

Network topology is another key consideration. MPLS networks inherently provide an any-to-any architecture; Layer 2 networks are often hub-and-spoke. Enterprises might prefer Layer 2 networks if they want control over communications between spokes or to maintain firewalls between them. For these same reasons, enterprises choosing an MPLS network should make sure that the provider can maintain a hub-and-spoke design within it. Alternatively, if the enterprise moves from a hub-and-spoke design to an MPLS network, Vigil points out, it may find that it reduces the number of physical and virtual circuits, depending on the existing deployment.

Routing and Routing Convergence

How many routing prefixes will the provider accept? Which routing protocol will be used—Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), or Border Gateway Protocol (BGP)? If the provider is managing the enterprise-provider link, the provider is responsible for choosing the protocol and maintaining the link. Does the provider impose limits on the number of routes your organization can use? What happens if the number is exceeded?

Routing convergence for your enterprise network will depend on how the provider handles convergence within its network. What are the approximate and maximum convergence times if a link fails or a route is unavailable? Does the provider offer a convergence guarantee in its service-level agreement (SLA)? Can the provider deliver remote access to the MPLS VPN for telecommuters, mobile workers, and other remote users? When an enterprise site has multiple links to the provider, who manages load balancing? Will the provider route IP VPNs via VRF tables from other companies if you want to create an extranet for suppliers or customers?

QoS, Security, and SLAs

Of course, you want your critical, delay-sensitive traffic to receive the appropriate priority across the provider's shared network as well as within your own MAN or WAN. What classes of service does the provider offer, and how many? Can the provider map your LAN classes, which are typically more numerous, to its own end to end across the network? If your traffic leaves the provider's network for another network the provider has interconnected with to increase its coverage area, will the end-to-end QoS and security be maintained?

If you need a specific capability such as multicasting, not all providers can support it. Some may still carry multicasts through Generic Routing Encapsulation (GRE) tunnels rather than in native form, according to Loo. If the provider supports multicast VPNs, how many multicast distribution trees does it offer?

Perhaps highest on the list of considerations is security. The level of security inherent in an MPLS network is the same as that in ATM or Frame Relay networks, but some questions are nonetheless in order. Given that the network infrastructure is shared, what practices are in place to prevent misconfiguration that could breach your network security? MPLS networks are inherently resistant to label spoofing and other attacks, but you should understand how the provider protects its core network from attacks of all kinds and whether the MPLS network and public Internet traffic use the same routers or other equipment. If so, can you get dedicated provider edge (PE) routers? What's the cost?

Finally, don't forget your SLA. Among the items the SLA should cover specifically are agreed-upon deliverables; consequences if the provider fails to deliver; definitions of network availability; resolution response times for problems that might arise and lead times for new service, VPN, and site requests; bandwidth; latencies; QoS; and SLA reporting. Cisco has tools that can help you monitor your provider's compliance with the SLA: IP SLA and the Optimized Edge Routing (OER) feature in Cisco IOS Software. For instance, if the MPLS cloud exceeds the permitted delay, OER

reroutes traffic using an alternate path such as a GRE tunnel over the Internet or another MPLS provider.

Deploying Your Own MPLS Network

According to Forrester Research, companies who choose to install and manage MPLS MANs and WANs themselves already have the hardware and expertise in house, want to retain control over their networks, and believe they can save costs by doing it themselves. Loo adds that the ability to segment their networks into closed user groups (CUGs) makes self-deployment particularly interesting to large organizations with multiple units that want to keep internal networks separate. For example, a large university might find that logically separate networks for individual departments work better administratively. Segmentation overall is increasingly important to large enterprises, says Loo. In addition to keeping critical applications separate and secure, some companies might want to ensure that top executives are on VPNs unavailable to other employees or set up CUGs that give limited access to customers, partners, or groups outside business unit or departmental boundaries.

An enterprise can create multiple virtual networks that share the same infrastructure, much as a service provider creates virtual networks when it manages MPLS networks for multiple enterprise customers. If, for example, a worm infects a PC on one virtual network, the threat will not spread to other virtual networks. It is also a good idea to establish address transparency across organizations for scalable shared, or virtualized, services such as firewalls, intrusion prevention, and other security capabilities.

Enterprise IT staffers deploying their own MPLS network should be schooled in routing protocols including BGP, EIGRP, OSPF, Label Distribution Protocol (LDP), and, of course, MPLS. An understanding of how these protocols interact with backdoor links is recommended as well. Enterprise IT will also need to handle load balancing, because MPLS networks commonly have multiple paths available between any two points. Several mechanisms can be used for efficient routing, such as Cisco Express Forwarding or unequal cost load balancing.

Convergence and its effect on network performance is as important in self-deployed MPLS networks as in outsourced ones. Of particular concern are convergence in the backbone, in a VPN site, and VPN route redistribution times. "Intelligent network design can help create faster convergence times," says Vigil. "Network engineers can optimize convergence times by tuning the timers listed for each routing protocol and monitoring network load and application requirements."

Some enterprises might use a combination of self-deployed and outsourced MPLS networks, reserving the self-deployed MPLS network for what it deems most critical or in need of segmentation.

Cisco Gear for Enterprise MPLS

With more than 300 MPLS customer deployments, including more than 70 enterprise installations, Cisco has the expertise to help businesses deploy and manage MPLS networks. Enterprises that manage their own MPLS networks have the choice of customer edge (CE), provider edge (PE), and provider (P) routers; those who outsource might be able to affect what CE routers their provider offers. Either way, the advantages of an end-to-end Cisco network demonstrated by a Cisco Powered Network provider extend through to the CE and into the LAN.

Recommended PE and P routers include the Cisco 12000, 7600, and 7200 series, and the 7304 Router. The 12000 Series offers a modular, distributed architecture that can scale from 2.5 Gbit/s to nx10 Gbit/s capacity per slot and support up to 1,000 VRF instances. One of the highest performance and density platforms in the Cisco edge/aggregation routing portfolio, the 7600 Series offers a complete Layer 2 or Layer 3 MPLS solution, full QoS, modular security services, and support for up to 1,000 VRF instances. The most widely deployed MPLS router, the 7200 Series offers flexibility with comprehensive IOS routing, QoS, and security services support at up to OC-3 speeds. The 7200 also supports up to 1,000 VRF instances and complete Layer 2 and Layer 3 VPN services. The Cisco 7304 Router supports complete AToM features, up to 1,000 VRF instances, and hardware-accelerated performance with QoS.

For outsourced MPLS networks, CE router choices include the Cisco 7200 Series as well as Cisco 3800, 2800, and 1800 Series Integrated Services Routers. Depending on the model, the Integrated Services Routers support five to 1,000 VRF instances and 10,000 to 1 million VRF routes, with concurrent services at up to T3/E3 speeds.

Whether you are outsourcing your MPLS WAN to a provider or deploying it yourself, with WAN traffic going over a provider network at some point, you should look into whether the provider has a Cisco Powered Network designation, advises Loo. Capabilities such as availability, QoS, and integrated security are more likely to be found with these providers who employ Cisco equipment and technologies end to end and meet Cisco's support standards (find them at cisco.com/go/cpn).

◆ ◆ ◆
Whether self-deployed or outsourced, usually the move to an MPLS network does not entail a forklift upgrade. Industry data shows that enterprises are already moving gradually to MPLS IP VPNs. A company might begin simply by upgrading the Cisco IOS Software feature sets on its routers and switches, which can provide some of the routing and security capabilities of an MPLS network. For many enterprises, MPLS WANs will eventually be a reality, and users and IT staffs alike will be glad to have them. ■



Top 5 Freeware Tools for Cisco Network Administrators

Would you like to efficiently manage your Cisco network using the tools commonly used by experts? These five freeware tools are easy to set up and use and can help you manage your Cisco network.

1 PuTTY for SSH access (putty.nl) is a Secure Shell (SSH) client that runs on Windows. You can use PuTTY to access remote devices using the SSH protocol. With an increased emphasis on security, most Cisco devices can now be remotely accessed via SSH. PuTTY supports SSH versions 1 and 2 along with Data Encryption Standard (DES), Triple DES (3DES), Blowfish, and Advanced Encryption Standard (AES). With a single click of a button, you can log a session for later review. PuTTY also supports Telnet.

2 PumpKIN TFTP server (kin.klever.net/pumpkin) Most Cisco devices use the Trivial File Transfer Protocol (TFTP) as the primary way to transfer system image files or configuration files, so a simple, stable TFTP server is an essential part of a network administrator's toolkit. PumpKIN TFTP server provides a simple, easy-to-use GUI and runs on all versions of Windows. PumpKIN also plays *.wav* files to provide audio alerts indicating the state of a TFTP transfer—a feature that is very handy for multitasking network administrators.

3 Kiwi Syslog server (kiwisyslog.com) Because most Cisco devices use the Syslog protocol to generate system messages, deploying a centralized Syslog server is often recommended. Additionally, network administrators might occasionally need a local Syslog server that can be quickly deployed for testing or troubleshooting purposes. Kiwi Syslog server is an easy-to-use but versatile tool that runs on all versions of Windows. Kiwi Syslog provides the built-in ability to listen to Simple Network Management Protocol (SNMP)

traps, too. An additional useful feature is that the Kiwi will monitor the free disk space on the server itself. When the disk space falls below a configured threshold, Kiwi generates e-mail and audio alerts—a very useful feature that prevents messages being lost due to lack of disk space.

4 Nmap Network scanner (nmap.com) Originally developed as a hacker tool, mainstream networking professionals now use Nmap for security testing of firewalls, routers, and intrusion detection system (IDS) devices. Nmap can generate ping sweeps to scan a range of IP addresses, verify the working of specific TCP or UDP ports on a target machine, quickly scan the TCP or UDP ports on a target host, and use OS fingerprinting techniques to make an educated guess about the OS running on the target devices. Nmap works with both Windows and Linux. Nmap use may be prohibited in many secure environments and can trigger IDS alarms, so always check your corporate security policies before using Nmap.

5 Protocol Analyzer or Network Analyzer (ethereal.com) Ethereal is a commercial grade network analyzer that can be installed over any Windows or Linux/UNIX platform. Ethereal can analyze more than 700 protocols, including Cisco specific protocols such as Cisco Discovery Protocol and Inter Switch Link.

—Submitted by Anand Deveriya, CCIE No. 10401, NEC Unified Solutions, (author of Network Administrators Survival Guide, Cisco Press, 2005)

Tech Tips

Configure the Cisco VPN Client for Third-Party Client Software. Learn how to configure a Cisco VPN Client to coexist with third-party client software, including Microsoft, Nortel, Checkpoint, Intel, and others. Compatibility includes the ability to use other VPN products while the VPN Client is installed.

cisco.com/packet/181_4e1

Learn how to Move Cisco CallManager to a New Location.

This document outlines how to move Cisco CallManager with or without an IP address change.

cisco.com/packet/181_4e2

Configuring a Wireless Domain Services Access Point as an AAA Server. Learn how to configure an access point to provide Wireless Domain Services (WDS) and perform the role of an authentication, authorization, and accounting (AAA) server. Use this setup when there is no external RADIUS server to authenticate infrastructure access points and client devices that participate in WDS.

cisco.com/packet/181_4e3

Configure Access Point ACL Filters. Get help configuring access control list-based filters on Cisco Aironet access points using the command-line interface (CLI).

cisco.com/packet/181_4e4

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 Configuring Large Switch Environments

In a large switch environment, to configure all or multiple interfaces on a switch with the same configuration parameters, do the following:

```
Switch(config)# interface range [ interface { port
range } ]
```

For example:

```
Switch(config)#interface range fastEthernet 0/1 - 30
```

To configure different ports with the same configuration:

```
Switch(config)#int range fa0/1 , fa0/12 , fa0/13
```

—Rajesh Kumar Ojha, Habib Bank Limited, Karachi, Pakistan

Troubleshooting

TIP Backup Solutions for Frame Relay Point-to-Point Networks

When implementing a Frame Relay point-to-point network, you might want to have a backup solution in place. Most backup methods are triggered by either a routing protocol failure or an interface going down. When using the interface method, problems can arise if the Frame Relay provider is using multiple switches between the two end points. A point-to-point interface on one end can be brought down due to a link failure, although it remains up on the other end. This can cause the backup solution to fail. Frame Relay end-to-end keepalives (EEK) can resolve these types of issues. For example:

```
Hostname R5
!
interface Serial0/0.30 point-to-point
 ip address 163.1.30.5 255.255.255.0
 frame-relay interface-dlci 504
 class EEK
!
!
map-class frame-relay EEK
 frame-relay end-to-end keepalive mode request
!
```

```
Hostname R4
!
!
interface Serial0/0.30 point-to-point
 ip address 163.1.30.4 255.255.255.0
 frame-relay interface-dlci 405
 class EEK
!
!
map-class frame-relay EEK
 frame-relay end-to-end keepalive mode reply
!
```

To verify that this is working properly:

```
R5#show frame-relay pvc 504
PVC Statistics for interface Serial0/0 (Frame Relay
DTE)
DLCI = 504, DLCI USAGE = LOCAL, PVC STATUS = ACTIVE
(EEK UP), INTERFACE = Serial0/0.30
```

```
input pkts 5161          output pkts 6022
in bytes 56771          dropped pkts 0
out bytes 331070        in pkts dropped 0
in pkts dropped 0      out pkts dropped 0          out bytes
dropped 0
in FECN pkts 0         in BECN pkts 0
out FECN pkts 0        out BECN pkts 0           in DE pkts 0
out DE pkts 0
out bcast pkts 860     out bcast bytes 294936
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
pvc create time 14:19:49, last time pvc status
changed 13:45:49
!
```

—Mike Griffin, Robert Half International Inc., Pleasanton, California, USA

TIP Pinging Multiple IP Addresses with the put Command

You can use the `put` command to ping multiple IP addresses simultaneously from the router. The following is an example of pinging IP addresses

192.168.26.1, 192.168.25.1, and 192.168.17.2:

```
tc1sh
foreach addr {
192.168.26.1
192.168.25.1
192.168.17.2
} {puts [exec "ping $addr"]}
```

—*Santosh Gamre, CCIE No. 13265, IPsoft, Inc., New York, USA*

Editor's Note:

This is indeed a good, simple use of TCL. However, there are better ways to ping multiple IP addresses. You can use either IP SLA (the old SAA or RTR feature) or the CISCO-PING-MIB. IP SLA has a very usable

command-line interface and the reliability and accuracy of the probes, whether through pings or another mechanism, is considerably better than from the command line or from the CISCO-PING-MIB. It is not difficult to create an IP SLA config to do this. While the CISCO-PING-MIB does not have a CLI interface, you can write one in TCL and do the same thing as the script above. The advantage of these strategies is that the process doing the pings is considerably lighter weight than TCL and, therefore, has much less impact on device performance. This is much more important as soon as you make the first obvious enhancement to this script, which is to ping every <x> seconds.

SUBMIT A TIP

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

EIGRP Efficiency

Scaling EIGRP Neighbors in Hub-and-Spoke Networks

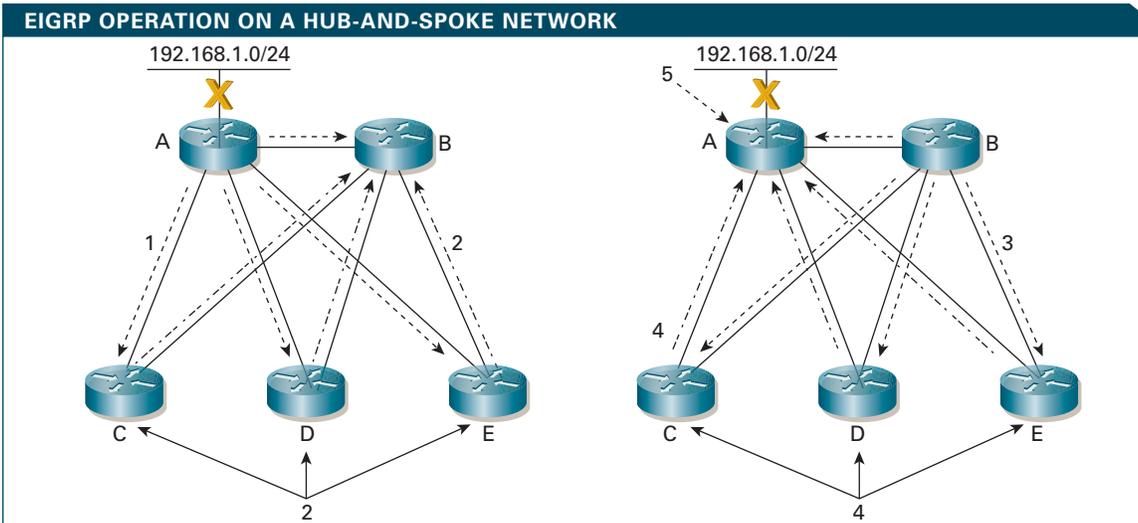


FIGURE 1 The query process typically has a heavy impact on EIGRP scaling in a hub-and-spoke network. The hub router must send and track queries for each route that it marks active.

By Russ White

While many people count sheep to get to sleep, most network engineers find other things to count, such as IETF RFCs, and how many remote sites they can connect to a hub-and-spoke Enhanced Interior Gateway Routing Protocol (EIGRP) network. How many EIGRP neighbors can you have on a hub-and-spoke network? Looking at current best practices and testing can help answer the question.

Native Mode Dual-Homed EIGRP Hub and Spoke

To understand how EIGRP normally operates in a hub-and-spoke network, refer to the simple set of network events in Figure 1.

1. A loses its connection to 192.168.1.0/24, marks the route active, and sends queries to each of its neighbors. B, C, D, and E mark the route as active and send queries to their neighbors.
2. At this point, the timing of the operations depends on the speed of the links, router types, router processor load, and other things. For this example, assume that C, D, and E query B before B can send its queries to the remotes.
3. B now has no neighbors from which it has not received a query, so it marks 192.168.1.0/24 as unreachable and sends replies to each of its neighbors.

4. C, D, and E now have no possible paths to 192.168.1.0/24, so they mark it as unreachable, and send replies to A.

5. When A receives these replies, it determines there are no other paths to 192.168.1.0/24, so A removes the route from the local routing table and sends an update for 192.168.1.0/24 with an infinite metric.

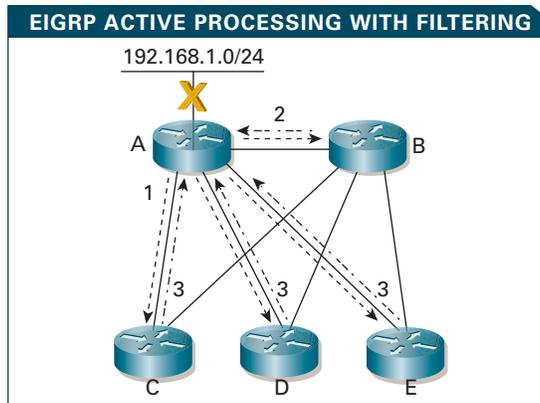
As you can see, EIGRP scaling on a hub-and-spoke network is heavily impacted by the query process. The hub router must send and track queries for each route that it marks active.

Filtering Toward Remotes

The first practice network administrators can use to increase EIGRP scaling on a hub-and-spoke network is to reduce the number of routes sent to the remote neighbors through filtering and aggregation. This reduces the number of queries and replies sent, even in this simple network (Figure 2). Assume that Routers A and B are filtering or aggregating routing information so that the default route 0.0.0.0/0 is the only route transmitted to remote routers C, D, and E.

1. A loses its connection to 192.168.1.0/24, marks the route active, and sends queries to each of its neighbors.
2. B determines its only path to 192.168.1.0/24 is through A. Why? Because C, D, and E are not even

FIGURE 2 Filtering and aggregation is a good way to increase EIGRP scaling on hub-and-spoke networks because this reduces the number of routes sent to the remote neighbors.



receiving this route, so they are not advertising it back to B.

3. C, D, and E have no alternate route to 192.168.1.0/24, so they reply to A. Router A receives these replies and notes that all queried neighbors have replied. A marks this route as unreachable and removes it from the local routing table.

Additionally, processing could be greatly reduced by configuring the remote routers as EIGRP stubs.

EIGRP Stub Routers

When considering this small network, one point is obvious: the remote routers will never be used to transit traffic between the two hub routers, A and B. Consequently, remote routers C, D, and E will never have an alternate path to any destination that the hub router has learned about through any other path. Therefore, there is no reason for the two hub routers to query the remote routers.

The EIGRP stub feature works from the assumption that a router configured as an EIGRP stub signals its neighbors that it will never have any valid alternate paths, so there is no reason to query the stub router. Using the same small network, with remote routers C, D, and E configured as EIGRP stub routers, what does the processing look like when A loses a route? Figure 2 illustrates.

1. A loses its connection to 192.168.1.0/24.
2. A marks the route as active, examines each of its neighbors, and determines which neighbor could have an alternate path to 192.168.1.0/24. Because



RUSS WHITE, CCIE No. 2635, is a technical leader in the Cisco IP Technologies Group, where he specializes in designing and implementing routing protocols and scalable networks. He is a frequent contributor to *Packet* and *IP Journal*, and can be reached at riw@cisco.com.

C, D, and E are advertising themselves as stub routers, A does not need to query these neighbors. A sends a query to B.

3. B has only one path to 192.168.1.0/24, through A, so it will mark the route as unreachable and send a reply to A. Why? Because C, D, and E are advertising themselves as stub routers, which means they will never have an alternate path.
4. Router A receives this reply from B, marks this route as unreachable, and removes it from the local routing table.

Configuring the remote routers as EIGRP stubs dramatically decreases the processing at the hub routers. How much of a difference does this make?

Neighbor Counts with and Without Stub Routers

This question has two answers, one based on real-world experience, and the other based on lab testing. In real deployments without the remotes configured as EIGRP stub routers, the largest deployments are in the 200 neighbor range on Cisco 7200 and Cisco 7600 series routers equipped with fast processors. The primary determining factors to scaling these deployments is generally bandwidth on the links between the hub and spoke routers and the processing requirements on the hub routers. The primary mechanism used to reach these neighbor counts is aggregation or route filtering.

By comparison, deployments where remote routers are configured as EIGRP stubs commonly reach as high as 800 or more remote routers using the same hub routers, and there are some exceptional larger examples. Again, scaling these networks is very dependent on aggregation or filtering toward the remote routers; more routes transmitted toward the remote routers, lower bandwidth available between the hub and the remote, and less router processing power at the hub and remote routers all translate to lower numbers of supportable neighbors.

Comparing these two deployment types, the ratios of neighbor support with and without the remote routers configured as EIGRP stubs is about 4:1.

Cisco's Routing Protocols Verification Lab, in Research Triangle Park, North Carolina, performed tests on large-scale hub-and-spoke EIGRP networks. Interestingly, the test results are very similar to the stability points in live networks:

Using standard EIGRP (no stub routers configured, no summarization or filtering), the network begins to destabilize when between 250 and 300 neighbors are attached. At this point, the network takes about 9 minutes to converge initially, and takes hours to converge if a single hub router fails.

Configuring the dual-homed remotes as stub routers, the network begins to destabilize when between 800 and 1,200 neighbors are attached. The network takes about 9 minutes to converge initially, and about 30 seconds to converge if a single hub router fails.

Essentially, lab tests conducted on large-scale hub-and-spoke EIGRP networks net results similar to real-world experience. The primary differences involve traffic levels between the remote sites and the hub routers. The lab tests do underscore the differences in convergence time after configuring the remote sites as stub routers; this is something we cannot test under real-world conditions, in most cases, with real traffic flowing through the network. We still see about a 4:1 ratio between stub and nonstub neighbors in a large-scale EIGRP hub-and-spoke network.

Remote Sites with Multiple Routers

One specific challenge arises when attempting to scale to large numbers of EIGRP neighbors on networks with two or three routers in a site that is considered a stub. Figure 3 illustrates this situation.

In this network, Routers A and B are hubs, C and D are connected together in one site, while E and F are connected together in another site. If Router C is configured as a stub router, it will not advertise 192.168.1.0/24 to Router D, so internally the site does not have full connectivity. In the same way, if the link from Router D to B fails, C will not advertise any routes learned from A, including the default route, so any hosts attached to Router D will not be able to be connected to any devices behind the hub routers.

Does this mean that multiple router sites cannot be scaled in the same ways as single router dual-homed sites, because the routers at these multiple router remote sites cannot be configured as stub routers? Actually, a new EIGRP feature in the Cisco IOS Software allows a router to be configured as a stub router, but to leak specific prefixes to a peer router. In this case, we could configure Router C as a stub router but also configure it to advertise the default route, learned from Router A, and the locally connected network 192.168.1.0/24, to Router D. In the same way, Router D can be configured as a stub router and can also be configured to leak locally connected networks and the default route learned from Router B to Router C.

The following configuration illustrates how to configure an EIGRP *leak-map* in this network.

```
router eigrp 100
  eigrp stub connected summary leak-map stubsite
  !
  route-map stubsite permit 10
    match ip prefix-list default
    match interface e0/0
  route-map foo permit 20
```

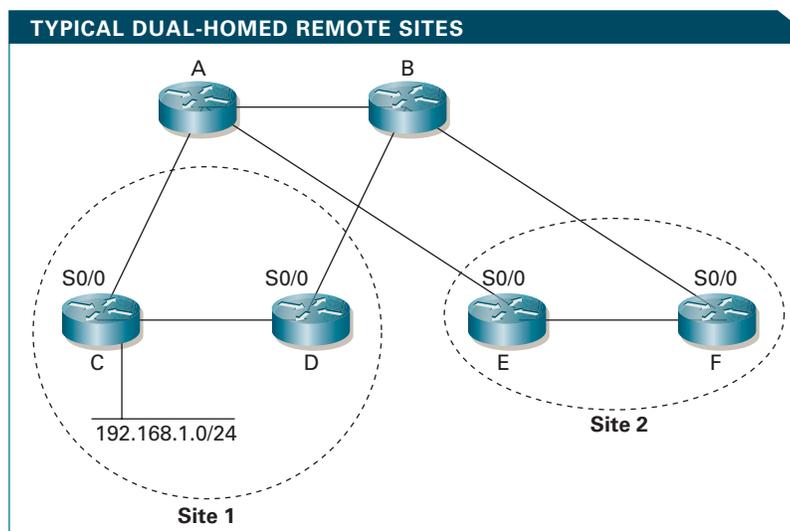


FIGURE 3 Site 1 does not have full connectivity because Router C is configured as a stub router and does not advertise the IP address to Router D.

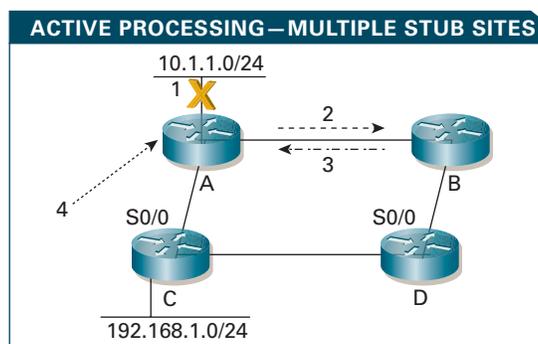


FIGURE 4 When multiple router sites are configured as stubs, the hub treats them the same as it does a remote site with a single router that is configured as a stub.

```
match ip prefix-list localroutes
match interface s0/0
!
ip prefix-list default permit 0.0.0.0/0
!
ip prefix-list localroutes permit 192.168.2.0/23 ge
/24
```

In this configuration, the prefix lists allow different parts of the address space, the default route from the hub routers toward the remote site routers, and the locally originated routes toward the hub. The interface matches tie the routes advertised to the interface facing the hub or the other remote site router.

Does this change the way the hub treats the remote site? Figure 4 illustrates the difference by showing the active process in a dual-homed multiple router site.

A loses its connection to 10.1.1.0/24. A marks the route as active, examines each of its neighbors, and determines which could have an alternate path to 10.1.1.0/24. Because C is advertising itself as stub routers, A does not need to query this neighbor. A sends a query to B.

Continued on page 73

Dynamic Buffer Limiting

Industry's First Hardware and Flow-Based Congestion Avoidance at Wire Speed

By Rupa Kaur

A Cisco innovation, Dynamic Buffer Limiting (DBL) is the first flow-based congestion avoidance quality-of-service (QoS) technique suitable for high-speed hardware implementation. Operating on all ports in the Cisco Catalyst 4500 Series Switch, DBL effectively recognizes and limits numerous misbehaving traffic flows, particularly flows in a network that are unresponsive to explicit congestion feedback such as packet drops (typically UDP-based traffic flows). It is a multiprotocol technique that can examine a flow's Layer 2/3/4 fields.

DBL provides on-demand Active Queue Management by tracking the queue length for each traffic flow in the switch. When the queue length of a specific flow exceeds its limit, DBL will drop packets or mark the Explicit Congestion Notification (ECN) field in the packet headers, so the flow can be handled appropriately by servers in the unlikely event of network congestion. Unchecked flows—also known as

belligerent or non-adaptive flows—use excessive bandwidth, and their consumption of switch buffers results in poor application performance for end users. This misbehaving traffic can also negatively affect well-behaved flows and wreak havoc with QoS.

Because DBL is implemented in ASICs, wire-speed packet manipulation is achieved without switching performance degradation. Up to 136-Gbit/s switch capacity and 102 million pps of wire-speed forwarding are supported on a single Cisco Catalyst 4500 Series Switch. This advanced QoS control is especially critical in Internet edge, distribution, and core networks.

High-Speed Networks = Greater Protection

More than ever, today's networks are built using equipment with very high bandwidth (Gbit/s) and performance (Mpps) characteristics. With more bandwidth comes greater responsibility to provide good QoS protection mechanisms. A Gigabit Ethernet network moves 1 billion bps compared to a WAN T1/T3 network (1.55 or 45 Mbit/s). Without proper

HOW DBL WORKS

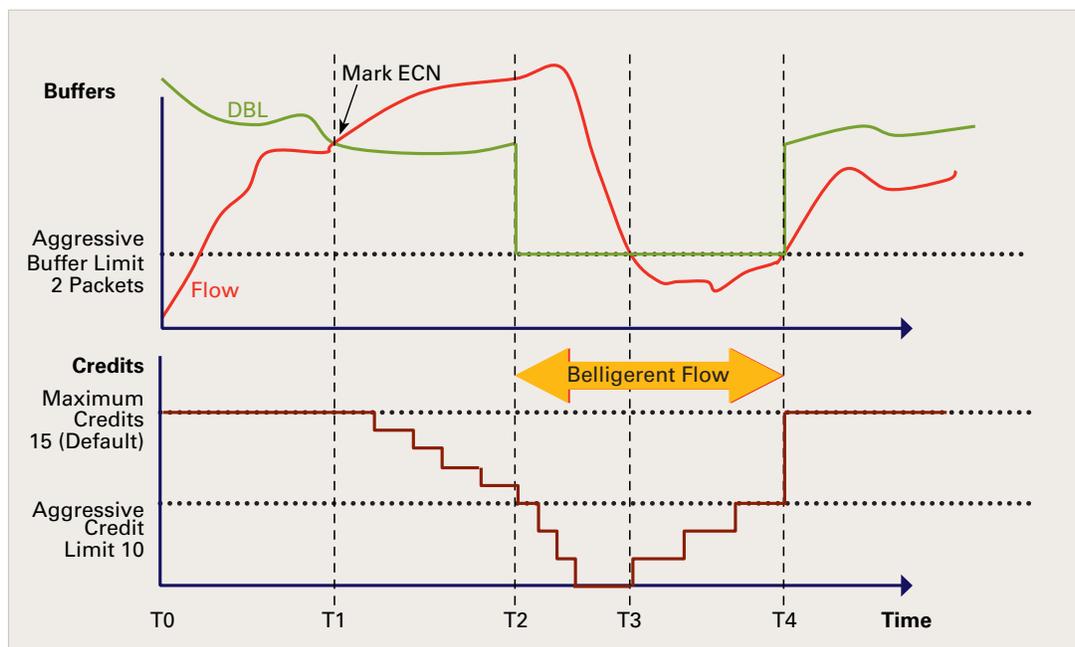


FIGURE 1 Shown here are the key ways DBL works on a given flow. The switch maintains the number of credits and buffers concurrently.

safeguards, high end-user bandwidth puts more stress on QoS. Voice quality suffers in the presence of belligerent flows, and QoS degrades when these flows cause denial of service (DoS) to well-behaved flows.

A belligerent flow travels at high speed (with multiple packets) and does not reduce its traffic rate in response to packet drops. Examples include:

- **Thick UDP flows.** Downloading a movie from the Internet is an example of a thick flow; it requires higher dynamic buffer utilization (high bps) compared to a user performing router terminal access using Secure Shell (low bps). Voice is an example of a thin flow and takes precedence using a strict priority queue.
- **Spanning Tree, IP Multicast loops.** A Spanning Tree Loop usually lasts 30 seconds, during which network bandwidth is consumed (on uplinks, for example).
- **Streaming multimedia,** common in many Internet applications.

DBL in Motion

After meeting age requirements and passing a written test, a person in the US must pass a driving test to obtain their Class C automobile license. Each candi-

date starts with 100 points. The more mistakes one makes (e.g., poor parking, exceeding the speed limit), the more points are deducted. If these points total 30 or less, the candidate is issued a license.

DBL's technique for congestion avoidance is similar. Instead of points, however, DBL uses *credits*. Credits are decremented the moment a flow misbehaves. In the beginning phase of congestion, no packet is dropped and the ECN is marked. (IETF RFC 3168 has more information on the ECN bits in the IPv4 ToS field.) A single packet is marked with ECN so the flow has a chance to adjust or retreat without losing more packets (a warning, unfortunately, the candidate in the driving analogy doesn't get).

Linux and Solaris 9 operating systems have support for ECN. Each flow starts with 15 credits (this



RUPA KAUR, a senior technical marketing engineer in Cisco's Gigabit Switching Business Unit, has been at Cisco for ten years. Prior to her role in technical marketing, she was a development engineer for ATM platforms. She can be reached at rupa@cisco.com.

parameter can be configured globally) and counts down. DBL is highly granular; it performs computation based on a flow and not on all the packets within a queue. This is the main distinction between DBL and Weighted Random Early Detection (WRED).

As shown in Figure 1, for every active flow, the switch maintains two parameters concurrently: the number of credits and buffers. When a flow consumes more buffers than the dynamically computed limit, DBL (starting at T1 in the plot) marks the ECN or drops one packet. If the flow does not respond to the single packet drop and continues to send packets at the same rate, it will lose its credits one by one.

The flow is considered aggressive when its credits fall below the aggressive credit limit. At this point, the buffer usage is also reduced to the aggressive buffer limit. On the other hand, at time T3 when the flow adjusts itself and uses fewer buffers than the dynamically computed limit, the number of credits begins to increase one at a time. Well-behaved TCP flows regain to full credits immediately.

DBL Configuration Overview

The Cisco IOS DBL configuration is very simple (see Figure 2). The switch uses the Modular QoS Command-Line Interface (MQC) with keyword “dbl” for the policy map. The ECN config is also global. All the above-mentioned buffer and credit parameters are displayed in the `show qos dbl` command.

The DBL transmit queue logic operates on each port’s four transmit queues (a total of 1,500 queues for 384 ports on a Cisco Catalyst 4510R Switch). In addition, DBL manages flows on any port type: switched, routed, trunk, access, or EtherChannel.

DBL Spanning Tree Loop Performance

Figure 3 depicts DBL’s network performance improvement in a Spanning Tree Protocol (STP) loop scenario. In this case, without DBL, the TCP flow slows and a non-looping, well-behaved 70-Mbit/s flow has 26.5-Mbit/s (37 percent) throughput. With DBL, a non-looping flow has 69-Mbit/s (99 percent) throughput in addition to TCP flows. The throughput is just one lab scenario to ensure DBL’s effect during an STP loop. Similar behavior of well-behaved and adaptive flows will be observed if the STP scenario is replaced by a Layer 3/4 deployment.

DBL, AutoQoS, and WRED

With more than 50 million ports deployed worldwide, the success of the Cisco Catalyst 4500 Series Switch lies in its centralized architecture and simple configuration macros. One such macro is AutoQoS, which automatically configures DBL and QoS on a particular port. Thousands of enterprises are using the AutoQoS feature and running DBL for their IP telephony deployments.

CISCO IOS DBL CONFIGURATION

```
4xxx(config)#qos dbl
4xxx(config)#qos dbl exceed-action ecn
4xxx# show qos dbl // Truncated
DBL flow includes layer4-ports
DBL uses ecn to indicate congestion
DBL max credits: 15
DBL aggressive credit limit: 10
```

FIGURE 2 All buffer and credit parameters are displayed.

DBL AND SPANNING TREE LOOP

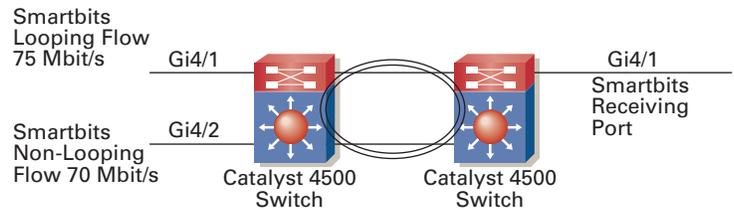


FIGURE 3 DBL improves network performance in an STP loop scenario.

WRED is another popular Active Queue Management technique; DBL and WRED can work together in any network deployment. The main difference between DBL and WRED is that DBL is flow based. If congestion occurs, the hardware logic drops a packet based on the flow.

Also, a flow’s contribution to congestion is measured. For example, if congestion occurs on an uplink wire-speed 10 Gigabit Ethernet port, the UDP-based audio packets are higher targets than the voice or Telnet packets because audio uses more buffers. DBL only operates when the pipe is full, e.g., if there is more than 20 Gbit/s of data on a 2-port 10-Gbit/s EtherChannel.

◆ ◆ ◆

DBL is a Layer 2-4 multiprotocol congestion avoidance technique that strengthens network security by limiting belligerent flows and preventing DoS attacks. Belligerent flows usually occur at the edge of the network where streaming media and other bandwidth-hungry applications wreak havoc on the network. Cisco Catalyst 4500 Series switches are positioned at the network edge to provide proper safeguards against network congestion. These switches support DBL in hardware on all Supervisors (Sup2+ and higher) with performance ranging from 64 Gbits/48 Mpps to 136 Gbits/102 Mpps. The 1-rack-unit Catalyst 4948 fixed switches also support DBL in hardware. ■

FURTHER READING

- IETF RFC 3168, “The Addition of ECN to IP”
ietf.org/rfc/rfc3168.txt
- Cisco Catalyst 4500 Series Switches
cisco.com/packet/181_5b1

Unleash Your Network



PHOTOGRAPH BY RICHARD KOH

Cisco Service-Oriented Network Architecture outlines how enterprises can evolve their network to increase efficiencies, lower costs, and strengthen business agility. The National University of Singapore knows. **By Janet Kreiling**

Services

AT LEAST ONE OF YOUR pressing business goals is probably on this list:

- Firing up your supply chain, all the way, before your stock gets thin
- Improving security or mobility
- Enabling locations anywhere in the world to work smoothly with headquarters and each other
- Making your branch office effective in supporting customers or remote employees
- Delivering training efficiently
- Ensuring fast recovery of operations and data after an outage
- Consolidating applications, data operations, and storage away from vertical silos
- Simplifying IT operations; saving capital and operating expenses

Achieving every one of these goals, and hundreds more, will be easier and less costly if you take full advantage of your network *and* also rethink how you view it. From the corporate data center all the way out to an office halfway around the world, your network can improve the way you serve your customers; the way you create and deploy new services or products; the way you empower employees; and the value you get from any kind of data, your production processes, inventory management and supply chain,

financial systems, and every other activity your company engages in.

Impractical? Not really. But first, you must embrace a new perspective on your network. Think of all your business processes as sitting atop a single broad architecture—an architecture of architectures, if you will—so they make use of the same network services.

The National University of Singapore (NUS) began to shift its perspective five years ago when it decided to build an integrated network and create an online portal for students, faculty, and staff to conduct their day-to-day activities. NUS wanted to create an environment where learning opportunities surrounded students 24 hours a day, everyday. It also wanted flexibility for the future. To support its online learning portal with vast capabilities that could be expanded as needs arose, NUS replaced its entire network with an integrated Cisco infrastructure (see sidebar, page 25).

SONA

Cisco has developed a framework called the Service-Oriented Network Architecture (SONA) that embodies the National University of Singapore's forward-thinking view of its network. SONA outlines how enterprises can evolve their IT infrastructure into an intelligent information network that accelerates applications and maximizes business processes and resources. The framework shows how integrated systems across a fully converged network

allow flexibility, while standardization and virtualization of resources increase efficiency.

Cisco SONA has three layers (Figure 1). The networked infrastructure layer is where all of the IT resources are interconnected across a secure and converged network foundation. This layer encompasses all places in the network: campus, branches, data center, WAN/MAN, and teleworker locations.

The interactive services layer enables efficient allocation of resources to applications and business processes delivered through the networked infrastructure. Residing here are services such as security, mobility, storage, authentication, policy management, virtualization, and segmentation. What defines these as services, rather than as applications, is that the systems providing them pervade the network with various components residing on different systems, and they are available to all users, according to Bridget Bisnette, Cisco's global director for enterprise solution partners. Security, for example, involves firewalls, Network Admission Control (NAC), intrusion detection and prevention, and much more. Some functions are router-based; others are on network appliances, but all are used collectively to provide security to users, applications, and systems enterprise-wide. "Computing, voice, identity, and storage services are others that began as applications but evolved

into services that can move into and be managed by the network,” says Bisnette.

The *applications layer* contains the business and collaborative applications that leverage efficiencies from the interactive services. These applications are also available enterprise-wide. Someone in a call center in India can pull up the same customer data on his or her desktop, from the same copy of the CRM application in the same data center, as a user at the company headquarters in Europe or North America.

Underlying the services and applications layers is the concept of *virtualization*, which goes beyond mere access to availability. For any user, applications and network services are as available as if generated in the nearest departmental or branch equipment closet, whether the user is at headquarters or halfway around the world. Given that 50 to 80 percent of employees are usually not located at headquarters, virtualization must intuitively improve productivity, emphasizes Paul McNab, vice president of marketing in the Integrated Networks Systems Engineering group at Cisco.

Virtualization supports the trend toward convergence: of voice and data networks, of services, and even of data itself—all of which must be protected through an integrated, end-to-end network with security services embedded into it. Having only one installation of a CRM, enterprise resource planning (ERP), or warehouse management program saves all the work of replicating and updating data in copies in different locations. When data is always up to the millisecond, it can be used by one application after another, and the network can ensure it flows from one application to others.

Moreover, data and applications can be located on any server or storage device that has the available capacity.

What virtualization brings to data in any enterprise, McNab says, is visibility—visibility to anyone in the enterprise or beyond who needs it. “A large retail chain may have 5,000 stores and employ 5,000 manufacturers. People or applications at the store, at headquarters, and all down the supply chain, wherever they’re located, can be enabled to see the information when a store in Los Gatos, California, sells a widget,” he explains. “That communication, especially with all links in the supply chain, is becoming crucial as companies need to control inventory tightly and cost effectively.” The task at hand, he adds, “is no longer managing the product, but managing data *about* the product.”

Making Silos Obsolete

It used to be that applications and data storage needed to be near their users, so network latency didn’t become a problem. There just wasn’t the sheer volume of applications as there is today. Now, says Greg Mayfield, senior manager in the Enterprise Solutions Marketing group at Cisco, it’s not uncommon for an enterprise to have hundreds of applications and databases in separate silos, with a lot of vacant space on servers, and as many as a thousand applications queued up to install and run.

Cisco has introduced several new products that optimize application performance. For example, latency can be cured by Cisco’s Application Velocity System (AVS), which minimizes both the number of transmissions across the WAN to use an application and their content. Cisco’s Wide Area Application Services, which

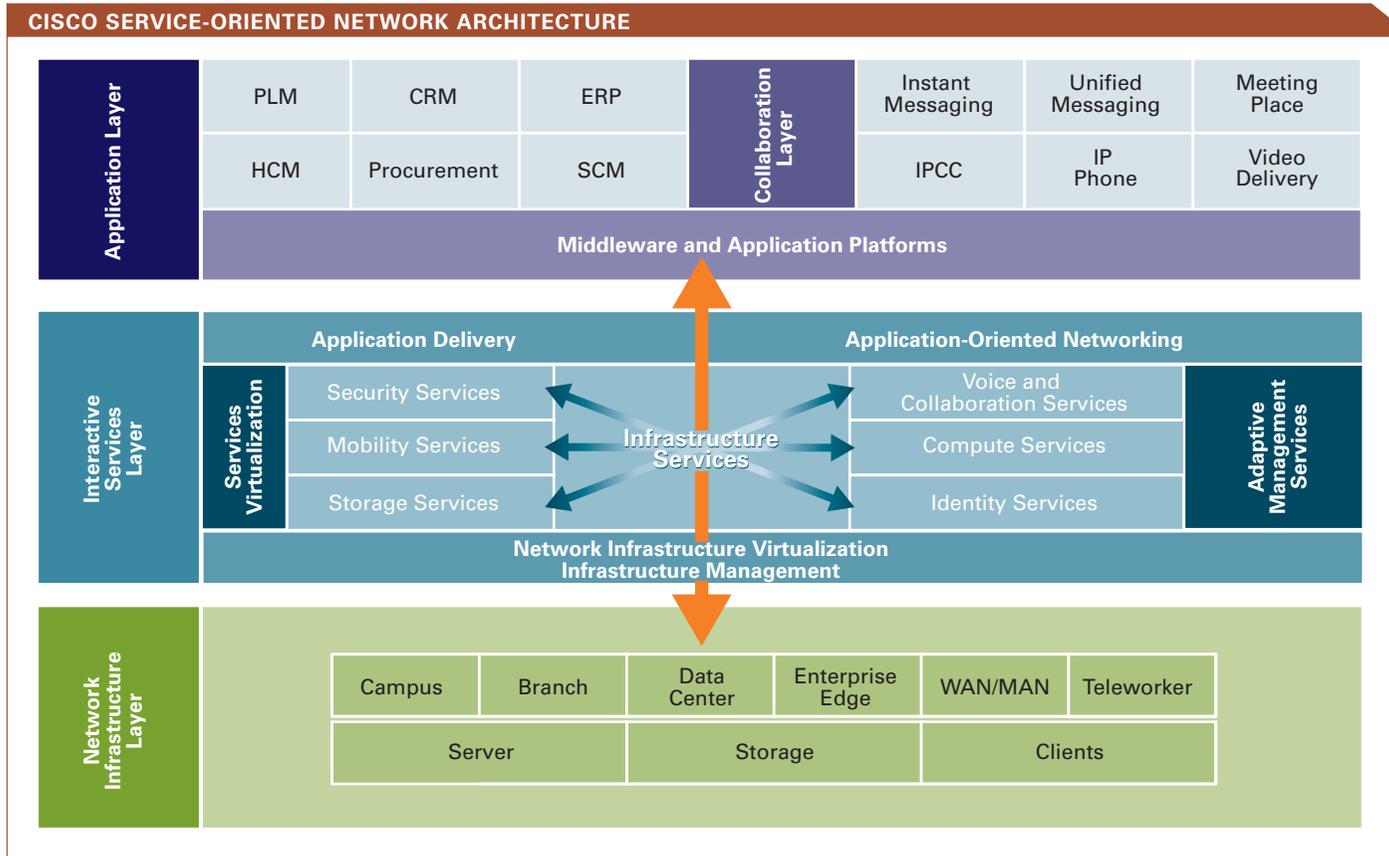


FIGURE 1 In Cisco’s three- to five-year vision for SONA, all IT resources are securely interconnected across a converged network foundation. In this new services infrastructure view, the network significantly enhances applications with its ability to deliver secure, optimized, resilient end-user services.

caches stable information locally, also minimizes WAN traffic, as do the Cisco Content Services Switch and Cisco Content Services Module, which balance request loads across multiple application servers according to policies that ensure requests go to the right server. For more on these and other products, see the *Packet* Special Report on Application Networking, Fourth Quarter 2005.

A sticking point in automating workflow processes has been the inability of applications to share data with each other because they employ many different languages and protocols. Cisco's Application-Oriented Networking (AON) products take care of that roadblock. An AON router blade functions as a universal translator for many enterprise applications, so data can move from one to another seamlessly. The AON module speaks many protocols and languages; it can read messages to see what they

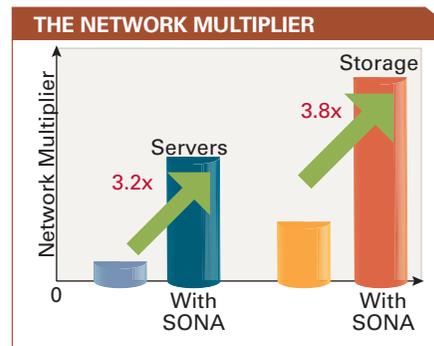


FIGURE 2 Virtualization boosts the percentage of network assets utilized, in turn boosting the effectiveness of IT spend. This chart represents the experience of Cisco's IT department with server capacity: the use of servers was 3.2 times more effective with SONA than without; the use of storage systems 3.8 times more effective.

contain and where the information should go. AON can also apply policies and priorities to messages.

Continued on page 27

The National University of Singapore: SONA in Practice

The National University of Singapore (NUS) comprises 13 faculties, 12 university-level research institutes and centers, 32,000 students, and more than 3,000 faculty members and researchers. NUS beared several goals when it decided to build a network for supporting myriad multimedia applications including e-learning, live lecture broadcasts, business warehousing, IP telephony, digital library and media archival, student administration, university admission, course registration, and GRID computing. Its network had to be scalable, robust, reliable, and achieve high performance to adapt dynamically to all demands for real-time and non-real-time, asynchronous or synchronous traffic.

NUS turned to Cisco for an integrated, converged network, and has in the five years since evolved its IT infrastructure into a highly available, intelligent network based on the Cisco SONA framework. The network distributes security, identity, wireless, and storage services across the organization. It enables all students and faculty members to obtain Webcast lectures, test scores, assignments, library materials, a messaging system, and other resources through the Integrated Virtual Learning Environment (IVLE). With campus-wide wireless service, NUS also enables students to access course materials anytime anywhere using notebook computers. With SONA, the wireless and wired experiences are the same, beginning with security. Authentication and authorization services are consistent between both networks.

"Applications such as IVLE must be built without knowing exactly what the demand will be or how many people will use it simultaneously," says Roland Yeo, network manager at NUS. "You have to base it on a scalable network and ensure that underlying network services will deliver the application as it is needed."

The network can manage very granular policies for all of its end users. For example, a student can log into the IVLE to take an online test through a secure VPN but meanwhile cannot access any other resource on campus or off. Coded access cards, required to open all doors on campus, can track individuals' movements on campus in

times of emergency, which was useful during the SARS crisis when the university needed to know who might have been exposed. The network has also enabled new, mission-critical applications such as Centralized Online Registration Systems (CORS) for module selection and allocation. "As the university advances towards a broad-based education, students will be required to enroll in modules across faculties, and a system must be established to facilitate fair, equitable, and responsible selection of modules and, at the same time, allow teaching departments to manage their resources optimally," explains Kwee-Nam Thio, manager for academic information systems at NUS. CORS has successfully allocated 1,200 modules each semester to the more than 20,000 students registering online from anywhere on or off campus.

"Our most important challenge is balancing security with openness," Yeo says. "The Cisco SONA architecture provides a framework for ensuring security, rather than ad hoc attempts at preventing a wide range of security threats." The flexibility has become crucial, he adds, given that the network is adept at fending off threats, such as denial-of-service attacks, that were unheard of when it was installed five years ago.

While the university has taken important steps towards establishing an extensive, e-enabled environment, Tommy Hor, director of the computer center at NUS, sees the following areas as important for future NUS IT capabilities: *a single-source environment*, where university reporting would all come from a central tool, creating a unified view across multiple departmental functions; *integration of voice, data, and video*, with more integrated interfaces for audio, video, and Web calls with simultaneous access to real-time databases and applications for intelligent, informed decision-making; *wireless technology devices and systems* for always-on access to resources; *personalized portals* for distributing information instead of pure e-mail dissemination.

"The Cisco SONA architecture will greatly facilitate NUS in this journey," says Hor.

An Architecture of Architectures

One size, of course, rarely fits all. The Cisco SONA framework must encompass all places across the entire enterprise. So, Cisco has created model SONA “sub-architectures” for each of these places. The campus architecture, for example, is further divided into access, distribution, and core areas. Design guides for each architecture ensure that the systems installed perform the functions needed in each part of the overall network and also work together network-wide. Take the variations for branch offices. “There are many different types of branch office. A call center is a branch office,” says Jeanne Beliveau-Dunn, director of marketing for enterprise routing and switching at Cisco.

“But most branches don’t have resident IT support, and their networks are kitted and sent out from headquarters. Cisco has tried to do as much integration of the services that might be needed as possible. The SONA branch architectures show customers a path for streamlining those offices to include security, reliability, service convergence, IP telephony, video, wide-area file sharing, content networking, QoS, application velocity, application-oriented networking, or whatever they need.” Beliveau-Dunn adds, “That’s what a big part of SONA is all about—showcasing best practices for designing any part of the network.”

The Role of Partners

The aforementioned AON products demonstrate Cisco’s role in the SONA applications layer: to facilitate their use. But the applications themselves and their integration into workflow processes are the task of Cisco partners, notes Bisnette. In fact, she says, SONA relies on Cisco partners for:

- Supplying the Cisco products associated with SONA, and providing the support services and overall IT systems management required to deliver a complete solution, from separate business applications to lifecycle network design and implementation.
- Consulting and business re-engineering as companies ready their networks for SONA, and providing managed services. In the future, the next-generation networks deployed by service providers will interface tightly with their customers’ SONA architectures, so hosted services and applications will pass transparently from one to the other.

How to Get Started

Evolving to a more intelligent, integrated network with SONA is done in phases:

Convergence and standardization of networks across the entire enterprise. Best-in-class IT organizations are designing single voice/data/video networks to handle all business communications, and standardizing on components in networks, desktops, and servers to optimize and simplify their infrastructure.

Consolidation of IT resources such as servers, which are notoriously underused when siloed and spread across the globe. Many are used to only 20 or 25 percent of their capacity.

Virtualization of IT resources, such as server clustering, where a given application or data store resides in just one or a few central places rather than in dozens or hundreds. Or virtual networking, which allows you to securely segment the network to address scalable and separately managed and billed network systems for handling multiple discreet businesses. In both cases, virtualization allows you to group systems or assignments of tasks together to maximize resources and reduce costs and overhead of management.

Automation, or deploying network-based services such as security and identification across the network so all applications can and do call on them. Enterprises also invoke application optimization

in this stage, adding capabilities such as AON, with its inter-application translation services and intelligent data management.

The key to an effective migration, says Mayfield, is thinking first about the business problem, rather than the technology. What is the business problem you want to solve? How can you best address this problem? And then, What solution fulfills the immediate need and builds a foundation for handling future needs? There are many ways to begin migrating to a SONA architecture, adds McNab. “Most people begin with the network they have, rather than undergoing a forklift upgrade. They might start with one service or application, such as file management, identity verification, or user presence. The difference is that now they’d look at how to create an identity service that works across their entire enterprise, rather than just for one department.”

The Adaptable Network

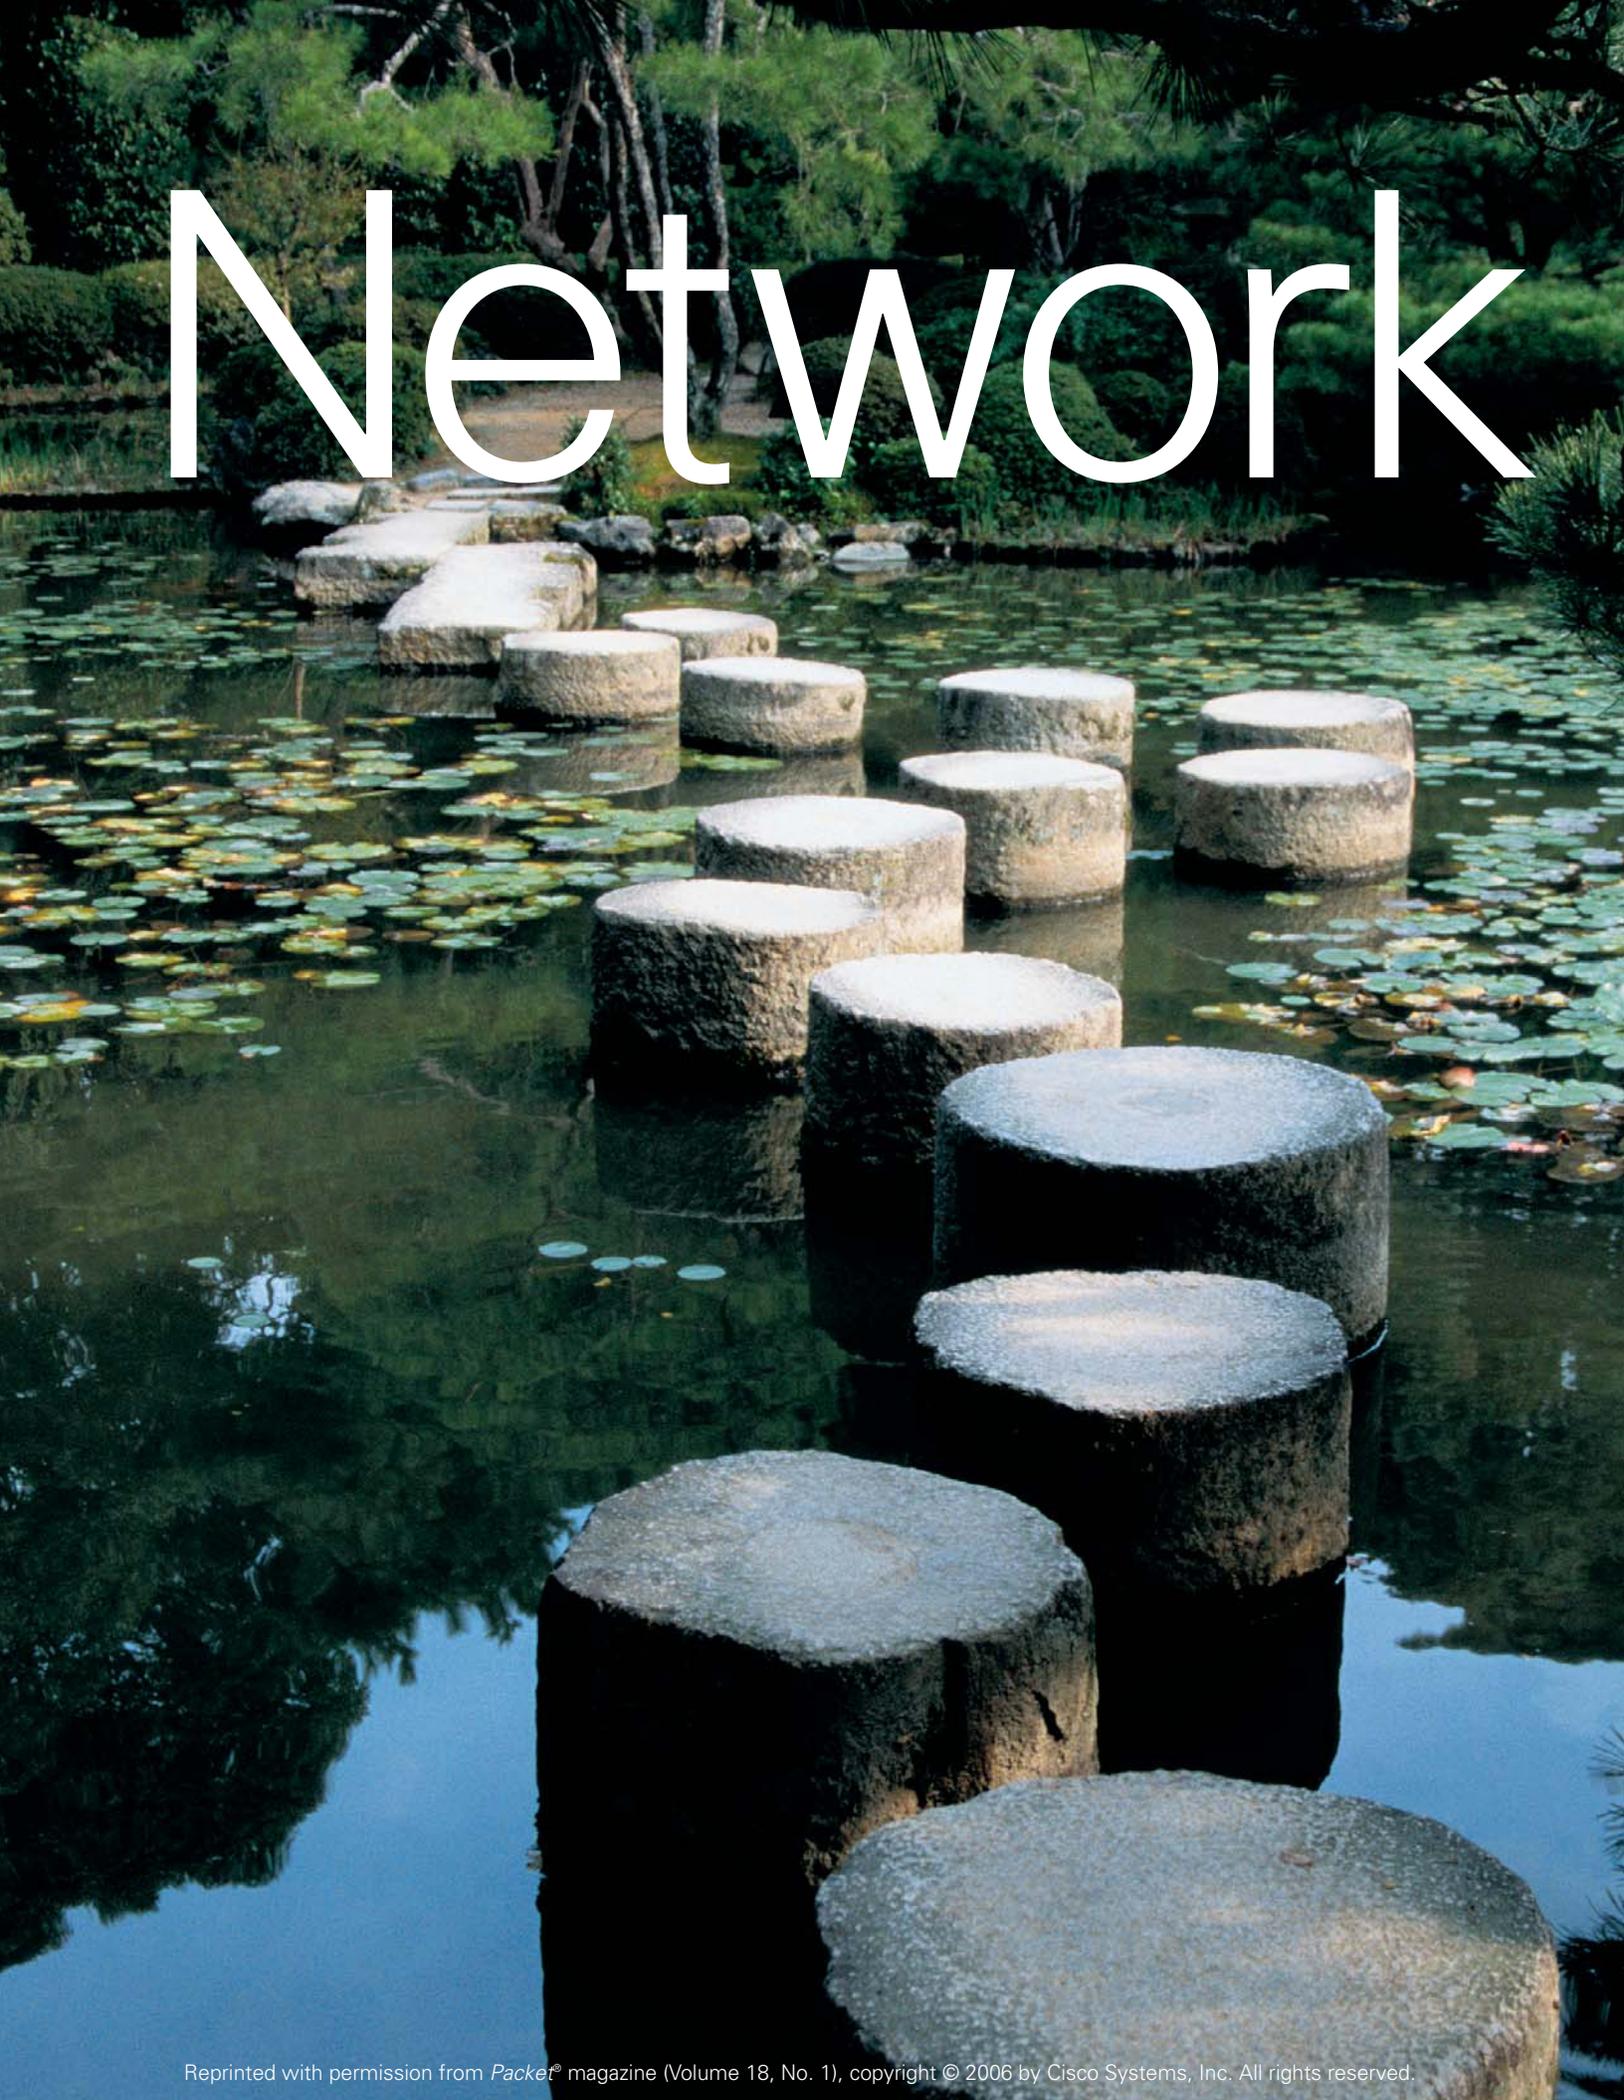
Cisco’s products already have a high degree of interoperability. You can build an integrated network based on SONA for all the sub-architectures of your enterprise from available products. In the coming years, McNab says, Cisco will work on ensuring that all the architectures and individual products interoperate from an applications perspective in providing network services. The bottom-line deliverable for SONA is an adaptive network, one that can provide a company the ability to react in real time to new business opportunities, unforeseen market changes, and customer demands. That’s exactly what Roland Yeo, network manager at NUS, believes he has. “Our Cisco infrastructure met our immediate needs, and now, five years later, we are able to add security services, deploy voice, and implement video applications simply by enabling QoS features in the network.”

The SONA-based network has changed the way NUS does business—the way it educates students and the way its employees interact with each other. Adaptable indeed. Oh, and the university is also saving US\$1 million in voice telephony costs alone. ■

FURTHER READING

- The Business Case for a Service-Oriented Network Architecture
cisco.com/go/sona
- Application Infrastructure Primer for Network Professionals
cisco.com/packet/181_6a1
- SONA Branch Architectures
cisco.com/go/branch

Network





Through the process of simplification, networks achieve new levels of resiliency, responsiveness, and lower total cost of ownership. By **Joanna Holmes**

Nirvana

t

THE CONCEPT OF NETWORK SIMPLICITY sounds paradoxical. Networks are by definition highly complex, globally dispersed, assembled through any number of technologies, and accessed through any number of media. And these networks deliver a wide array of sophisticated services. It's safe to say that no enterprise network will ever be simple, but new technologies, products, and processes enable it to be *simplified*, and the transformation can make it a much more lucrative business investment.

As network technologies have matured over the past decade, the corporate network has undergone major evolutionary phases, drastically changing the nature of its business role. As new advances such as firewalls and Web-caching technologies arrived on the market, they were sold as separate boxes, network add-ons. The network itself became more complex as new boxes from a variety of vendors were connected. This “network by committee approach,” while able to satisfy certain new discreet demands, drives operating costs higher, renders networks less reliable, and lengthens the reaction time to each new onrushing requirement.

Adding to the complexity of what was once simply the company's “data network,” IP telephony integration, video, and wireless access began to redefine the backbone as the overall communications network, raising the stakes for dependability.

Then, with the widescale economic downturn in 2001, the network integration of new technologies continued, but most companies' IT departments were understaffed in the operational

resourcing supporting these networks, even as functionality and importance was thrust upon them. Today, as the economy recovers, the challenge for most companies is that the resources allocated to their IT staffs are not keeping pace with network expansion. The result: longer hours and high stress levels—and increasingly, failure to meet intensifying business requirements.

Fortunately, as network services mature, they are progressively becoming integral components of the network equipment itself—the routers and switches—giving network managers the ability to pack more of the capabilities they need into a single box. “It’s a natural trend for features to be pulled into new equipment,” says David Willis, chief of communications at research firm Gartner. “For example, we used to add on firewalls at the network edge. Nowadays you get firewall features that are quite strong built into the same edge device that you need there anyway.”

That type of integration is a key component of the next phase of network evolution—the age of the simplified network that efficiently and effectively delivers an ever-expanding set of high-impact services. In this phase, companies can begin to streamline their networks, pulling in all the extraneous pieces and automating functions, then standardizing on single plat-

these areas with automation technologies such as NAC [Network Admission Control], IPS [intrusion protection systems], and dynamic VPNs provides huge benefits.”

It’s vital to take a step back and thoroughly assess where you are, what you have out there, and whether this network is going to serve you for the long haul,” says Beliveau-Dunn. “Even simple things like updating and standardizing IOS software releases can make a big difference.”

IT managers should develop plans to help them achieve the service-oriented networks they need. Often, the baseline network health assessment is driven by a particular requirement. It might be the need to increase security for all protected devices, or the rollout of a particular application. The impetus might be as simple as a general spending review and a cost-cutting plan—or, more specifically, the need to streamline operations for an over-taxed IT staff.

Stretched thin as the IT staff is, the network evaluation should be outsourced to a knowledgeable Cisco partner, or done through an advanced service from Cisco. A good first step is the Cisco Discovery Tool (cisco.com/go/partner-discovery). This PC

“Complexity is the enemy. If you have an old infrastructure, you can give it new functionality by introducing new devices, but over time, everything collapses under its own weight because of the diversity of the equipment.”
—David Willis, Chief of Communications Research, Gartner

forms and virtualizing network and network-connected resources. The resulting network is not only more resilient and responsive, but has a much lower cost of ownership than its more complex predecessor.

Baseline Assessment of Network Health

“It’s not that networks are getting simpler,” says Jeanne Beliveau-Dunn, a marketing director in Cisco’s Product and Technology Marketing Organization. “In fact,” she says, “by definition they’re getting extremely complex, because today they’re carrying everything—secure data, voice, video, wireless communications, and more.” Now that the network has become the central nervous system of the business there is more emphasis on things that make it resilient such as security and high availability built into every part of the network. Here, as business demands have intensified and technology has advanced, the network itself has evolved to being more service delivery platform than basic connection utility (see related story, page 22). To achieve the best results through that service-oriented network, Beliveau-Dunn explains, perform a step-by-step assessment of how you can consolidate, standardize, and integrate many of its functions to provide, for example, access control, intrusion prevention, and stateful failover—which in turn deliver better security and more reliable connection performance. “Our customers are struggling with automating policy controls and deploying multisite secure network connectivity. Developing a scalable strategy to address

application can create a record of all connected network devices, including product platforms and operating system versions. Cisco customers can get access to this application from their Cisco sales reps or channel partners. Those with advanced service contracts can also take advantage of network assessment and security audits provided by Cisco’s Customer Advocacy organization or its partners.

New functionality often drives the need for a baseline health assessment, says Willis. “But even in a stable environment, after several years’ time, it’s valuable to do an overall cost review, an overall audit [of functionality]—and in a WAN, an assessment of whether you’re optimizing your carrier services so you’re not spending more than you need to.”

“One challenge that frequently affects our customers,” says Beliveau-Dunn, “is that IT managers get hit with requirements they don’t even know if they can meet, because they don’t know what’s in their networks.” A baseline health evaluation can give them a two-to-three year outlook and help clarify the steps necessary to achieving their business goals.

“For example, we see hospitals looking at ways of locking down their networks, while maintaining or increasing network performance over many different connected devices,” Beliveau-Dunn continues. Hospitals characteristically must provide

Network Simplification Tips from Cisco IT

"In the USA, we call it 'eating our own dog food.' In Europe, it's more politely known as 'drinking our own champagne.'" Either way, says Vlada Marjanovic, senior director of the Cisco on Cisco program, it refers to the company's reliance on its own network technologies to support its mission-critical corporate network.



Cisco's core backbone supports 40,000 employees, as well as countless partners. It integrates voice, video, SAN, and wireless access for all offices and visitors, with 11 VPN points worldwide for access by some 11,000 Cisco employees.

Growing network complexity continuously raises the bar for Cisco's IT team. "As you start putting more eggs in one basket, that basket becomes increasingly important—particularly from an availability and security perspective," says Marjanovic. With a network that receives about US\$18,000 per second in online transactions, Cisco can be greatly affected by any outages. "When you start running business-critical applications on your network, it has to provide better quality and security than when you're just Web surfing."

The more you can simplify your network, the more you can secure it and increase its availability. The Cisco IT organization takes a service view of its long-term architecture and works continuously to simplify the corporate network, reducing complexity and improving resiliency

even as the network takes on increasingly more services and applications. One way the IT staff does this, says Marjanovic, is vigilant attention to "the next phase." In fact, at any given time, Cisco IT actually maintains two architectures. "We have a single architecture that lasts for three years," Marjanovic says, "and in the meantime, we are designing a new architecture.

Every new capability gets assigned to the new architecture." He refers to the venerable IT adage, "the less you touch the box, the higher your availability." The vast majority of network problems, says Marjanovic, generally results from change.

Change management is Marjanovic's number-one recommendation for network health. "If you let the engineers be the engineers," he says, "they will experiment with technology and equipment and there will be problems. Managing changes requires no CapEx or OpEx, but it can still be very difficult for IT managers."

Another recommendation from Marjanovic: Always link your business to your network. "You have to ask yourself, where is my strategic advantage?" Imagine, for example, that you manage videoconferencing services for your entire network and you meet your service-level agreements (SLAs), but you fail on your company's quarterly CEO/analyst call. You need to understand the business impact of your network, and prioritize accordingly.

extensive wireless support—and a high level of security is vital. "Everyone these days is increasing security spending, but for the most part they're not optimizing their security systems, which may be open and vulnerable." A general assessment or security audit can go a long way toward improving the network's health.

Achieving Nirvana

After assessment, the heavy lifting of the network simplification process begins. The company's IT organization must review the assessment results to ensure that all the company's mission-critical applications and the teams that rely on them have rock-solid assurances of high network availability. Mark Leary, a product and systems marketing manager for foundation technologies at Cisco, recommends that IT managers work toward building "the four pillars of simplicity": *standardization, integration, automation, and virtualization*.

Standardization

Particularly in large networks, the single most important step toward reducing network complexity is standardization. A network might have tens of thousands of connected devices, but by standardizing on a few platforms and running uniform versions of software on those platforms, businesses can reduce OpEx by

lowering training costs, decreasing the requirement for spare parts, and eliminating service gaps. Just as one US-based low-fare airline attributes its success to its business model of flying only one type of airplane, businesses can streamline their operations by standardizing on a single networking vendor and a minimal number of routing or switching platforms within their networks.

"Imagine if you were in charge of a 100-branch network," says Leary. "Now multiply those 100 branches by a handful of devices at each location—router, firewall, intrusion prevention system, voice system—and you've got a management nightmare. In contrast, you could simply have 100 Cisco 3800 Series routers, all with the same management system, the same operating system, the same hardware components—and all offering the same set of intelligent services. Suddenly it's much more manageable."

Integration

The integration aspect of network simplification involves moving features and functionality into a single box or, optimally, a single platform. Here, Cisco routers and switches serve as a foundation for a services-oriented network. "For example, most high-impact Cisco network services, such as NAC, VPN, firewall, intrusion prevention, IP telephony, wireless, and application networking, are

Continued on page 33

Network Nirvana, Continued from page 31

now offered as integrated services within our routing and switching platforms,” says Leary.

An example of an integrated platform is Cisco’s Integrated Services Router product line. Serving as a prime service-delivery platform within Cisco’s SONA, these routers deliver multiple concurrent services, while maintaining consistent performance, providing an infrastructure that enables fast, secure access to essential business applications, and readily accommodates future applications.

These routers provide a convenient platform for branch offices, helping IT staff by replacing functionality that was previously provided by other external devices. For example, IT engineers require much less training in networks that deploy a common set of multiservice networking platforms from one vendor, versus a variety of more specialized networking devices from different vendors. A study conducted by Sage Research found that Cisco Integrated Services Router users spend, on average, seven hours installing and configuring

slowdowns and failures while maximizing availability networked resources. As the single common denominator for all IT resources—servers, storage, desktops, mobile devices, applications, etc.—the network should be central to any organization’s move toward virtualization. Here, gains in user productivity, information protection, and resource availability and utilization are just now beginning to be realized.

Benefits of Simplicity

Simplicity is the model—and perhaps a good mantra—for the “Nirvana Network,” but simplicity in itself is not the goal. The real aim is a smooth-running network with designed-in resiliency, sustainable performance, service readiness, and a lower cost of operation. “You may not be able to control or standardize on everything the users are adopting for connecting to the corporate network,” says Beliveau-Dunn, “but you can design the network to provide sustainable resiliency and performance, and you can lower its cost by following the key planning principles of simplicity to assess and plan the network.”

“Once you’ve transformed your network into an integrated system, then you can find ways to automate and virtualize many of its functions so that you’re not increasing your OpEx and staffing costs as you create this Nirvana Network.”
—Jeanne Beliveau-Dunn, Cisco Product and Technology Marketing Organization

devices when adding a new site to a network. In contrast, non-users spend an average of 12 hours to add a new site. “That’s just the time savings in deployment for one site. Imagine the savings possible in maintaining the network or resolving a network problem,” says Leary.

Automation

Largely enabled by standardization and integration, a highly automated system minimizes network and service disruptions, and boosts productivity for end users and—perhaps more importantly—freeing up time for IT staff. For example, on the Cisco Catalyst 6500 Series Switch, the Embedded Event Manager (EEM) provides a method of triggering preprogrammed local actions upon detection of specific events, resulting in increased manageability, control, and resiliency. And that leads to good things for the network, says Leary. “Most IT teams today are just happy to have their networks up and running on a consistent basis. There’s no time to focus on network improvements.”

A recent Gartner study reports a 30/70 rule, where 30 percent or less of IT staff time goes to planning and proactive IT activities, and 70 percent is focused on day-to-day infrastructure operations, or what Gartner calls “keeping the lights on.” The report contends that businesses need to watch that ratio closely so that they can achieve healthy balances of proactive work versus reactive work, without neglecting operational demands. Increased automation is one solid way that businesses can bring that goal into sight.

Virtualization

As the final “pillar of simplicity,” the process of virtualizing the network enables companies to control IT costs, and avoid systems

To support the process of simplifying the network and creating more resilient and responsive services, Cisco is developing secure, high-availability components and automated systems. A recent example is Cisco IOS Software Modularity on the Catalyst 6500 Series Switch. This capability sets new standards for network availability. It reduces the complexity of the software certification and upgrade process by allowing network administrators to apply incremental patches to address time-sensitive requirements, such as critical security fixes, without impacting their network availability.

“Application adoption is happening today at twice the speed it was five years ago,” says Beliveau-Dunn. “With that acceleration, simplifying the network is more important than ever before. Network simplification is a best practice in the IT industry. It comes down to getting more thoughtful and more assessment-based about your network—and then getting ahead of the planning curve.” ■

FURTHER READING

- Cisco IT@Work website
cisco.com/go/ciscoitwork
- LAN operations white paper
cisco.com/packet/181_6b1
- WAN operations white paper
cisco.com/packet/181_6b2
- Abercrombie & Fitch case study
cisco.com/packet/181_6b3
- *Computerworld* article on State Street Corporation
cisco.com/packet/181_6b4

Cisco IPICS creates communications interoperability
by joining radio systems with IP networks. *By David Barry*

Push to Talk



TWO-WAY RADIOS, ALSO KNOWN AS push-to-talk radios, have been a steady fixture in public safety, utilities, manufacturing, recreation, and warehousing industries for more than 50 years. From public agencies to emergency operations to global businesses, many workers depend on these ubiquitous devices to enable communications among their field and mobile workforces.

Until recently, however, two-way radio systems have been isolated. Based on proprietary technologies, push-to-talk radios, which include Land Mobile Radio (LMR), cellular, and wireless LAN, have been unable to connect outside their own networks. Not only do these networks lack interoperability with other voice networks, they are also incapable of handling new communications modes such as messaging, presence, and video. Such lack of interoperability greatly limits the usefulness of these crucial communications tools—exemplified most dramatically during catastrophic events when fire, police, and local emergency workers are unable to share critical information due to radio incompatibilities.

Everywhere



This lack of communications interoperability extends to any industry where enterprises conduct business-critical voice communications on traditional communications systems, including transportation, financial services, retail, and the public sector (see sidebar, “IPICS in Action: Maher Terminals”). Large-scale replacement of these systems is disruptive and impractical. Regardless of the industry, a new network-based solution from Cisco—the Internet Protocol Interoperability and Collaboration System, or IPICS—aims to close this communications interoperability gap seamlessly and economically.

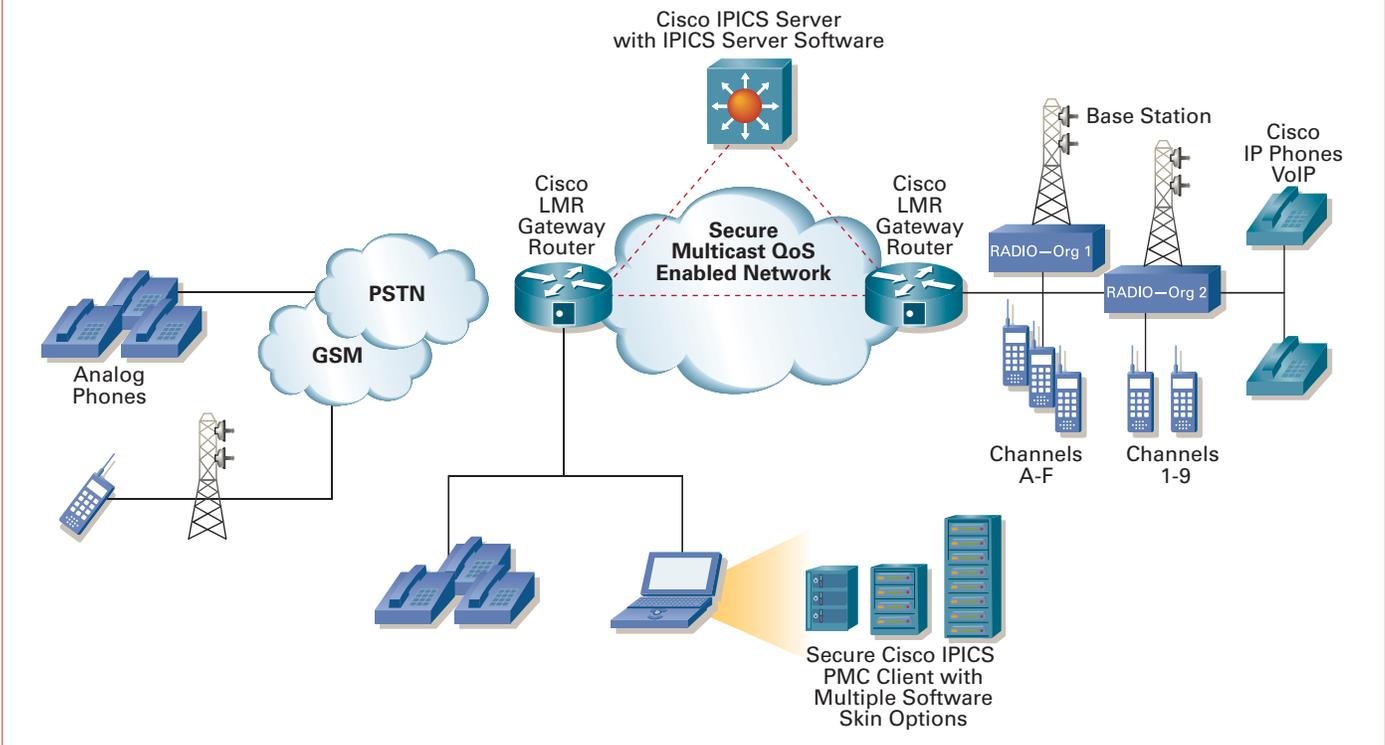
How IPICS Works

Cisco IPICS is a systems-level, network-based solution for integrating traditional communications systems with other disparate voice, video, and sensor networks. In addition to providing scalability and investment protection, Cisco IPICS takes full advantage of IP standards and the network infrastructure for greater resilience, scalability, and security, says Ken Chen, product manager in the Safety and Security Systems Business Unit at Cisco.

The Cisco intelligent network is the foundation for IPICS, providing the quality of service (QoS) and IP multicast that is critical for real-time communications. To bring radio traffic onto the IP network, companies deploy LMR gateways—Cisco Integrated Services Routers with special voice services, interface cards, and digital signal processor (DSP) functionality installed. The Cisco LMR gateways convert analog radio traffic to IP traffic, thereby extending radio’s reach to other IP-based devices while preserving the investment in traditional radio systems.

Each radio channel (or talk group as is the case with hoot-and-holler systems) is mapped to an IP multicast address. Users on IP-connected devices, such as the Cisco IPICS Push-to-Talk Management Center (PMC) client application on a PC or laptop, can also participate in these channels, enabling users that were previously blocked from communications to join the push-to-talk network.

CISCO IPICS NETWORK ARCHITECTURE



COMMUNICATIONS EVERYWHERE Cisco IPICS enables communications interoperability between radio and cell phone push-to-talk networks with a converged IP network.

Additionally, PSTN, cellular networks, and cellular push-to-talk networks (such as Sprint Nextel Push to Talk) can be seamlessly integrated into the Cisco IPICS architecture. This comprehensive voice interoperability delivers push to talk everywhere—interoperability from any push-to-talk voice device to any other push-to-talk voice device regardless of the underlying networks.

Security is also part of the Cisco IPICS solution. Users on a Cisco IPICS PMC client, for example, must log in securely and be authenticated by Secure Sockets Layer (SSL) before admission to a talk group. PMC client users can also access various channels based on user privileges or on demand via dispatch.

The first phase of the Cisco IPICS solution focuses on putting the foundation in place for basic connectivity and management between IP networks, various voice communications systems, and two-way radio networks. The focus will later shift to integrating other resources into the Cisco IPICS collaboration environment, such as standard telephones, cell phones, video feeds, remote sensors, and GPS devices.

Competing Interoperability Solutions

Two competing approaches for achieving interoperability among radio systems include using the same radio system or

using gateway devices. While using the same radio system among all organizations and agencies can be an ideal solution, it isn't practical. For public safety alone, a

IPICS in Action: Maher Terminals

Maier Terminals is one of the world's largest shipping container operators, handling about 1.2 million containers a year at its 450-acre headquarters in the Port of New York and New Jersey in Elizabeth, New Jersey. Maier has long used push-to-talk radios among its 250 employees and several hundred contract workers—including field personnel operating cranes moving 40-ton containers on and off ships and personnel on the ground coordinating the container movement.

Increasingly, Maier needed to enable communications between workers in the field and in offices where radio reception was spotty. By deploying the new Cisco IPICS, Maier has enabled instant communications between field and office personnel on IP phones (and PC-based softphones), and with others on 700 Sprint Nextel Push-to-Talk cell phones. According to Steve Rummel, vice president of data systems at Maier, "Integrating the Sprint-Nextel phones leveraged an already substantial investment in the phones." Beyond cost savings, IPICS provides immediate flexibility that the previous communications system lacked. For example, Maier can now easily patch together channels to create talk groups among any assortment of devices as required by a project—such as an incoming ship—for instant communications among that group. In the future, Rummel hopes to link key employees with US Customs and the Department of Homeland Security under a single push-to-talk system.

IPICS in Action: Incident Management

Imagine this: A fire breaks out in a high-rise building, and the fire chief needs to be updated quickly on what has transpired as he heads to the location. Using Cisco IPICS and the incident management application, the incident commander at the scene (one of the first-responder firefighters) can communicate with dispatch to orchestrate real-time communication among not just the fire chief but everyone involved in squelching the fire as fast as possible. The application's intuitive, drag-and-drop functions allow dispatchers to set virtual talk group (VTG) templates, activate VTGs to begin conferences, add or remove participants in VTG templates and active VTGs, and monitor active VTGs and events—all based on users' roles and policies and privileges assigned by an IPICS operator. The Cisco IPICS PMC client helps end users participate, through an IP network, in one or more VTGs simultaneously.

Meanwhile, the fire chief receives a page from the incident management application that contains a URL address. On screen, the chief is taken to a rich conferencing environment. Instantly, he can begin talking directly with the incident commander on location, firefighters en route on their push-to-talk radios, and other officials on their cellular phones. He also can bring up a display of the building's assets—location of stairwells, exits, etc. (IPICS will integrate GIS technology and perform database lookups to match addresses to building information). Using presence technology and GIS, the chief can also see the location of the fire and the resources in transit.

common radio communications system in the local, state, and federal governments would cost approximately US\$20 to \$40 billion, according to industry estimates. Additionally, the time to complete an infrastructure replacement and installation would be 20 years or so. Also, these radio systems typically have a lifecycle of 15 to 20 years, and in some cases they have only very recently been deployed. Any new solution, therefore, must be able to use the existing radio systems.

A second option is to use gateways to provide limited interoperability between two otherwise incompatible systems. But these devices should be considered a tactical, interim solution, says Dean Zanone, customer solutions manager in the Safety and Security Systems Business Unit at Cisco. This approach does not scale well as the number of radio devices increases. Management is especially difficult with large, incompatible radio systems. Thus, radio gateways are best used in a locally focused, limited role to ease interoperability issues in the short term.

Most significantly, analog-to-analog gateways do not take advantage of the latest technologies that greatly enhance communications. For example, these gateways do not provide a means to dynamically respond and adapt to ad hoc events and emergencies. And they can't

tap into a converged voice, video, and data IP network where rich communications extend everywhere. Ultimately, an IP-based network solution such as Cisco IPICS will render these gateway devices obsolete, says Zanone.

Radio Interoperability Using IP

An IP network-based interoperability solution for push-to-talk, LMR, and hoot-and-holler systems is preferred over the aforementioned alternatives because it connects communications paths together so that people can talk using their existing systems and devices. (Hoot-and-holler systems are hard-wired radio networks widely used in the financial industry for instant communications among stock brokers and analysts; these systems can be very costly, however, as they use separate leased line circuits to connect remote offices. With IPICS, they become part of the converged network.) In addition, IPICS is flexible and allows for dynamic linking of networks, organizations, and users on a case-by-case or emergency basis, a fundamental requirement for interoperability during a catastrophe.

Zanone offers an example of a critical event within the financial services industry. Financial analysts and brokers continually monitor worldwide events to gauge their impact on futures markets such as oil, grain, or other commodities.

If a natural disaster, such as a hurricane, develops and threatens oil platforms in a region, this information or "situational awareness" about the storm must be relayed to many people within a brokerage firm so they can reach consensus on how best to advise their clients. With Cisco IPICS, the brokerage firm can quickly bring together onto a conference call a broad group of people on different systems, including push-to-talk radios, hoot-and-holler, PSTN, cell phones, and brokers on laptops.

Real-Time Operations Management

The power and flexibility of an IP-based network approach is demonstrated by the Cisco IPICS solution that integrates an incident management application—which virtualizes resources such as users, user groups, or radio channels—across multiple networks and operational domains for dispatch or incident command (see adjacent sidebar). This IPICS solution provides dynamic orchestration of various resources for the event, and allows for graceful escalation and de-escalation as the situation unfolds or policies, roles, or responsibilities change, says Zanone. Companies and public safety departments can bring in resources on an as-needed basis and easily remove them when the incident is over. What's more, additional data such as Geographic Information Systems (GIS) and presence and database information (e.g., the location of stairwells in an individual high-rise building or surveillance camera images) can be incorporated and conveyed contextually to anyone who needs it in real time. All of these capabilities make for more efficient, collaborative incident management control.

◆ ◆ ◆

While providing an immediate tactical solution to voice interoperability, an IP-based network solution such as Cisco IPICS builds on the most widely deployed, scalable technology that will be driving innovative communications solutions in the future. Radio systems will become another application, like voice, video and data, on the IP network—and will take advantage of application convergence to achieve new, powerful capabilities. ■

FURTHER READING

- Cisco IPICS
cisco.com/go/ipics
- Cisco IPICS Deployment Options
cisco.com/packet/181_6d1

Minding the Store

A SHIFT IS TAKING place in the way businesses deploy storage networks. The reason is that these networks have come to play a significant role in helping organizations to ensure business continuance during system failures and site outages. Where data backup was once confined to a few servers, tape drives, and switches—easily controlled and secured in a single data center—the spate of recent natural disasters and other emergencies has caused most businesses to rethink this design.

To better protect themselves, more organizations are backing up data in two or more locations and using TCP/IP protocols for fast, versatile access from distributed sites. The now common use of IP-based storage technologies such as Internet Small Computer System Interface (iSCSI) and Fibre Channel over IP (FCIP) allows users to be automatically redirected to backup resources in geographically diverse locations, in the event that data should become inaccessible in a primary site. This strategy is a boon to data

availability; however, it brings with it some new considerations for the backup network.

These issues involve ensuring the performance of storage data across long-haul links, as well as the security of that data in transit and the scalability of the storage-area network (SAN) footprint. In short, IT managers are beginning to face many of the same issues with their storage networks that have confronted them in their data networks. As with data networks, increasing volumes of storage data are traversing the WAN, which introduces distance-driven delay and new security exposures into the design equation.

IT professionals should consider the following issues as they build SANs that now might reside many thousands of miles away from the sites attempting to gain access to them:

- Are there ways to offset distance-induced delay to accelerate SAN performance?



- How can storage data, now leaving the data center and transiting common data networks, be secured against eavesdropping, alteration, or theft?
- How can strong authentication and authorization of users, devices, and IT management personnel be ensured?
- Is there a way to partition access to SAN resources, much in the way that Ethernet virtual LANs (VLANs) partition access to live servers, using logical user groups, for scalability, fault isolation, and security?
- How do organizations manage a heterogeneous SAN environment?

A combination of industry standards and interfaces, along with features in vendor equipment, help IT managers ensure performance acceleration, security, and management in the SAN for improving business continuance.

Accelerating the SAN

Acceleration in the SAN is conceptually similar to application performance acceleration in data networks. Instead of proxying application protocols for local acknowledgement, however, it is SAN control protocols that are locally acknowledged to reduce the number of round trips and associated time it takes to move blocks of data from point A to point B. For example, the SCSI protocol requires two roundtrips of acknowledgements for every write issued. When local devices request data from a distant SAN, those devices can be acknowledged locally to reduce WAN-induced latency.

There are two types of acceleration, or local acknowledgements, prominently in use: *write acceleration* and *tape acceleration*. Their use depends on which type of storage media is to be accessed. Both are supported in the Cisco MDS 9000 Family of multiprotocol storage switches. These devices simultaneously support Fibre Channel, FCIP, iSCSI, and mainframe Fibre Connection (FICON) connections. They switch Fibre Channel data among

Storage Security and Management Lexicon



FC-FS-2 Fibre Channel-Framing and Signaling-2. Specification defined by the T11 Technical Committee of the InterNational Committee for Information Technology Standards (INCITS) to transmit SCSI command, data, and status information between a SCSI initiator and a SCSI target.

FC-GS-3 Fibre Channel-Generic Services-3. Inband management standard defined by the INCITS T11 Technical Committee for transferring status and configuration information, including VSAN information, among Fibre Channel devices.

FC-SP Fibre Channel-Security Protocols. Draft standard by the INCITS T11 Technical Committee for securing Fibre Channel storage data in transit using data encryption, cryptographically secure key exchange, and device authentication. FC-SP is supported by a variety of SAN switch vendors and by all major host bus adapter vendors. Targeted for approval in March 2006.

SMI-S Storage Management Initiative Specification. A standard developed by the Storage Network Industry Association (SNIA) that is intended to facilitate the management of storage devices from multiple vendors in SANs, allowing a single management application to handle multiple tasks that would otherwise require multiple applications.

like ports and also encapsulate Fibre Channel data in IP and send it out an Ethernet interface for IP transit.

Both types of acceleration boost performance.

Write Acceleration. This acknowledgement enhancement used for disk-to-disk and host-to-disk transmissions reduces SCSI's two roundtrips to one, thereby doubling performance. In this case, an acknowledgement of the receipt of intact data is sent after the second roundtrip of the handshaking process.

Tape Acceleration. This acknowledgement enhancement builds on write acceleration, described above, to accelerate moving storage data from a media server to a tape drive. Performance is first enhanced by local acknowledgements that reduce the roundtrip WAN acknowledgements by half. A configurable file mark mechanism in tape systems, however, also allows the IT administrator to set the long-distance acknowledgement mechanism to take

place after a desired number of data blocks rather than after every single one to reduce the number of required acknowledgements even further. Data is buffered between acknowledgements. Because tape media performance is notoriously sluggish, reducing the acknowledgements required to every X number of data blocks buys significant performance benefits. For example, in some Cisco customer tests, a 100-millisecond acknowledgement has had just a 15 percent impact on performance.

Compression

As in data networks, compression can also be used to increase the effective WAN bandwidth, avoid congestion, and improve performance. Cisco storage switches support different data compression algorithms, selectable depending on configuration, that allow compression ratios as high as 30:1, depending on data compressibility of the data block. Typical ratios for common database traffic are 2:1 to 3:1.

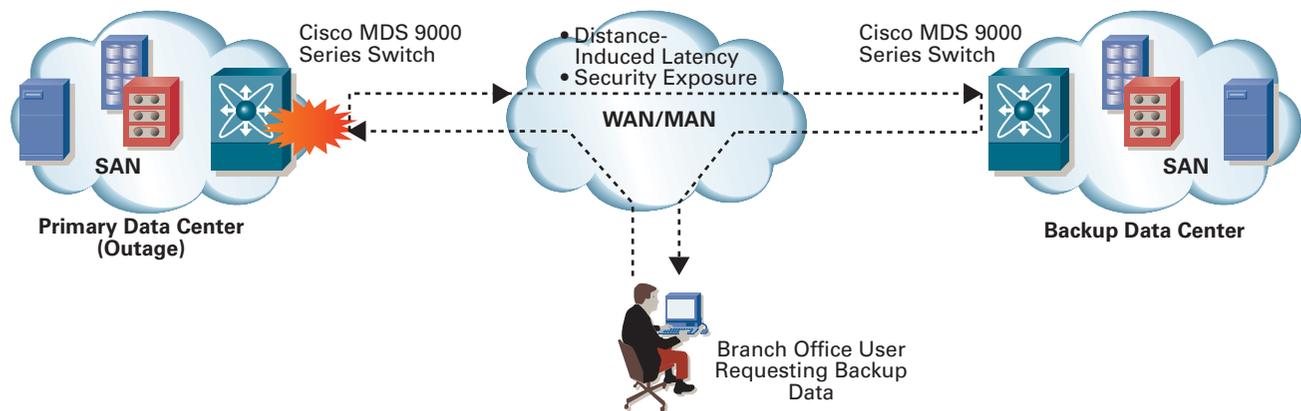
Securing Storage Data

Typical network security concerns are now beginning to apply to SANs. SANs



TOM NOSELLA, CCIE
No. 1395, is Director of Engineering in Cisco's Data Center, Switching, and Security Technology Group. He can be reached at tnosella@cisco.com.

WAN-INDUCED BUSINESS CONTINUANCE CHALLENGES



IN THE CLEAR Considerable distances between data center SANs, and between branch offices and SANs, drives the need for SAN performance acceleration using local SAN protocol acknowledgements. Traditional security mechanisms should also be added to SAN data carried over the WAN.

have generally been small and localized within a single data center. Now, however, long-haul networks involving several service provider infrastructures might be used to move critical storage data that may never before have left a data center except on a piece of physical media in a truck.

The result of this shift in the treatment of storage data is the need to apply the security features prevalent in IP network elements to the Fibre Channel environment. This involves protecting data in transit, securing against unauthorized user and device access, and guarding against malicious management misconfiguration. In a network of storage switches such as the Cisco MDS 9000, also called a *storage fabric*, this involves encryption, authentication, and securing the SAN management infrastructure.

Encryption. Data encryption is important for preventing intruders from viewing or modifying confidential information. Cisco storage switches use the IPsec protocol to help ensure confidentiality and data integrity of storage data in transit. Cisco MDS 9000 multiprotocol SAN switches, for example, include integrated hardware-based IPsec encryption/decryption supporting Advanced Encryption Standard (AES), Data Encryption Standard (DES), and Triple Data Encryption Standard (3DES) algorithms for iSCSI and FCIP storage traffic.

Authentication and Authorization. These functions are now necessary to avoid accidental corruption and malicious attacks on SAN data. They enable only certified users and devices to connect to stored data. Storage switch-to-switch authentication and authentication of other switches connecting to a Cisco storage fabric use the cryptographically secure key exchange and device authentication components of the draft Fibre Channel-Security Protocols (FC-SP) standard of the American National Standards Institute's InterNational Committee for Information Technology Standards (INCITS) T11 Technical Committee (see sidebar, "Storage Security and Management Lexicon," page 40). Organizations can authenticate users and devices locally in the storage switch, reducing latency, or remotely through centralized authentication, authorization, and accounting (AAA) servers.

Secured Management Infrastructure. The data center management functions of network and storage devices must also be secured to thwart unauthorized access. Malicious users with access to the console of a networked storage device can easily alter configuration. As with other Cisco network elements, Cisco MDS 9000 switches provide secured management functions, including Secure Sockets Layer (SSL) and Secure Shell (SSH) Protocol Version 2, which secure remote access using authentication and encryption. SSHv2 can be used in conjunction with backend user authentication protocols such as TACACS+ and RADIUS that

may already be in place in the organization. In this case, the storage switch acts as a client to the back-end AAA servers running these protocols.

Finally, Simple Network Management Protocol version 3 (SNMPv3) support provides authentication and authorization services for accessing SNMP management information bases (MIBs).

VSANs for Scale and Fault Isolation

A well-planned virtual SAN (VSAN) architecture reduces the total number of SANs (or fabrics) deployed, while enabling businesses to separate their backup, recovery, and remote data mirroring domains from application-specific SANs.

Cisco technology was chosen by the INCITS T11 Technical Committee last year as a standard for VSANs. VSANs allow network administrators to segment a single, physical SAN fabric into many logical, completely independent SANs. As with VLANs in an Ethernet data network, this approach enables the creation of separate SAN domains without having to build out multiple separate and costly physical infrastructures.

Continued on page 50

FURTHER READING

- Cisco MDS 9000 Family Fabric Management Solutions Guide cisco.com/packet/181_6c1
- Storage Networking Industry Association snia.org/home

Primary Wireless

At Intel's Jones Farm campus, employees will use a wireless LAN as the primary access method for data, voice, and video.

By Rhonda Raider

When Intel employees walk through the doors of the Jones Farm campus, near Portland, Oregon, their laptops, which use Intel Centrino mobile technology, automatically detect the best available RF signal and log onto the wireless network. Employees open their laptops and are immediately connected.

With nearly 6,000 employees, the Intel Jones Farm campus is among the first in the world to adopt wireless as its primary access method, successfully addressing issues and challenges such as better performance, system management (RF coverage), increased security (RF interference and rogue detection), client roaming, and quality of service (QoS) for voice over wireless. The Intel IT group accomplished this engineering feat using some of the features that will soon be available in the Cisco Business Class Wireless Suite, which combines the Cisco Unified Wireless Network and Intel Centrino notebooks with Intel Centrino mobile technology. The company expects that when deployment is complete later this year, 75 percent of campus residents will use primary wireless exclusively.

The Ascent of Wireless

In the three years since Intel deployed its first WLAN, usage has skyrocketed—from 1,500 mostly US users in 2002, to 55,000 global users in 200 locations at the end of 2005. “When we first introduced wireless networking, employees relied primarily on the wired LAN, using the wireless LAN only for secondary access such as connectivity in conference rooms,” says Sylvia Stump, Intel's IT wireless program manager. But Intel's IT surveys revealed that nearly all Intel employees using laptops actually preferred wireless LAN connectivity. “Our employees enjoy the freedom of working anywhere—conference rooms, common areas, cafeterias—just as productively as if they were at their desks,” Stump says.

The idea of using wireless for primary access arose when the Intel IT wireless team began investigating ways to increase employee productivity and reduce network costs. In 2004, Intel IT was managing three separate networks—LAN, WLAN, and telephony—which essentially tripled operational costs. And as the number of wireless access points surpassed 5,000, the WLAN management burden was becoming overwhelming. “With our original wireless architecture, we managed each access point individually,”



WIDESPREAD ADOPTION This year 75 percent of the employees at Intel's Jones Farm campus will be using wireless on the job as their primary means of connectivity.

Stump explains. “As the WLAN continued to grow in size and complexity, we realized we would need centralized management and more automation to prevent a dramatic rise in operational costs and staffing requirements.”

Intel IT decided to build a next-generation network that would both continue to increase wireless network access and simplify network management. “By integrating the LAN, WLAN, and telephony networks, we could put together a more robust, stable, and efficient architecture that delivers voice, video, and data,” says Yossi Bar-El, a wireless LAN engineering manager for Intel. Using wireless as the primary access method will reduce capital expenses for cabling while satisfying employees' preference for mobility.

Jones Farm, where the Intel Centrino Group is headquartered, became the model for primary wireless LAN because the majority of employees is highly mobile and runs business applications that are better suited for wireless access.

Cisco Unified Wireless Network

Intel is developing a primary wireless solution with the Cisco Unified Wireless Network. The strategy is to build a unified wired and wireless solution to cost-effectively provide bandwidth for voice, video, and data, as well as the QoS needed to ensure call clarity. Intel's Cisco Unified Wireless Network architecture comprises Cisco Aironet 1240 AG Series

lightweight access points, Cisco Wireless LAN Controllers, and the Cisco Wireless Control System (WCS) for simplified monitoring and management. The Cisco Aironet access points communicate with the Cisco Wireless LAN Controllers using Lightweight Access Point Protocol (LWAPP, see sidebar, page 45). The underlying network is based on Cisco Catalyst 6500 Series switches at the distribution layer, and Cisco Catalyst 3550 and 4500 series switches at the access layer.

Intel Centrino mobile technology enables laptops to automatically detect and connect to the network with the best bandwidth available—for example, IEEE 802.11g at home, and 802.11a at work. “The client performs all the engineering work needed to associate with the network, so the user doesn’t need to do anything but log into their notebook,” says Bar-El.

Intel is currently midway through a three-phase deployment for primary wireless access:

- Phase 1: Implement the wireless architecture in one building, measuring wired and primary wireless networks to ensure comparable performance within a control group of 200 employees. Intel completed this phase in mid-2005, and is currently analyzing the results.
- Phase 2: Extend primary wireless data access to the entire campus. Completion is planned for mid-2006.
- Phase 3: Add voice over wireless capabilities with IP telephony softphones, Wi-Fi phones, and dual mode (Wi-Fi-cellular) devices. Completion is planned for the end of 2006.

Planning Bandwidth Capacity

Intel selected the 802.11a architecture because it provides the most bandwidth in actual practice. “802.11a and 802.11g both provide theoretical bandwidth of 54 Mbit/s, but 802.11a provides higher actual bandwidth because it has 12 channels compared to three for 802.11g,” says Bar-El. At Intel, for example, 802.11a provides 24 Mbit/s to 54 Mbit/s per user, depending on the user’s distance from the access point. Each access point contains both types of radios, and the 802.11b/g radios provide redundancy

and enable connectivity with older laptops that lack 802.11a support.

To plan capacity, Intel is currently fine-tuning its estimates of how much bandwidth individual users need for voice, video, and data. “We’re looking for the best balance between access point density and costs: the shorter the distance between a user and an access point, the greater the available bandwidth,” says Bar-El.

Centralized Management

Historically, Intel IT has manually configured, installed, and managed its thousands of access points. “But manual configuration introduces the risk of error because technicians might interpret our design standards differently,” says Stump, who adds that she could detect a difference in WLAN performance when she worked at different Intel sites.

The Cisco Unified Wireless Network ensures consistency through centralized management. The majority of network intelligence now resides in the wireless LAN controllers. The Cisco Aironet 1240 AG lightweight access points provide RF services that support dynamic RF monitoring, IDS detection, and location tracking. The Cisco WCS sets up configuration templates for all wireless LAN controllers, eliminating the human error that can occur when technicians manually enter settings. Newly installed access points have their RF channel and transmit power settings automatically configured by the wireless LAN controllers based on the Cisco WCS configuration templates. The system monitors coverage and adjusts configuration as needed to help ensure consistent coverage and performance. And if one access point becomes unavailable, the WLAN self-heals: Neighboring access points are configured to compensate for the unavailable access point and then resume their normal configuration when the faulty access point is back online.

“From support and engineering perspectives, the Cisco Unified Wireless Network is very manageable and easy to support,” says Stump. “Instead of dedicating an access point to monitoring, as we used to do, we have one system that manages and inventories all solution components: how they’re doing, what they’re doing, the levels of roaming, and types of users connecting.”

Security

The principal security concern for the Intel Jones Farm campus is RF interference and loss of employee network access, according to Bar-El. The 802.11a/g redundant radios and the Cisco Aironet 1240 AG access points, which dynamically select RF channels with transmit power, help to ensure business continuity. If RF interference occurs in the 802.11a network, employees’ laptops automatically reconnect to a different channel or the 802.11b/g network.

Talk About It

Want to share your expertise on wireless LANs with your peers? Get answers to your questions from Cisco experts? Join the Networking Professionals Connection discussion at cisco.com/discuss/wlangeneral.

LWAPP: Enabler for Centralized Intelligence in Wireless Networks

The LWAPP is a draft IETF standard. Authored initially by Airespace (acquired by Cisco in March 2005) and NTT DoCoMo, LWAPP standardizes the communications protocol between access points and WLAN systems such as controllers, switches, and routers. Its goals are to:

- Reduce the amount of processing within access points, freeing up their computing resources to focus exclusively on wireless access instead of filtering and policy enforcement
- Enable centralized traffic handling, authentication, encryption, and policy enforcement for an entire WLAN system

- Provide a generic encapsulation and transport mechanism for multivendor access point interoperability, using either a Layer 2 infrastructure or an IP routed network

The LWAPP specification accomplishes these goals by defining:

- Access point device discovery, information exchange, and configuration
- Access point certification and software control
- Packet encapsulation, fragmentation, and formatting
- Communications control and management between access points and wireless controllers

Another of Intel IT's tactics to combat denial-of-service attacks is using LWAPP to identify the attack source.

Intel's current secondary access WLAN uses VPN technology to encrypt and protect data in transit. However, VPNs also added an extra layer of infrastructure to manage, another point of failure, and the potential for bottlenecks. Today, Intel has eliminated the need for VPNs at the Jones Farm campus by using Wi-Fi Protected Access (WPA) 2 and the Advanced Encryption Standard (AES), part of the 802.11i encryption capabilities in the Cisco Unified Wireless Network.

And to authenticate users as part of network admissions control, Intel takes advantage of the 802.1X authentication in the Cisco Unified Wireless Network to communicate with a centralized RADIUS server like the Cisco Secure Access Control Server (ACS).

Quality of Service

"QoS has become essential now that the WLAN is used for primary access, because wireless bandwidth limited," says Bar-El. "And when we begin providing voice over wireless, QoS is indispensable because time-sensitive voice traffic needs to receive priority over data traffic."

To provide QoS from the desktop client to the access point, Intel uses Wi-Fi Multimedia (WMM), a subset of the 802.11e QoS standard defined by the Wi-Fi Alliance. WMM is built into the Cisco Unified Wireless Network. Intel Centrino mobile technology incorporates Cisco Compatible Extensions that take full advantage of the QoS capabilities by using Differentiated Services Code Point (DSCP) marking on packets based on transport layer filtering translating

to wireless mobile markings of the Cisco Unified Wireless Network architecture, as well as its roaming and security features.

Roaming Between Network Segments and Buildings

Roaming capability ensures that an employee's voice or data connection is not interrupted when the employee crosses subnet boundaries, as happens when he or she walks across campus while talking on a wireless Cisco IP Phone 7920, for example.

Intel IT enabled roaming via the WLAN controllers use of the LWAPP. LWAPP enables the client to maintain its IP address as it crosses subnet boundaries in a single RF domain. The client is not aware that it's changing subnets. Instead, the back-end controllers manage the traffic among themselves, ensuring that traffic reaches its destination no matter where the client has moved within the network.

Continued on page 73

FURTHER READING

- Cisco and Intel Alliance
ciscointelalliance.com
- Cisco and Intel Wireless and Mobility solutions
cisco.com/packet/181_7a1
- Cisco Unified Wireless Network
cisco.com/packet/181_7a2
- Lightweight Access Point Protocol white paper
cisco.com/packet/181_7a3

DMVPN Deployment

Large British broadcaster moves from a mixed-media architecture to MPLS with the help of mGRE technology.

By Robert Thompson, Andy Murkin, and Tim Taverner

ITV plc, Britain's largest commercial broadcaster, needed to replace its 140 Mbit/s switched, video-only network. ITV needed a network with flexibility at the cornerstone—a design that could react to and evolve with the changing needs of ITV's broadcast lifecycle. In particular, the design needed to accommodate ITV's plan to migrate, over time, many of its video-centric operations into a data-centric configuration by exploiting the emerging standards for file-based video encapsulation, such as MXF (Media eXchange Format).

With the maturation of MXF and similar standards, coupled with IP class of service extensions, it was

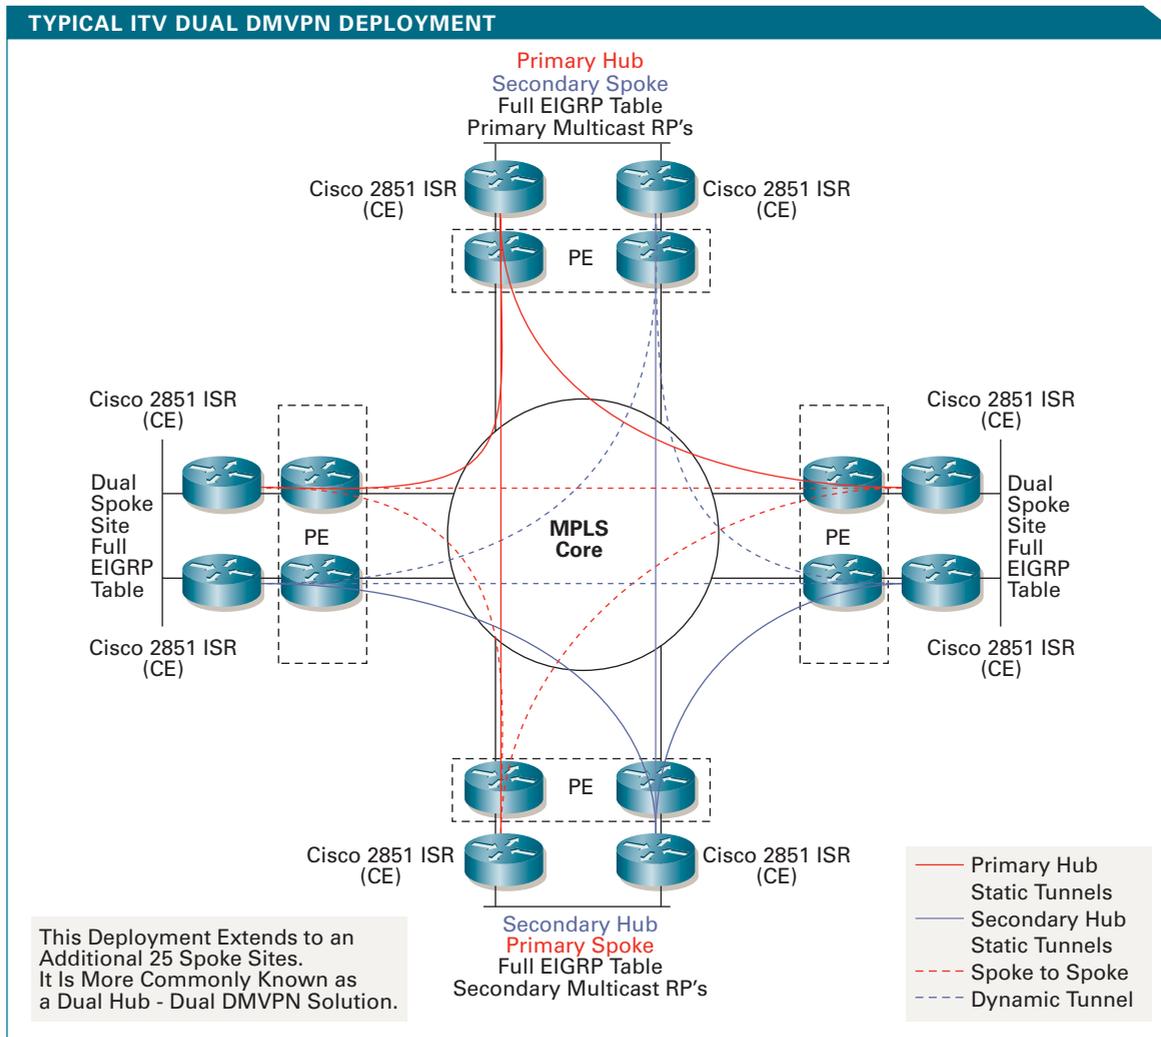
reasonable to assume that convergence for connectivity needs would be delivered over an IP-based topology. This new topology, which ITV calls its Video Contribution Network, entailed building a Multiprotocol Label Switching (MPLS) network to deliver studio-to-studio video connections and transmission feeds from multiple regional studios to central transmission sites in London and Leeds in the UK.

So, how did ITV do it?

Why Dynamic Multipoint GRE Tunnels

ITV initially had a multitude of complex networks

DESIGN FOR TV OF THE FUTURE ITV's Video Contribution Network is based on core and access trunks with capacities from 45 Mbit/s to n x 155 Mbit/s. Serial Data Interface (SDI) video at 270 Mbit/s is encoded using MPEG2. The resulting DVB-ASI transport streams are converted to IP and carried across the MPLS network. The MPLS cloud allows every node to communicate with every other node in either a point-to-point or point-to-multipoint configuration. The bandwidth on the MPLS cloud is managed by a web-based scheduling system available to ITV users.



joined together to form a resilient, redundant data network. With ATM, Frame Relay, and point-to-point leased lines, the routing and redistribution between different protocols was a major hurdle ITV needed to clear in migrating to MPLS.

To implement MPLS with little or no impact on availability, ITV introduced networks individually on a site-by-site basis, which also required the introduction of MPLS customer edge (CE) routers. In introducing the CE routers, ITV faced a huge routing challenge: how to handle redistribution and the possible introduction of Border Gateway Protocol (BGP) or Open Shortest Path First (OSPF) on the ITV core network. Redistributing different protocols between each other (mutual redistribution) was not an option in the network because of the many redundant and backdoor links to each point of presence (POP) and local office (hundreds of links and redundant paths). Another challenge: The MPLS core network did not contain only Cisco gear and, therefore, ITV's use of Enhanced Interior Gateway Routing Protocol (EIGRP), from provider edge (PE) to CE, was not a viable option.

Scalability and cost were the main criteria in selecting not only the CE hardware but the suitable means to interconnect multiple sites. Because ITV's core network was made up of a large number of sites, it was not financially viable to deploy large, high-end CE routers. In addition to being scalable and cost-efficient, ITV needed a network that could:

Move data, voice, and video traffic simultaneously; quality of service (QoS) was imperative.

Provide an "unwritten service-level agreement (SLA)" to end users, who needed to operate with the same, or better, level of service yielded by their legacy network.

Interconnect a large number of sites at all times; a full mesh design was needed.

Provide a level of transparency to traffic entering the MPLS core; the CE devices needed to be the policing and shaping point of the voice, video, and data traffic prior to core ingress.

Ensure that ITV's data network would continue to operate as usual during the migration, and be highly available upon deployment. We needed a phased

configuration, with changes and downtime considered at each step of the install to limit downtime to already-migrated sites. Failover and resilience mechanisms, e.g., Next Hop Resolution Protocol (NHRP) tunnel failover and Hot Standby Router Protocol (HSRP), also needed to be factored in.

Be as standardized as possible (from a configuration and hardware perspective) to facilitate later adds, moves, and changes.

What's more, ITV needed CE routers with a high level of performance to ensure extremely low latency on the voice and data packets and deliver the horsepower to effectively police and shape all the traffic without disabling the CE.

After research and testing, we chose a derivation of Cisco Dynamic Multipoint VPN (DMVPN) tunnel technology—namely, multipoint Generic Routing Encapsulation (mGRE) tunnels with QoS, failover (via NHRP), HSRP, etc. This solution would give us the ability to build the required full mesh between sites while retaining scalability and ease of configuration.

For the testing phase, ITV chose Cisco 2651XM multiservice routers for the CE devices. No additional interfaces were needed, because the PE routers were physically sitting in ITV's data communications rooms. The CE to PE routers were interconnected using a standard Category 6 crossover cable.

In subsequent testing and deployment, the final CE hardware solution—Cisco 2851 Integrated Services Routers—were chosen after the first deployment of the CE 2651XM routers had already taken place (see "The Implementation" section for more on how this developed).

The Test Network

The test network was built in the modern laboratories of BTSkynet. The core MPLS and PE network was designed and built by BTSkynet and integrated into the PE-CE network with Cisco 2651XM multiservice routers. The core was simulated with two Cisco Catalyst 6500 Series switches, running Supervisor Engine 720-3B technology and native IOS with MPLS support. A combination of line cards was used to interconnect to PE devices. The logical routing protocol used between the core and PE devices was OSPF. Multiprotocol BGP (MBGP) was used between the PE devices.

In October 2004, at the time we were developing and testing ITV's network, DMVPN was a relatively new technology from Cisco. To use this technology, we had to build it and understand it. ITV, BTSkynet, and Cisco worked together to build the test network as a proving ground for CE configurations based on IP Security (IPSec) DMVPN deployments that had been developed for other projects involving voice and data transmission across multipoint VPNs.

Continued on page 51

ROBERT THOMPSON, CCIE No. 10302, is a consulting engineer at BTSkynet, a Cisco Certified Gold Partner. He has more than 10 years experience in communications.

ANDY MURKIN is a senior network engineer at ITV plc. He has more than 15 years experience in communications and originally trialed DMVPN as a possible ITV solution.

TIM TAVERNER is a senior systems engineer with the UK and Ireland Channels Network Integrator team at Cisco. He has more than 18 years of experience in networking and telecommunications.

SANs, Continued from page 41

Cisco MDS 9000 devices can create up to 256 isolated VSAN topologies (the hardware supports expansion up to 4096 within the same physical infrastructure). This allows administrators to use simple zoning to restrict access and traffic flow among devices by securing access at the edge. Businesses can segregate even a single storage switch into multiple virtual environments, or domains. They can completely separate different VSANs to help ensure that fabric instability or a device outage is isolated within a single VSAN and does not cause a fabric-wide disruption.

Managing Diversity

As storage network environments continue to grow, organizations are deploying

Talk About It

Want to share your expertise on storage technology with your peers? Get answers to your questions from Cisco experts? Join the Networking Professionals Connection storage discussion at cisco.com/discuss/storage.

storage solutions using equipment from multiple vendors, which each arrive with their own separate SAN management program. Administrators require a way to effectively manage the heterogeneous storage environment in a way that ensures maximum performance and cost-effectiveness.

The Cisco Fabric Manager for SANs lets administrators view and manage the heterogeneous fabric as a collection of devices, recreating the entire topology and representing it as a customizable map. Any device in the fabric that supports the INCITS T11 Fibre Channel-Generic Services-3 (FC-GS-3) standard for in-band management can be discovered and mapped as part of the topology. A topology window displays the discovered devices for customization and navigation, while an inventory window displays a tree-like structure of both physical and virtual elements. Yet another window displays the tools administrators can use to configure, monitor, and troubleshoot devices.

Cisco Fabric Manager also supports open interfaces with access to raw performance

and configuration information within switches that can be used by third-party management applications. Support for the Storage Networking Industry Association's Storage Management Initiative Specification (SMI-S), for example (see sidebar, page 40), enables element management across multiple vendors' SAN management products.

◆ ◆ ◆

Organizations tackling business continuance build out their SANs, they are finding themselves face to face with many of the WAN-centric performance, security, and management issues that have confronted them in their data networks. Because increasing volumes of storage data are traversing the WAN, distance-driven delay and new security exposures are rearing their heads. Enterprises should look to support for SAN acceleration techniques, multifaceted security support, and support for industry-standard management interfaces and capabilities to ensure that their SANs perform well and remain secure, cost-effective, and manageable.”■

DMVPN Deployment, Continued from page 49

Different releases of Cisco IOS Software were tested during this phase, because IOS had to support the EIGRP extensions that allowed EIGRP to work across the mGRE tunnels. In the early testing phase, the IPsec features of DMVPN were extracted from ITV's solution. Because ITV has a private, wholly-owned network, the additional overhead of IPsec headers was deemed unnecessary.

QoS was tested to a degree in the lab, and multicast free-ware utilities were used to simulate multicast traffic across the DMVPN cloud. The end solution had to support voice, video, and data streams across the available bandwidth without end users experiencing any difference from their current mixed-media WAN deployment.

The Implementation

The network implementation was conducted in phases. ITV Meridian, the ITV franchise holder for the south and southeast of England, was relocating

to a new purpose-built digital television facility based in Whiteley, Hampshire, UK. Along with its regional offices at Maidstone, Newbury, Brighton, and the London transmission center, this provided an ideal testbed for the DMVPN design and hardware. Rolling out the Cisco 2651XM multiservice routers, pointing the spoke routers to hub routers in London and Whiteley, proved to be the best test we could have.

In every network design, something surfaces that the network engineer hadn't counted on during the initial planning stage. In ITV's case, very high resolution compressed video images were being sent by FTP in nonreal-time from one end of the news network to the other at any time of day—which basically overwhelmed the Cisco 2651XM multiservice routers. Now, ITV had to find an alternative CE router.

With additional testing in the lab network at BTskynet, the then newly-available Cisco 2851 Integrated Services Router was put through its paces—and passed. BTskynet shipped the Cisco 2851 Integrated Services Routers directly to the Whiteley office, and they were deployed within hours.

The fourth-generation architecture of the Cisco Integrated Services Router (integrated advanced services at line rate) along with its built-in VPN and QoS acceleration features solved the performance problems—so well, in fact, that ITV and BTskynet ran a battery of tests to ensure the validity of the vastly improved performance data! As the new routers came online, the network performance increase was so dramatic that the video tests were conducted again and again, with added clips to prove that no caching on any end system was taking place to fool the statistics we were seeing.

The Cisco 2851 Integrated Services Routers, running their initially released (non-GD) IOS code, saved the implementation. This model was then used as a template to roll out the main network, which consisted of an additional 20 sites, each with two Cisco 2851 Integrated Services Routers.

In addition to the impressive performance of the Cisco Integrated Services Routers, throughout the testing and implementation phases, we found the DMVPN solution to be a configuration engineer's lifesaver. DMVPN really is a “build once, use many” technology that saved ITV time and effort throughout the project lifecycle. ■

Implementation Tale: Resolving a Performance Issue

During the latter implementation phases, ITV's network server teams encountered a bizarre (at the time) performance issue. One particular network operating system vendor was implementing Path MTU discovery (a technique for avoiding fragmentation) each and every time a new client-server session started, and then timing out the routing cache at regular intervals, resulting in further Path MTU discovery dialogues. The DMVPN tunnels had a 1,472 byte MTU (derived from the 1,500 byte MTU of the underlying MPLS LSP less 24 bytes of standard GRE header and 4 bytes of optional GRE header, used for the key field). The DF bits of the IP packets were being seen by the routers, but not all of the Internet Control Message Protocol (ICMP) replies were being sent by the routers. By default, the routers limited the number of replies they sent out for ICMP to 2 pps as a basic denial-of-service defense mechanism. This meant that Path MTU discovery was failing some of the time.

ITV solved the problem by rate limiting the ICMP replies for “DF unreachable” to 1 ms, thus allowing for a maximum of 1,000 replies per second. This was achieved using the global config command **ip icmp rate-limit unreachable DF**. ITV also tested the **ip tcp adjust-mss** command on some routers, but it was unnecessary after the ICMP rate limit was set.

FURTHER READING

- Dynamic Multipoint VPN
cisco.com/packet/181_7c1
- Cisco 2800 Series Integrated Services Routers
cisco.com/packet/181_7c2

Say It with IPTV Services

**Tapping the Flexibility of a Converged Network Infrastructure
That's smart. Ask HKBN.**



© Justin Quarriglia/National Geographic Image Collection

By Janet Kreiling

“It’s not what you have that counts, but how you use it,” as the saying goes. Actually, what Hong Kong Broadband Network Limited (HKBN) has is quite impressive—a Carrier Ethernet network capable of delivering 1 Gbit/s to any residence or business. But it’s what the company is doing with the network that earned HKBN the 2005 Global Entrepolis @ Singapore award for innovation, sponsored by the *Asian Wall Street Journal* and the Singapore Economic Development Board. HKBN is extending 10- and 100-Mbit/s services to all subscribers who want them at considerably lower prices than its competitors; 1-GB service is available for a premium price. And most recently, HKBN has been building a healthy market for IPTV.

A subsidiary of City Telecom (H.K.) Limited, HKBN has grown from 12 employees to Hong Kong’s second largest alternative service provider in a little more than ten years. How? HKBN

reaches out to underserved markets, starting with people who have never used broadband, offers them a year’s service free, and then continually introduces new services. By the end of 2005, its network, built at US\$130 per home passed, had a reach of approximately 2 million residences, 90 percent of those in the city. Symmetric 100-Mbit/s and 1-Gbit/s services are US\$27 and \$172 per month, respectively.

Underpinning HKBN’s successful content strategy for IPTV and its other innovative content and marketing strategies is a Cisco IP Next-Generation Network (IP NGN) infrastructure. “HKBN’s greatest strength remains its Cisco IP NGN converged infrastructure, which enables the carrier to stay flexible with its business strategies,” according to an October 2005 case study conducted by research firm IDC.

Network Supports Business Strategy

Cisco equipped HKBN's IP NGN Ethernet infrastructure, and has also counseled the carrier on "creating a network architecture that delivers a better customer experience, enables it to provide any and all services it wants to, and integrates seamlessly with the rest of its video infrastructure," says Pankaj Gupta, senior manager for service provider marketing at Cisco. "Cisco helped to architect a network that complements and helps enable HKBN's business strategy."

With its Cisco IP NGN, HKBN has deployed an intelligent network that is scalable and resilient to changing subscriber usage patterns via its dynamic bandwidth capabilities. For example, network usage and subscriber traffic patterns can vary widely between mostly broadcast TV and mostly video-on-demand (VOD) services. HKBN can set its network to reallocate bandwidth dynamically if the traffic mix in a given area of subscribers varies across services, over time and subscriber growth, or around major sporting or other events. In this way, capacity is as available and used as efficiently and economically as possible, explains Wayne Cullen, senior manager in the Product and Technology Marketing Organization at Cisco.

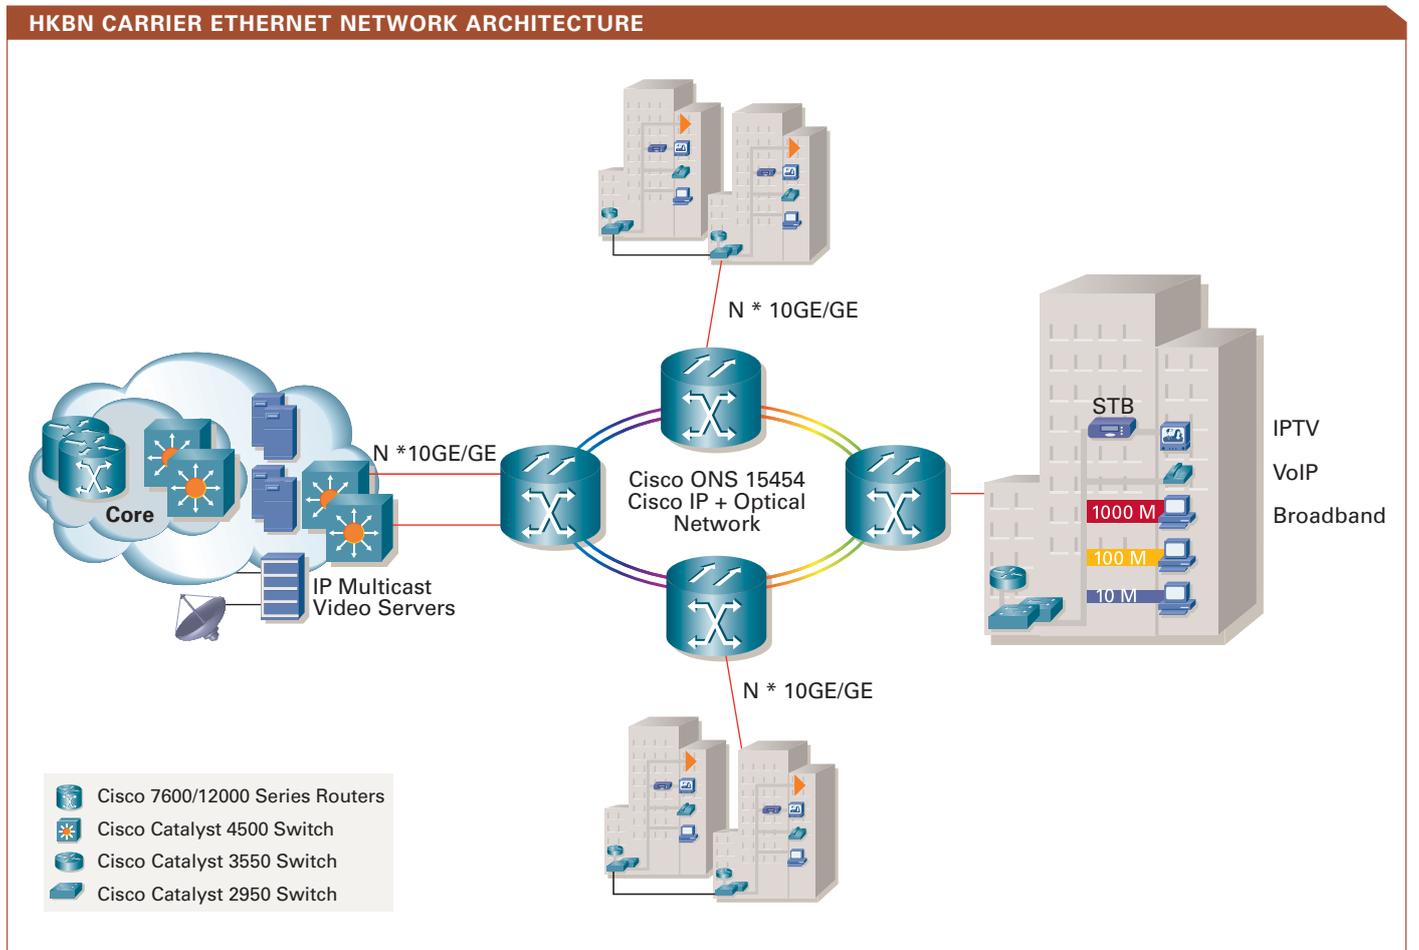
HKBN has essentially created a city-wide Ethernet

infrastructure (see figure). The core transport network is built on the Cisco ONS 15454 SDH Multiservice Provisioning Platform (MSPP), which supports various Ethernet speeds as well as sub-50-millisecond Resilient Packet Ring (RPR) protection. The latter is essential for voice and video quality of service (QoS) as well as HKBN's service-level agreements (SLAs). At the network core is a Multiprotocol Label Switching (MPLS) backbone with Cisco 7600 and 12000 series routers. Linked to the core are Cisco Catalyst 4500 Series Ethernet switches, which act as aggregation nodes; from there links fan out to the city's many multitenant high-rise buildings. Within these high rises, HKBN deploys Cisco Catalyst 3550 and 2950 series switches, the former in the basement and the latter on individual floors. Most of the network runs over fiber; copper might handle the last few hundred meters to an individual apartment or condominium. The network delivers 100 Mbit/s over copper or fiber, and 1 or more Gbit/s over fiber.

"With Cisco IP NGN, HKBN is enabling an enhanced subscriber experience," says Cullen, "so it is better able to monetize its services and ensure the needed QoS for paying IPTV customers."

"HKBN has been able to roll out a succession of new services without making any changes in the

CERTIFIED NETWORK
HKBN was one of the first service providers to earn the new Cisco Powered Network QoS Certification, which requires third-party verification that the network meets Cisco's best practices and standards for QoS.



infrastructure,” adds Gupta. “Moreover, the network is future-proof: the IP NGN architecture will be able to handle whatever HKBN demands of it, such as high-definition TV.”

If HKBN chooses to compete for premium pay TV subscribers, HDTV, or support online multiplayer gaming, it can do so. As noted in the October 2005 case study prepared by IDC, “Essentially, with an IP NGN converged network in place, HKBN can treat these issues simply as business decisions and not as technology-related roadblocks.”

IPTV: Marketing to the Masses

The fastest-growing segment of HKBN’s triple play services, IPTV exemplifies the provider’s smart marketing strategy. Specifically, it is:

- Expanding the market, rather than poaching subscribers from other providers. Most of its IPTV customers have not previously subscribed to any pay TV service.
- Bringing IPTV to markets not served by competitors, especially the Chinese-speaking population in Hong Kong, which until 1997 was a British colony whose official language was English.
- Offering desired TV programming not provided by other companies, presented in Chinese rather than English.

The market is wide open because pay TV penetration is low. “Hong Kong has a population of 7 million, but fewer than half of the city’s residences subscribe to pay TV services,” notes Gupta.

At the heart of HKBN’s IPTV value proposition is its savvy market segment analysis—its ability to identify target markets and serve them with the right content at the right price. In this case, HKBN can offer IPTV at the aggressively low price point of US\$16 per month, a figure that is possible because of the carrier’s low operating expenses and flexible IP NGN converged infrastructure, according to Gupta.

Content Caters to Diverse Personal Interests

Hong Kong residents in general are technologically sophisticated, embracing new information services and gadgets. The more than half who don’t already subscribe to pay TV are either not interested in premium channels or cannot afford them. HKBN wins them over with low prices, 50-plus channels, 23 of them interactive (its network has a total capacity of 200 channels), and content not likely to be delivered by its competitors. For example:

- Highly popular Japanese football (soccer)
- Live interactive healthcare programs, with physicians to answer questions

New Cisco ME 3400 Series Ethernet Access Switch

In its Carrier Ethernet network, HKBN beta tested this new switch, which Cisco engineered specifically for service provider metro access. The Cisco ME 3400 Series Ethernet Access Switch operates at Layers 2 and 3, and is optimized for protected delivery of high-bandwidth Ethernet to the home (ETTH) triple play and Ethernet to the business (ETTB) VPN services to multiple customers.

Installed at the customer’s premises, the Cisco ME 3400 Series can be deployed economically with a pay-as-you-grow approach. The switches support multiple software images, with feature sets for triple play, premium triple play/Layer 2 VPN services, or Layer 3 VPN services. Providers pay only for the features they need now and can add new ones with a simple upgrade for optimizing their CapEx and OpEx. A single platform for the ETTH and ETTB also reduces the provider’s training, maintenance, and sparing costs.

What’s more, the Cisco ME 3400 Series switches provide the most comprehensive, enhanced security for Carrier Ethernet deployments. Among their features are access control lists (ACLs) and IEEE 801.X support at the network level; control plane protection, storm control, and port security at the switch level; and UNI/NNI, Dynamic Address Resolution Protocol Inspection, and IP Source Guard at the subscriber level. For more on these new switches, see cisco.com/packet/181_8a1.

- The “Drama Buffet,” launched in September 2005, features episodes of popular Chinese, Korean, and Japanese soap operas and other programs in near video-on-demand (NVOD) format. Viewers choose from episodes that air on the next hour or half hour, cycling through multiple episodes if they wish.

“Apart from delivering content which caters to the general public, HKBN IPTV also addresses the audience’s diverse personal interests,” says Haily Leung, senior vice president at HKBN Digital TV. “This drives us toward more interactive and personalized services for our audience, making full use of the superiority of our IP network infrastructure.”

Cisco’s Commitment to Video

Video, as HKBN has aptly discerned, is one of today’s biggest opportunities for service providers.

“Cisco has made a commitment to provide networks on which service providers can offer video services in a reliable, profitable model,” says Jeffrey Spagnola, vice president for global service provider marketing at Cisco. “We have learned how video architectures need to scale, whether we are talking about multicasting, VPLS [Virtual Private LAN Services], or MPLS in the infrastructure. Cisco has learned how to deliver economic value and how to optimize the use of bandwidth and enable a better customer experience.”

Underscoring Cisco’s commitment is its intent to acquire (pending regulatory approval) Scientific-Atlanta, a global provider of set-top boxes, end-to-end video distribution networks, and video system integration. “The assets Scientific-Atlanta will bring to Cisco’s solution set are very complementary to our ability to transport video triple play over a digital infrastructure,” says Spagnola. ■

FURTHER READING

- IDC Case Study: “HKBN Implements IPTV Solution on Cisco IP NGN Converged Infrastructure”
cisco.com/packet/181_8a2
- White paper: “Building the Carrier-Class IP NGN”
cisco.com/packet/181_8a3

Is the Hong Kong Market Unique for IPTV?

Hong Kong might seem like a unique market, with 7 million people in some 2.2 million residences, densely packed into 422 square miles of land. Concentrating residents into multitenant buildings makes a fiber network relatively inexpensive, even when it includes fiber to the home. But, according to Gupta, “The underlying architecture for IPTV across access technologies, such as fiber, DSL, and Ethernet, are similar. The major components of IPTV architectures are set-top boxes, access, aggregation, core, video head-ends, and middleware.”

So, what makes HKBN unique is not its fiber, but its savvy market segment analysis, content choices, and IP NGN infrastructure. Profitable IPTV does not depend on a high-density market such as Hong Kong, notes Gupta, but rather on providing an enhanced customer experience, meeting the content needs of consumers, and differentiating yourself from competitors.

Gaining the Ethernet Edge

New OAM protocols enhance a carrier's service deployments and make managing and monitoring Metro Ethernet networks easier.

By Chiara Regale

Ethernet is a widely deployed technology not only in LANs but also increasingly in metropolitan- and wide-area networks (MANs and WANs). Ethernet in the MAN/WAN has amplified the need for flexible, comprehensive management and monitoring functionalities to increase end-to-end service operational efficiency. To foster wide-scale adoption of Metro Ethernet and broadband services, equipment vendors and service providers must jointly consider ways to facilitate and expedite service deployment. Some of these ways can be found within protocols such as Ethernet Operations, Administration, and Maintenance (Ethernet OAM) and Ethernet Local Management Interface (E-LMI).

The Ethernet OAM standard brings to Ethernet much of the OAM functionality found in traditional carrier

network technologies. Advanced capabilities such as link monitoring, fault detection/isolation, and remote loopback control give carriers the end-to-end OAM tools required to maintain and monitor their Metro Ethernet networks in a manner consistent with other carrier technologies.

OAM gives network operators the ability to monitor the health of a network and, more importantly, the health of the services being offered, and quickly determine the location of failing links or fault conditions (see Figure 1). This article details the functionalities of Ethernet OAM protocols, specifically E-LMI and new OAM capabilities within IEEE 802.3 Clause 57 (formerly known as 802.3ah) and 802.1ag. MPLS OAM falls outside the scope of this discussion.

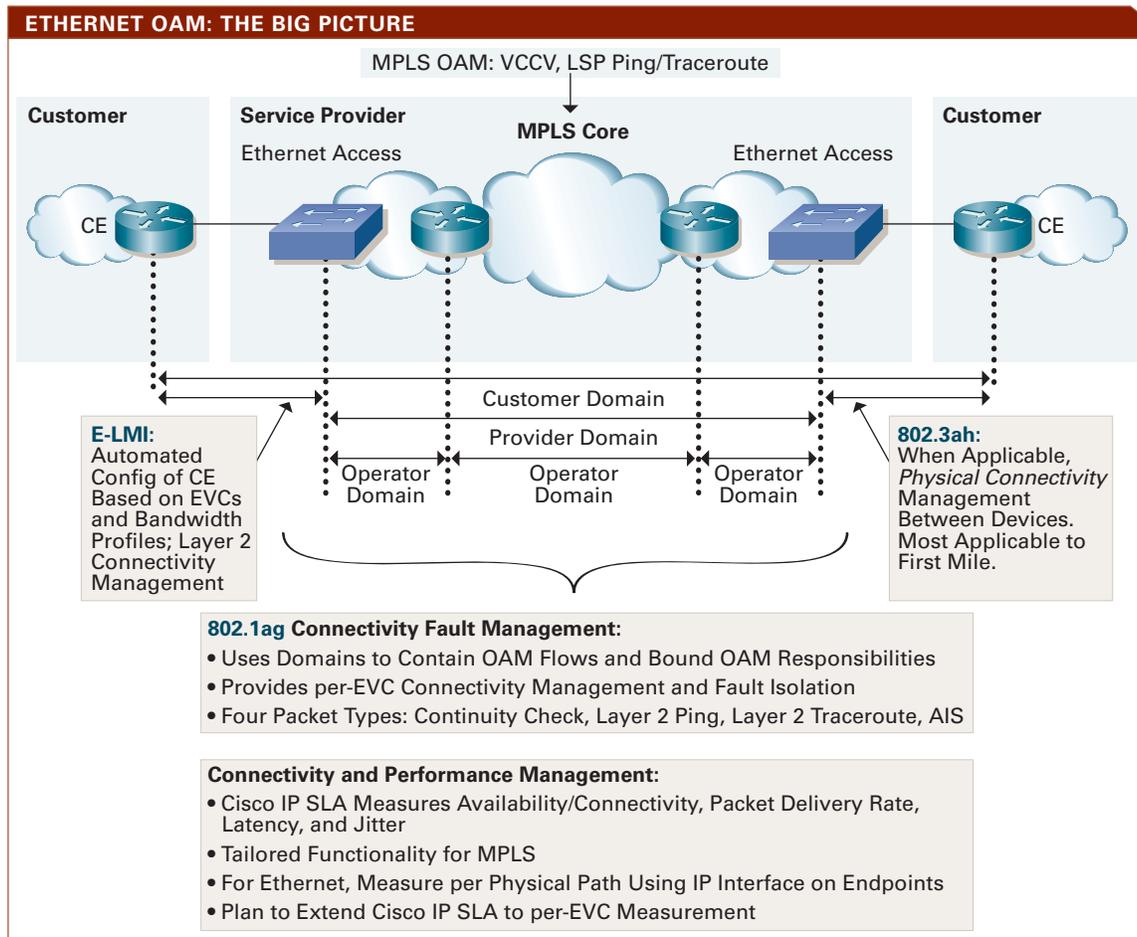


FIGURE 1 With Ethernet OAM, every network domain (from access to aggregation to core) and every application layer (Layer 2 or IP/MPLS) is characterized by its own OAM protocol and recovery mechanism.

FIGURE 2 The E-LMI framing structure is based on the IEEE 802.3 untagged MAC frame format. E-LMI messages are encapsulated inside Ethernet frames, making E-LMI implementation easy for organizations that already have Cisco routers and switches.

E-LMI Framing Structure				
Destination Address	Source Address	EtherType	PDU (Message)	CRC
6 Octets	6 Octets	2 Octets	46-1500 Octets (Data + Pad)	4 Octets

E-LMI

Based on ITU-T Q.933, X.36, the E-LMI technical specification enables customer equipment (CE) to request and receive status and service attribute information from the provider’s Metro Ethernet network, so that the CE can automatically configure itself to access Metro Ethernet services. E-LMI has local significance at the User Network Interface (UNI) between the Metro Ethernet access device and the CE, and also provides UNI and Ethernet Virtual Connection (EVC) status information to the CE. This information enables automatic configuration of the CE operation based on the Metro Ethernet network configuration. The E-LMI protocol notifies the CE of an EVC addition; notifies the CE of an EVC deletion; notifies the CE of the availability state of a configured EVC (active, not active, or partially active); and communicates UNI and EVC attributes to the CE. To transfer E-LMI information, a framing or encapsulation mechanism is required (see Figure 2).

IEEE 802.3 Clause 57 OAM at the Link Level

IEEE 802.3 Clause 57 provides new OAM standards for Ethernet that contain useful mechanisms for monitoring link operation, such as remote failure indication and remote loopback control. The OAM described in this standard provides data link layer mechanisms that complement applications that may reside in higher layers. Link Level OAM is intended for point-to-point and emulated point-to-point 802.3 links. OAM information is conveyed in frames called OAM protocol data units (OAMPDUs), which contain the appropriate control and status information used to monitor, test, and troubleshoot OAM-enabled links. OAMPDUs traverse a single link, being passed between peer OAM entities. They are intercepted by the MAC sublayer and cannot propagate beyond a single hop within an Ethernet network.

The three main functions of Ethernet OAM are discovery, link monitoring, and remote failure indication.



CHIARA REGALE is a product manager for the Catalyst 6500 Series Switch in the Internet System Business Unit at Cisco. A regular presenter at the annual Networkers conference, she focuses on product strategy and solution roadmap for the Catalyst 6500 in Metro Ethernet/broadband aggregation. She can be reached at chiara@cisco.com.

Discovery

This is a means for detecting the presence of an OAM sublayer at the remote peer and establishing OAM on the link. In the Discovery phase, the following information is advertised in Type Length Values (TLVs) embedded within periodic Information OAMPDUs:

OAM mode: Conveyed to the remote OAM entity, this mode can be either active or passive, and can also be used to determine device functionality.

OAM configuration (capabilities): Advertises the capabilities of the local OAM entity. With this information, a peer can determine which functions are supported and accessible, such as loopback capability.

OAMPDU configuration: Includes maximum OAMPDU size for receipt and delivery. This information along with the rate limiting of ten frames/sec can be used to limit the bandwidth allocated to OAM traffic.

Platform identity: A combination of an Organization Unique Identifier (OUI) and 32 bits of vendor-specific information. OUI allocation is controlled by the IEEE, and OUIs are typically the first 3 bytes of a MAC address.

Discovery includes an optional phase wherein the local station can accept or reject the configuration of the peer OAM entity.

Link Monitoring

Ethernet Link OAM provides a mechanism to support event notification that permits link monitoring for detecting and indicating link faults under a variety of conditions. Link monitoring uses the Event Notification OAMPDU, and sends events to the remote OAM entity when problems are detected on the link. Because 802.3 Clause 57 OAM does not provide a guaranteed delivery of any OAMPDU, the Event Notification OAMPDU may be sent multiple times to reduce the probability of a lost notification. A sequence number is used to recognize duplicate events.

Remote Failure Indication

Ethernet Link OAM provides a mechanism for an OAM entity to convey failure conditions to its peer, via specific flags in the OAMPDU. Failure conditions include *link fault*—loss of signal is detected by the receiver; for instance, the peer’s laser is malfunctioning;

Support for Ethernet OAM

Cisco's premier switching platform for Metro Ethernet access and aggregation networks, the Catalyst 6500 Series provides increased end-to-end service operational efficiency. Management and monitoring features such as E-LMI and IEEE 802.3 Clause 57 OAM capabilities are already available in Cisco Catalyst OS Software Release 8.5, and 802.1ag OAM functionality will be supported in Catalyst OS Software Release 8.6. By the end of 2006, Ethernet OAM protocols will be introduced in Catalyst 6500 Series Cisco IOS Software Releases, and will not require a hardware forklift upgrade. Also by the end of 2006, Cisco plans to introduce support for Ethernet OAM protocols in the Cisco 7600 Series Router, Catalyst 4500 Series Switch, Catalyst 3750 Metro Series Switch, and ME 3400 Series Ethernet Access Switch. Also underway is the addition of E-LMI functionality to CPE including Cisco Integrated Services Routers.

applies only when the physical sublayer is capable of independent transmit and receive. *Dying gasp*—an unrecoverable, vendor-specific condition (e.g., a power failure); may be sent immediately and continuously. *Critical event*—an unspecified, vendor-specific critical condition; may be sent immediately and continuously.

Remote Loopback and MIB Variable Retrieval

Link OAM provides an optional data link layer, frame-level loopback mode, which is controlled remotely. Remote loopback can be used for fault localization and link performance testing. Statistics from both the local and remote peer can be queried and compared at any time while the remote entity is in OAM remote loopback mode. An implementation may analyze loopback frames within the OAM sublayer to determine additional information about the health of the link (e.g., determine which frames are being dropped due to link errors). In addition, Ethernet Link OAM provides a read-only access to remote MIB variables limited to a specific MIB branch and leaf. The request-response nature of variable retrieval can also be used to estimate the link capability to support a service-level agreement (SLA), similar to IP Ping for measuring delay, jitter, and throughput. This function assumes that the time accessing the variable is negligible compared to propagation and queuing delay of the request and response.

IEEE 802.1ag OAM Network-Wide

Whereas IEEE 802.3 Clause 57 OAM monitors individual Ethernet links, 802.1ag focuses on the

monitoring and fault detection capabilities across a large Ethernet network, which help in troubleshooting problems network-wide and at the service layers rather than just at the link layer. IEEE 802.1ag provides service OAM capabilities for monitoring and troubleshooting end-to-end Ethernet service instances. IEEE 802.1ag specifies the Ethernet Connectivity Fault Management (CFM) functionality that can detect, verify, and isolate connectivity failures in virtual bridged LANs. To support rapid detection of faults and accurate fault isolation without excessive consumption of network resources, CFM functions are partitioned as follows:

- **Fault detection.** Continuity Check protocol is used to detect both connectivity failures and unintended connectivity between service instances. Connectivity Check Messages (CCMs) are multicast frames that can be transmitted at a high rate, but are simply forwarded as data within the network and, thus, do not impose a CPU processing load.
- **Fault verification and fault isolation.** These are administrative actions, usually performed after automatic detection of a fault or receipt of some other error report. Fault verification is also used to confirm successful restoration or initiation of connectivity. Fault verification uses the acknowledged Loopback protocol to verify connectivity. Fault isolation uses the Linktrace protocol to determine the path from one CFM entity to another. Each Linktrace Message is sent to a multicast address to allow it to be readily intercepted on the path to the destination entity that returns unicast Linktrace Replies.
- **Fault notification.** This is provided, possibly by using the Continuity Check protocol, when a connectivity fault is detected.

◆ ◆ ◆

Equipped with advanced OAM capabilities, it might soon be possible for carriers to verify the SLA requirements of a Metro Ethernet offering across an entire network. And metrics such as uptime, latency, and jitter could be continuously monitored. This monitoring can prove and improve SLAs between a carrier and its customers, as well as bolster the carrier's revenue potential with enhanced services and reduced customer churn. These and other performance monitoring standards are underway. ■

FURTHER READING

- Metro Ethernet Forum
www.metroethernetforum.org
- IEEE 802.3ah
ieee802.org/3/ah/
- IEEE 802.1ag
ieee802.org/1/pages/802.1ag.html

The Winding Road to IMS

Non-IP Multimedia Subsystem applications are still key in the service provider evolution toward IMS.

By David Barry

Recently, the IP Multimedia Subsystem (IMS) has generated quite a buzz within the service provider industry. It's been hailed by many as the next-generation service provider architecture standard that will usher in the era of seamless multimedia communications across wireline, wireless, and cable networks. Meanwhile, numerous skeptics emphasize that the complexity of IMS could ultimately make its road to deployment steep.

So, just what is the status of IMS today and its promise for the future? While the road to IMS might prove to be long and winding, Cisco believes it will have significant relevance and value for some providers and, accordingly, has dedicated resources and established strategic partnerships to help companies navigate this road.

IMS and the Cisco IP NGN Architecture

IMS is an open, standardized architecture that aims to merge multimedia services across the cellular world and IP networks using the same standard protocols for both mobile and fixed IP services. Based on Session Initiation Protocol (SIP), IMS defines standard control plane interfaces for creating new applications. IMS essentially takes the place of the control infrastructure in the traditional circuit-switched telephone network; the key difference is that IMS separates services from the underlying networks that carry them. In this way, presence-based services such as instant messaging and push-to-talk (PTT), and other services such as voice mail and e-mail, can reside on application servers anywhere and be delivered by multiple wired and wireless providers.

In December 2005, Cisco introduced enhancements to its open Service Exchange Framework (SEF) that enable support for IMS. A key component of Cisco's IP Next-Generation Network (IP NGN) architecture, the Cisco SEF provides ways to analyze, optimize, secure, and meter application and content-based services already being deployed by service providers across wireless, wireline, and cable networks worldwide.

The IP NGN architecture defines three fundamental layers of convergence. Cisco SEF operates at the service convergence layer:

- *Application convergence*—integrating new, innovative IP data, voice, and video applications for a rich subscriber experience on any device anywhere, and delivered over a seamless, intelligent, high-performance infrastructure.
- *Service convergence*—enabling providers to deliver “triple play on the move,” which combines voice, video, data, and mobility services for the overriding applications. Service convergence includes policy-based network access and control that is technology-agnostic and seamlessly compatible with any underlying networking medium: mobile, fixed, wireless, cable, DSL, optical, or Ethernet.
- *Network convergence*—enabling providers to migrate from deploying, managing, and maintaining multiple service-specific networks to delivering all services across a single network, which provides advanced multicast, quality of service (QoS), and security capabilities end to end. Most often this network is based on IP Multiprotocol Label Switching (IP MPLS).

The new Cisco Service Exchange Solution for IMS provides a comprehensive foundation for deploying IMS-based services such as PTT and fixed mobile convergence. The solution consists of Cisco and partner products that map to IMS specifications. (For more on these products, see the white paper, “Cisco Service Exchange Solution for IMS,” at cisco.com/packet/181_8c1.)

A cornerstone of an IP NGN network is providing the ability for service providers to analyze, optimize, and meter application and content-based services in their existing IP networks. This might require tracking services that traverse multiple network types, each with its own unique capabilities, and the services could originate and terminate on many different devices.

For example, cellular providers may evolve PTT services so their subscribers can push to access a voice-enabled portal containing applications such as Mapquest. Directions could be spoken or displayed. Today, however, most providers have very little of the detailed subscriber information or control they need to deliver such services and applications.

Subscriber awareness and identity management are a few of the key capabilities now being enabled by the Cisco Call Session Control Platform (CSCP), an integral component of the Cisco Service Exchange Solution for IMS. Cisco CSCP supports the IMS reference architecture as a 3GPP-defined Call Session Control Function (CSCF). Cisco CSCP enables innovative IMS applications such as voice over broadband, PTT, presence-based services, video telephony, and fixed mobile convergence applications.

Among its other capabilities, Cisco CSCP helps to identify users and their devices, determine a person's location, and establish presence for that person, including sharing his or her status (on or off network) with other subscribers. With these capabilities, providers can deploy presence-based services such as PTT, instant messaging, call routing and screening, and 3G+ mobile applications such as streaming audio and video and interactive gaming.

The Cisco CSCP also helps to simplify application development by handling tasks common to all applications only once. This approach allows providers to attain greater customer control than was previously allowed in a traditional applications environment. The subscriber database, profile information, presence and location, and other information that binds a subscriber to a particular service are all managed and controlled through a single mechanism.

“IMS is still in its early days, and therefore non-IMS-based service deployments are still important and driving a lot of the service revenue.”

—Mark Bieberich, Director of Communication Network Infrastructure,
The Yankee Group

With the Cisco CSCP, providers can mix and match their applications to create new customer services packages. A mobile carrier might, for example, build a contact list application that simply shows which users are online. Although this list might initially be used to support an instant messaging service, the carrier can make the same list available for a PTT or a “find me” service (allowing a subscriber to locate all the members of a predefined group simultaneously). By aggregating the capabilities of individual applications, the Cisco CSCP makes such applications quite simple to develop—in sharp contrast to the cumbersome process found in a traditional mobile applications environment.

Non-IMS Services Continue to Drive Revenue

Service providers are waging intense competitive battles as the explosion of new services and consumer end devices spur a thriving market. In the consumer space, gaming, network-based personal video recorders, video on demand (VoD), Wi-Fi networks, and mobility are especially high-growth areas. In this highly competitive, dynamic arena, providers must be nimble and fast. They need to deploy revenue-generating applications and services that they can offer to their subscribers *today*.

In fact, many services that providers might want to offer—from IPTV to Web to business IP VPNs and even messaging—do not necessarily need to be offered over an IMS infrastructure. And it isn't merely a coincidence that many of these services are non-IMS-based.

“IMS is still in its early days, and therefore non-IMS-based service deployments are still important and driving a lot of the service revenue,” according to Mark Bieberich, director of Communication Network Infrastructure at The Yankee Group.

A Brief History of IMS

An open, standardized architecture, IMS defines how IP networks should handle voice calls and data sessions in the emerging arena of multimedia communications across wireline, wireless, and cable infrastructures. At its core is SIP, the signaling system for setting up and handling calls and data sessions, which already is the standard for voice over IP products.

IMS was initially developed by the Third Generation Partnership Project (3GPP) to meet the requirements of GSM operators seeking to deploy IP applications over their 3G wireless networks. Standards bodies for CDMA wireless as well as wireline networks have since adopted specifications based on IMS.

CableLabs, the standards-setting organization for the North American cable industry, has also adopted the signaling core of the IMS specification for PacketCable 2.0. This specifies requirements for real-time, interactive, multimedia services over cable DOCSIS access networks.

IMS in Action: Fixed Mobile Convergence

Fixed mobile convergence is a highly promising application that has gained high interest among service providers. Fixed mobile convergence requires that a handset be able to move seamlessly from a cellular environment to a Wi-Fi networking environment such as a home or business. With such a dual-mode handset, users would have service on the cellular GSM network while driving in their car, but on the walk into their house, the handset would automatically sense Wi-Fi availability and switch to the Wi-Fi network.

At the Worldwide Analyst Conference in December 2005, Cisco demonstrated such an application. While carrying a dual-mode (Wi-Fi/SIP and GSM) handset to make and receive voice calls, the demo presenter roamed across two access networks: Wi-Fi/802.11g and cellular GSM, with the call remaining active throughout the handoff. Specifically, this demonstration enabled the interworking of a variety of network and signaling protocols including IMS, SIP, GSM, MAP, Media Gateway Control Protocol (MGCP), and Wi-Fi. The demo also highlighted the standards-based interoperability of the Cisco SEF with critical partner products and services including handset/clients, HLR/HSS, and policy servers.

In addition to showing how Cisco's Service Exchange Solution for IMS can deliver QoS over any access network while meeting operator-defined policy enforcement requirements, the demonstration validated the robustness of the key products underlying the solution, in particular the Cisco CSCP, Cisco BTS 10200 Softswitch, Cisco uBR10012 PacketCable-qualified cable modem termination system (CMTS), and the Cisco PGW2200 Media Gateway Controller.

While SIP traffic will likely increase significantly over time (the next 1 to 3 years), non-SIP/non-IMS applications and services dominate the vast majority of IP network bandwidth today and will likely remain present in many provider networks for some time to come. All of these non-SIP/non-IMS applications represent sizable revenue streams for service providers. So, it's important for providers to have networks that support both IMS and non-IMS traffic so that they can provide the most attractive, profitable service mix for their customers.

To that end, a service provider might decide to forego a "formal IMS" implementation altogether and instead deploy an alternative, IP-based control mechanism. Or a provider might decide on a combined IMS and non-IMS implementation within its architecture. The Cisco SEF supports providers in both cases.

With comprehensive support for non-IMS applications, the SEF can help providers deliver:

- *More services* via capabilities such as personalization and differentiation through self-subscription; content filtering through deep packet inspection; more granular charging models with extensive pre- and post-paid options.
- *Greater efficiencies* via service prioritization through deep packet inspection; preservation of video QoS via efficient management of oversubscription; greater scalability through content virtualization; network-based service control and charging multiple access technologies.
- *Better control* via fair use enforcement through deep packet inspection; higher availability through enhanced security; transparent mobile data networking across multiple access networks.

◆ ◆ ◆

Indeed, IMS provides the fodder for a provocative, forward-looking vision and merits consideration by service providers today. But it is not a homogenous standard. Rather it's an evolving series of specifications that is still being crafted, digested, and interpreted by vendors and carriers across the diverse segments of the wireline, wireless, and cable landscape.

In the end, some providers might not pursue IMS at all, while others might decide to follow a customized, mixed architecture to address their service control needs. Whether they decide to take the IMS road or not, Cisco has the solutions and expertise to help them be successful on their IP NGN journey. ■

FURTHER READING

- White paper: "Cisco Service Exchange Solution for IMS" cisco.com/packet/181_8c1
- White paper: "Cisco Service Exchange Framework: Supporting IMS for Mobile, Wireline, and Cable Providers" cisco.com/packet/181_8c2
- White paper: "Service Exchange Framework: Providing Greater Control for Cisco IP NGNs" cisco.com/packet/181_8c3

Modular to the Core

Modular features and flexibility make the Catalyst 6500 Series Switch an affordable option for midsized networks.

By Gene Knauer

LAN switching features that were once available only to the world's largest enterprise networks are now available and affordable to midsized networks with 250 to 1,500 end users—thanks to product innovations in Cisco's Catalyst 6500 Series. Smaller form factors, modular components, and lower prices have garnered great interest in the Catalyst 6500 Series product line among midsized organizations. A wide range of port densities and performance and feature options make the platform equally attractive for multiple deployment scenarios.

Modularity, Options, and Affordability

While smaller than enterprise networks, today's midsized networks often have many of the same requirements as networks with thousands of end users. Midsized organizations want converged data, voice, and video services. They consider high availability and resiliency essential. Ease of use is also a big plus, as networks become more complex and traffic volumes grow. And advanced services such as firewalls, intrusion detection, wireless LAN, and virtual private networking (VPN) are also highly desirable.

“Previously, midsized businesses would deploy the Catalyst 6500 in high-density and high-performance locations,” says Marie Hattar, director of Enterprise Switching in Cisco's Product and Technology Marketing Organization. “But now they can get customized models of the Catalyst 6500 that deliver the performance, features, and form factor they want, with the ability to add additional, integrated service modules later because of the platform's modular design.”

Cisco has enhanced the affordability of Catalyst 6500 models, now with an installed base of hundreds of thousands of chassis and millions of ports worldwide, through the introduction of a Supervisor 32 Engine, giving midsized customers an alternative to the Supervisor 720. The 32 Gbit/s of throughput in the switching fabric in the Supervisor 32 is ample for most midsized networks, compared with the 720 Gbit/s Supervisor Engine, which is standard for large enterprises. The Catalyst 6500 Series chassis also comes with the choice of three, four, six, nine, or 13 slots.

“Midsized networks include many different kinds of businesses with different needs. They are definitely not ‘one size fits all,’” says Gautam Roy, Cisco product manager for the Catalyst 6500. “Giving the customer a lot of choices in how to customize LAN switches, and giving them a full range of features to choose from, has proven very successful.”



Columbia Association

CATALYST OF CHOICE Midsized organizations like the nonprofit Columbia Association are beginning to use the Cisco Catalyst 6500 Series for their converged networks.

Local Government Chooses Midsized Catalyst 6509

The network staff serving the Columbia Association, a planned community in Columbia, Maryland, chose the Cisco Catalyst 6509 switch for the core of the association's first LAN WAN to connect 42 facilities across Columbia. With 450 regular users and 1,500 seasonal employees accessing the network, the Columbia Association began planning for the move to a converged IP environment with data, voice, and video services in 2001. At that time there was no LAN or WAN, and none of the association's

42 facilities had e-mail. Only an aging minicomputer connected to point-of-sale locations at recreational facilities such as skating rinks, pools, golf courses, parks, and association offices using 56K leased lines. The phone system was equally antiquated: the private branch exchange (PBX) required 10-digit dialing between each member facility.

Program and Project Manager Nagaraj Reddi realized that an IP network could meet the Columbia Association's requirements for converged services, security, and resiliency. "We decided on the Catalyst 6509 for the core because it was much more economical to get one big switch to connect all of the facilities together rather than buying many smaller switches," says Reddi. "It occupies less space and gives us room to grow; if we need more ports or services we just add another service module blade."

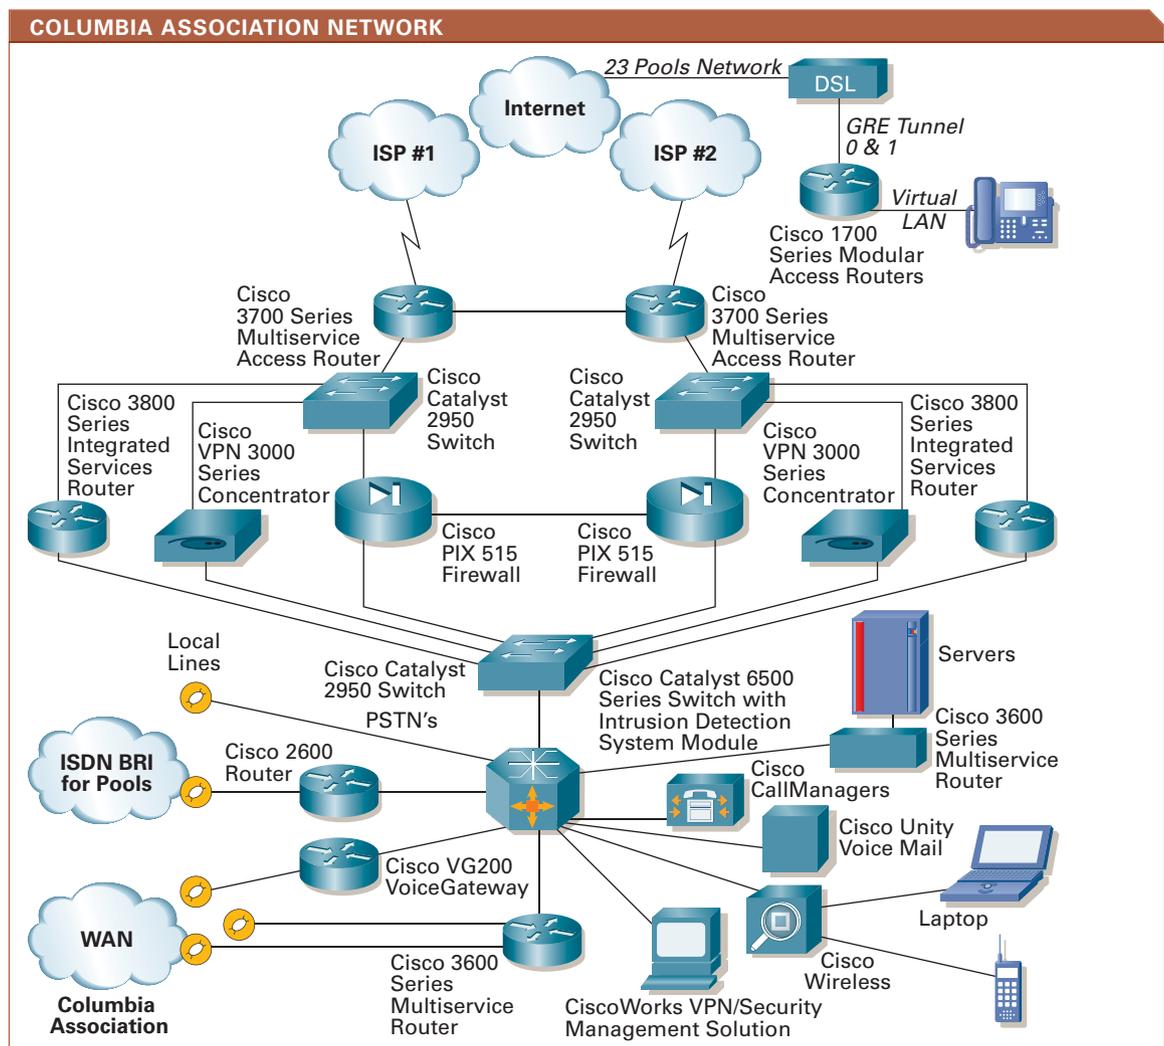
Cisco Catalyst 2950 Series edge switches in the main office and remote sites provide Gigabit Ethernet connectivity. Redundant Cisco 3800 Series Integrated Services Routers direct data, voice, and videoconferencing traffic over T1 lines, fractional T1 lines, DSL, and ISDN links to edge routers in Columbia Association facilities.

The Columbia Association has swiftly taken advantage of IP features in its Cisco end-to-end network. According to Reddi, it has installed nearly every model of Cisco IP phone, along with Cisco Aironet access points for wireless LAN connectivity in hotspots at many of the association's facilities. Segmented virtual LANs provide separate access for Columbia Association employees and visitors.

"The entire network was plug and play; we haven't had a single issue, and we continue to add new features without a problem," says Reddi, who manages the network with a staff of seven. "The phones became very popular once we showed people how to use all of the different features, and now every staff person who has a laptop can access the network through secure VPNs from their homes. The police pull up near WLAN hot spots throughout the city to access the network from laptops in their cars."

Reddi also likes the durability of the Cisco IP phones, which have endured extreme heat and cold in outdoor locations and chlorine beside indoor pools. For greater redundancy, Reddi will soon add another Catalyst 6509, although the network has thus far maintained its goal of 98 percent uptime.

INTEGRATED INFRA-STRUCTURE The Columbia Association network includes a Cisco Catalyst 6509 with an integrated Intrusion Detection System module in the network core.



Talk About It

Want to share your expertise on LAN and WAN routing and switching with your peers? Get answers to your questions from Cisco experts? Join the Networking Professionals Connection discussions at cisco.com/discuss/lan and cisco.com/discuss/wan.

High Availability and Resiliency

Availability and resiliency are crucial to networks with converged voice and data services. The Catalyst 6500 has redundant power supplies that use independent circuits to lower the risk of outages due to circuit failure. This helps ensure that power-over-Ethernet (PoE) devices like IP phones always remain on. Each Catalyst 6500 Series Switch can support redundant supervisor engines with Layer 3 subsecond failover to help ensure application continuity. Integrated online diagnostics monitor the system's vital signs.

"Many midsized customers think that if a switch goes down they can rely on a spare to provide redundancy, but with the Catalyst 6500 you have built-in high availability and resiliency to eliminate downtime and lost productivity," says Cisco's Hattar. "We have also introduced Cisco IOS Software Modularity, which makes the network even more resilient and available. You can restart processes, apply patches, and perform subsystem in-service upgrades without shutting the switch down."

Easier Than Ever to Deploy, Manage, and Maintain

A variety of tools for diagnostics and troubleshooting ease the burden of deploying, managing, and maintaining the smooth operation of applications on the Catalyst 6500. Smart ports, AutoQoS, and AutoSecure tools automate the consistent configuration of multiple ports for deployment of advanced services. Web-based tools like Cisco Network Assistant and Cisco View Device Manager, help in configuration, management, and troubleshooting of the Cisco Catalyst 6500 Series.

Integrated Internal and External Security

Midsized business networks need the same level of security as enterprise and service provider networks. Self-defending network features in the Catalyst 6500 protect the network from attacks in a variety of ways. Identity-based networking services allow network managers to identify users based on the IEEE 802.11 wireless LAN specifications and either allow access, disable access, or place guest users in a separate and secure VLAN. As users move from port to port, access control lists (ACLs), quality of service (QoS) settings, and other settings move with them. Network Admission Control (NAC) on the Catalyst 6500 Series

enforces access privileges of a device based on its level of antivirus software and software patch level, and ensure policy compliance.

Malicious attacks, such as Dynamic Host Configuration Protocol (DHCP) snooping, the flooding of the Address Resolution Protocol (ARP) table, and the use of spoofed IP addresses, are all mitigated on the Catalyst 6500 Series with the Cisco Integrated Security Toolkit. Should the switch itself become a target, hardware-based control plane rate limiters and policers intercept malicious traffic directed at the CPU to counter denial-of-service attacks. Integrated Cisco NetFlow support provides enhanced packet-capturing to detect anomalous traffic behavior. Cisco NetFlow is also useful in traffic monitoring and network capacity planning, and for applications such as granular accounting for user-based billing.

Longevity and Lower Total Cost of Ownership

As midsized networks retool to add new features such as converged services, advanced security, VoIP, and WLANs, they want to benefit from an integrated infrastructure that reduces their total cost of ownership. They also want the gear they buy today to last beyond the typical three-year product refresh cycle. Product manager Gautam Roy believes that the Catalyst 6500 Series will create investment protection for customer deployment needs for the next five years or more.

Prime candidates for the platform, according to Roy, are midsized data centers of organizations with up to 1,500 employees that may support thousands of online users. A special content switching module for the Catalyst 6500 provides Layer 4–7 services for faster Web response times. ■

FURTHER READING

- "Cisco Switching Solutions for the Midsized Business" video
cisco.com/packet/181_9a1
- Cisco Catalyst 6500 Series Switches
cisco.com/packet/181_9a2
- Cisco Solutions for Small and Midsized Businesses
cisco.com/packet/181_9a3

SPOTLIGHT ON:

Cisco Catalyst 6500 Series Wireless Services Module Cisco Wireless LAN Controller Module

The new Cisco Catalyst 6500 Series Wireless Services Module (WiSM) helps network administrators easily scale and manage wireless networks. The Cisco WiSM works with Cisco Aironet lightweight access points, the Cisco Wireless Control System, and the Cisco Wireless Location Appliance. Designed for mid-sized and large enterprise facilities, the module provides clustering capabilities of up to 3,600 lightweight access points per roaming domain and support for more than 10,000 wireless client devices per module.



cisco.com/go/wism

The Cisco Wireless LAN Controller Module (WLCM) allows small and midsized businesses and enterprise branch offices to cost-effectively deploy and manage secure wireless LANs. As a Cisco Integrated Services Router module, it delivers centralized security policies, wireless intrusion prevention capabilities, RF management, quality of service (QoS), and Layer 3 fast secure roaming for wireless LANs. The Cisco Wireless LAN Controller Module manages up to six Cisco Aironet lightweight access points and is supported on Cisco 2800 and 3800 Series Integrated Services Routers and Cisco 3700 Series Multiservice Access Routers.

cisco.com/packet/181_npd1

Core Routing Cisco CRS-1 Carrier Routing System: New DWDM Modules

Two new modules support service provider regional and long-haul transport applications: the Cisco CRS-1 single-port OC-768c/STM-256c Tunable WDMPOS Interface Module that provides up to 40 Gbit/s of data throughput across existing 10-Gbit/s dense wavelength-division multiplexing (DWDM) systems, and the Cisco CRS-1 4-Port 10GE Tunable WDMPHY Interface Module that is compatible with existing SONET/SDH operations support systems. Both modules are completely tunable across the C band with 50-GHz spacing and support high-gain Enhanced Forward Error Correction (EFEC), extending reach up to 1,000 km without requiring signal regeneration. The modules allow service providers to increase efficiencies, improve reliability, and reduce operational and capital costs by eliminating expensive, bulky transponder gear, even as video-based applications rapidly increase traffic in their existing DWDM networks.

cisco.com/go/crs

Cisco XR 12000 Series Routers: New Card and Adapters

The Cisco XR 12000 Packet Services Card (PSC-1) provides Session Border Control (SBC) functions integrated in the router, eliminating the need for overlay networks and standalone SBC appliances. The PSC-1 can be used by cable, wireline, and wireless service providers in deployments for peering or customer access. Separately, new Cisco 12000/XR 12000 Series shared port adapter (SPA) interface processors (SIPs)—the SIP-401, SIP-501, and SIP-601—host the common SPAs used to interconnect routers with each other or with gateways, Web servers, storage devices, and switches. The Cisco 12000/XR 12000 Series SIPs offer multiple rates up to 10 Gbit/s and support selected router models and specific SPAs.

PSC-1: cisco.com/go/12000

SIPs: cisco.com/packet/181_npd2

Edge Routing, Access, and Aggregation

Cisco Routers: New Shared Port Adapters

Five new shared port adapters (SPAs) offer enhanced feature support and accelerated service delivery for specific Cisco core and edge routers. The Cisco I-Flex design combines SPAs and SPA interface processors (SIPs) that can be used to prioritize voice, video, and data services for intelligent, flexible, secure networking. The OC48C/STM-16C ATM SPA for the Cisco 7600 Series Router includes a comprehensive ATM feature set for cost-effective routing in service provider POP and core applications. The Channelized STM-1/OC3 SPA supports channelized, T3/E3, and T1/E1 aggregation on Cisco 7600, 12000, and XR 12000 series routers. The 2-Port Gigabit Ethernet SPA (for the Cisco 7600 Series) and the 8-Port Fast Ethernet SPA (for the Cisco 12000 and XR 12000 series) are suitable for POP aggregation, Metro Ethernet, and Internet peering applications. The 2-Port OC-48C/STM-16C Packet over SONET/SDH and Resilient Packet Ring SPA are installed in the Cisco 12000 and XR 12000 series routers for applications including access and aggregation, WAN uplinks, and Internet peering.

cisco.com/packet/181_npd3

Cisco 7600 Series Routers: New Supervisor Engine and SIP 600

The Supervisor Engine 32 for Cisco 7600 Series routers delivers security, availability, and manageability services for Metro Ethernet access and smaller POP provider edge and enterprise WAN aggregation, deployed as a price-optimized, small form factor edge device. The Supervisor Engine 32 includes the policy feature card 3B (PFC3B) for architecture and feature consistency with the Supervisor Engine 720, and supports two uplink options: 8-port Gigabit Ethernet Small Form Pluggable (SFP)-based uplinks and 2-port 10 Gigabit Ethernet XENPAK-based uplinks. The Cisco 7600 Series SIP 600 supports bandwidth up to 10 Gbit/s, connects a variety of Cisco SPAs, including the 10 Gigabit Ethernet SPA and OC192/STM64 Packet over SONET SPA, and combines Layer 2/Layer 3 services for Metro Ethernet and VPN applications.

Supervisor Engine 32:
cisco.com/packet/181_npd4
SIP 600: cisco.com/packet/181_npd5

Switching

Cisco ME 3400 Series Ethernet Access Switch

The Cisco ME 3400 Series Ethernet Access switches are specifically engineered for service provider metro access, and optimized for protected delivery of Ethernet voice, video, and data services to residential customers and Ethernet VPN services to businesses. The Cisco ME 3400 Series supports multiple software images, giving service providers a pay-as-you-grow deployment model. Separate models provide an AC or DC power supply; both models include 24 Ethernet 10/100 ports and two 1000BASE-T Small Form-Factor Pluggable (SFP) uplinks. For a related article, see page 53.

cisco.com/packet/181_npd6

Cisco Catalyst 6500 Series Switch: New Supervisor Engine

Supervisor Engine 32 for Catalyst 6500 Series switches serves applications such as core functions, distribution, and access for small and midsize LANs as well as enterprise LAN/WAN access and service provider metro access and provider edge applications. Now available with Cisco IOS Software, Supervisor Engine 32 includes the policy feature card 3B (PFC3B) and Multilayer Switch Feature Card 2A (MSFC2A), and supports two uplink options: the 8-port Gigabit Ethernet SFP-based uplinks and 2-port 10 Gigabit Ethernet XENPAK-based uplinks. Supervisor Engine 32 provides architecture and feature consistency with the Cisco Catalyst 6500 Series Supervisor Engine 720, supporting all Catalyst 6500 Series classic modules and CEF 256-based modules.

cisco.com/packet/181_npd7

Security and VPNs

Cisco NAC Appliance

The Cisco NAC Appliance uses Cisco Clean Access technology to provide Network Admission Control (NAC) functions. The appliance allows network administrators to authenticate, authorize, evaluate, and remediate wired, wireless, and remote users and their machines

before allowing access to the network. The Cisco NAC Appliance integrates authentication, posture assessment, and remediation into a single device installed in a single rack unit. It extends NAC to all network access methods, including access through LANs, remote-access gateways, and wireless access points.

cisco.com/go/cca

Cisco Secure Access Control Server Version 4.0

Available in both appliance and software options, the Cisco Secure Access Control Server (ACS) products offer comprehensive, identity-based solutions for controlling network access. The Cisco Secure ACS version 4.0 software delivers many new capabilities, including a policy decision point for NAC deployments. This version also supports access by more devices and users than the previous version, enables profile-based access policies, and provides enhanced replication capabilities. Additionally, version 4.0 supports Cisco wireless LAN controllers and Cisco adaptive security appliances.

cisco.com/go/acs

Cisco Intrusion Prevention System Version 5.1

Among the numerous new features in the Cisco Intrusion Prevention System software version 5.1 is its ability to collaborate with edge routers and switches to preserve bandwidth through rate-limiting functionality. A single interface allows inline services on up to 255 virtual LANs across the network. Version 5.1 also provides a dedicated antivirus engine and inspection capabilities for Generic Routing Encapsulation (GRE) traffic, IP-in-IP traffic, and IPv6 traffic to detect and stop network attacks.

cisco.com/go/ips

Application Networking

Cisco 2600/2800/3600/3700/3800 Series Routers: Content Engine Network Module with WAFS Software

A content engine network module (NMC) with Cisco Wide Area File Services (WAFS) software installed is now available for Cisco access routers that are commonly deployed in branch offices.

The WAFS software enables infrastructure consolidation, simplifies data management, and reduces costs by offering centrally managed file and print services for users in remote offices. The NM-CE module is available for the Cisco 2600, 3600, and 3700 Series multiservice routers, and the Cisco 2800 and 3800 Series Integrated Services Routers.

cisco.com/packet/181_npd8

Storage Networking

Cisco Server Fabric Switch Portfolio

The Cisco Server Fabric Switch (SFS) product portfolio uses InfiniBand technology to create a high-performance, unified fabric for connecting servers into computing grids. The Cisco SFS 7000 Series InfiniBand Switch includes four models with up to 96 4X InfiniBand 10-Gbit/s or 20-Gbit/s full-duplex ports. The Cisco SFS 3000 Series Multifabric Server Switch supports up to 24 ports of 10 Gbit/s InfiniBand with up to 12 expansion modules. The Cisco InfiniBand Host Channel

Adapter offers high-performance 10-Gbit/s InfiniBand connectivity to servers based on PCI-X or PCI-Express. Cisco VFrame Server Virtualization Software is a data center provisioning and orchestration product for utility computing.

cisco.com/go/servernetworking

Cisco MDS 9020 Fabric Switch

The Cisco MDS 9020 Fabric Switch supplies 20 ports of 4-Gbit/s connectivity for a storage-area network (SAN) based on Fibre Channel technology. Offering simplified deployment and administration, the Cisco MDS 9020 is well suited for small and midsized businesses, branch offices, and enterprise workgroup applications. Included with the Cisco MDS 9020, the Cisco Fabric Manager tool provides fabric-wide discovery and simplified SAN management. Product features include topology discovery, fabric configuration and verification, provisioning, monitoring, and fault resolution.

cisco.com/go/mds

Network Management

Cisco Performance Visibility Manager

The Cisco Performance Visibility Manager (PVM) software offers integrated, end-to-end visibility into the performance of network and application resources. Used with Cisco Network Analysis Modules (NAMs), Cisco PVM delivers capabilities for traffic analysis, monitoring of application response times and bandwidth usage, and proactive monitoring of other key network metrics. The product's GUI simplifies troubleshooting, analysis, monitoring, and capacity planning. Suites of preconfigured reports present a comprehensive assessment of network and application performance.

cisco.com/go/pvm

Voice and Video

Cisco IP Phone 7900 Series: New Models

New models in the Cisco IP Phone portfolio offer choices for both high- and

low-volume environments. The Cisco IP Phone 7961G-GE is a manager model, and the Cisco IP Phone 7941G-GE is a business model for users with high-volume phone traffic or users running bandwidth-intensive (Gigabit) applications on collocated PCs. The Cisco IP Phone 7961G-GE provides six programmable backlit line and feature buttons. The Cisco IP Phone 7941G-GE provides two buttons. The Cisco IP Phone 7911G supports cubicle, retail, classroom, and manufacturing users with low-volume call traffic. This single-line basic model includes four dynamic softkeys and options for inline or external power. cisco.com/go/ipphones

Service Exchange Framework Products

The Cisco Service Exchange Framework includes several enhanced products for a service provider's voice, video, and data offerings. Among these new versions, the Cisco Service Control Application release 3.0 delivers new features for creating and monitoring services, as well as enhanced security, scalability, and integration choices. The Cisco Call Session Control platform release 3.0 supports innovative services such as push-to-talk (PTT), presence-based communications, video telephony, and fixed mobile convergence. A new software version for the Cisco PGW 2200 Softswitch supports media gateway control functionality for the IP Multimedia Subsystem (IMS) architecture. The Cisco BTS 10200 Softswitch release 4.5 provides enhanced operational capabilities, subscriber-focused telephony features, new platform support, and IMS integration. For a related story, see page 61.

SCA: cisco.com/go/servicecontrol
CSCP, PGW, and BTS:
cisco.com/go/sp-voice

Cisco IP Interoperability and Collaboration System

Cisco IP Interoperability and Collaboration Systems (IPICS) technology integrates PTT and other two-way radio systems with voice, video, and data devices. The Cisco IPICS Server software manages communications resources and provides authentication and security services, enforces user roles and policies, administers Push-to-Talk Management Center (PMC) clients, and collects audit information for training and operations

management. The Cisco IPICS PMC client is a PC-based application that lets users monitor and participate in up to eight PTT channels simultaneously. Cisco IPICS is covered in greater detail on page 34. cisco.com/go/ipics

Networked Home Linksys Wireless-G Broadband Router with SRX400

The Linksys Wireless-G Broadband Router with SRX400 (WRT54GX4) delivers a new generation of MIMO (Multiple Input, Multiple Output). When used with the new Linksys Wireless-G PC Card with SRX400 (WPC54GX4), this router supports faster throughput for home applications such as streaming content and voice over IP. The MIMO technology also reduces dead spots and increases range compared to traditional Wireless-G networks. The WRT54GX4 includes an Internet-sharing router, 4-port 10/100 Ethernet switch, and an enhanced Wireless-G Access Point. cisco.com/packet/181_npd9

Linksys Wireless-G USB Adapter with Wi-Fi Finder

The Linksys Wireless-G USB Network Adapter and Wi-Fi Finder (WUSB54G) is a pocket-sized device that gives users a wireless network scanner to locate hotspots and Wi-Fi connections. The WUSB54G LCD screen displays a located network's Service Set ID, signal strength, 802.11 mode, channel, and security. The WUSB54G can then connect to the wireless network via the USB port and client software on the user's notebook PC. cisco.com/packet/181_npd10

Linksys Internet Telephony Kit

The Linksys Internet Telephony Kit (CIT200) enables users of the Internet-based Skype phone service to place and receive calls on a handset instead of a PC. The kit includes a cordless handset, charger, and a base station that connects to a USB port on the user's PC. The base station handles communication between the handset and the Skype application on the PC. The handset also supports a variety of features for calling and messaging. cisco.com/packet/181_npd11

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 November 2005 and January 2006. For real-time announcements of the most recently released products, see "News Archive, News Releases by Date" at newsroom.cisco.com/dlls/.

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

Cisco IOS Software Metro Ethernet Enhancements

Selected releases of Cisco IOS Software now provide enhanced capabilities for Metro Ethernet deployments that use Cisco 12000 Series routers, Cisco 7600 Series routers, Cisco Catalyst 6500 Series switches, or Cisco Catalyst 3750 Metro Series switches. Among these enhancements is Hierarchical Virtual Private LAN Service (H-VPLS). VPLS is a multi-point, Layer 2 VPN technology that allows connection of multiple sites over a service provider-provisioned Multiprotocol Label Switching (MPLS) network. VPLS enables service providers to deliver popular new services such as multipoint Ethernet, Ethernet point-to-point Layer 2 VPN, and Ethernet access to Layer 3 VPNs. H-VPLS improves the scalability of VPLS by significantly reducing signaling overhead and packet replication requirements for the provider edge. Benefits derived from the new feature are simplified Layer 2 and Layer 3 access networks for multipoint transparent LAN services (TLS), the ability to integrate with an MPLS network, and improved scalability due to a tiered hierarchical approach. cisco.com/packet/181_npd12

Implementing and Troubleshooting IPSec Redundancy

The Networking Professionals Connection is an online gathering place for Cisco experts and networking colleagues. Following are excerpts from a recent Ask the Expert forum, “Implementing and Troubleshooting IPSec Redundancy,” moderated by Cisco’s Jazib Frahim. To view the full discussion, visit cisco.com/packet/181_10a1. To join other live online discussions, visit cisco.com/discuss/networking.

Q: *Is it possible to use high availability and load balance without a routing protocol?*

A: You do not have to use a routing protocol to use virtual tunnel interfaces (VTIs). Routing protocols such as Open Shortest Path First (OSPF) simplify manageability of the routes. Optionally, you can use static routes and point traffic over the tunnel interface.

Q: *Referring to IPSec Virtual Tunnel Interface documentation at cisco.com/packet/181_10a2, how is the load balancing done?*

A: You can load balance if you have two VTI tunnels defined. This is very similar to using the Generic Routing Encapsulation (GRE) tunnels.

Q: *Is it possible to set up the two hub routers as active/active for IPSec traffic instead of active/standby?*

A: You might be able to set up multiple Hot Standby Router Protocol (HSRP) groups on the inside and outside interfaces and then make one group active on Router 1 and the other group active on Router 2. You might need to do some policy routing to be able to force traffic to take one virtual IP (VIP) or the other.

Q: *I have the following IPSec redundancy using different ISP scenarios: Headquarters is connected to two different ISPs via two routers, namely ISP-A and ISP-B routers using different public IP. Behind ISP-A and ISP-B is a Layer 3 switch running OSPF. The remote branch is forming an IPSec tunnel to the headquarters ISP-A and ISP-B using a Cisco PIX firewall. What is the recommended IPSec config for both sites? The setup from top to bottom is:*

1. HQ network to L3 switch
2. L3 switch to ISP-A and ISP-B router
- 3a. ISP-A router to Internet
- 3b. ISP-B router to Internet
4. Internet to PIX
5. PIX to branch network

A: On the remote PIX firewall, you can have multiple set peer addresses in your crypto map. The PIX will try to connect with

the first address and if it does not respond, it will try to connect to the other address. Additionally, you will need to run HSRP on the inside network to ensure that the preferred router becomes the active routing/VPN termination device.

Q: *Referring to your advice on configuring “multiple set peer” on PIX, are both configs (a) and (b) below workable?*

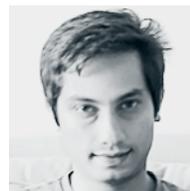
```
(a)
crypto map mymap 10 ipsec-isakmp
crypto map ios 10 match address acl
crypto map ios 10 set peer ISP-A
crypto map ios 10 set peer ISP-B
crypto map ios 10 set transform-set trans
```

```
(b)
crypto map mymap 10 ipsec-isakmp
crypto map ios 10 match address acl
crypto map ios 10 set peer ISP-A
crypto map ios 10 set transform-set trans
crypto map mymap 20 ipsec-isakmp
crypto map ios 20 match address acl
crypto map ios 20 set peer ISP-B
crypto map ios 20 set transform-set trans
```

How is the behavior of setting multiple peer for IPSec? PIX will peer to ISP-B if ISP-A fails. If Security Association (SA) to ISP-B timeout, will PIX peer to ISP-A or stick back to ISP-B? Assuming headquarters and branch can initiate an IPSec tunnel, how do I route the branch network from the Layer 3 switch to headquarters? Can I use Reverse Route Injection (RRI) at the ISP-A and ISP-B router?

A: The configuration in option (a) is correct. If the PIX is not able to connect to ISP-A, it will try ISP-B. If ISP-B is not available either, it will try ISP-A again, and so on. To route traffic from headquarters to branch, you can use HSRP on the inside and make your Layer 3 switch send traffic to the VIP. So, whichever ISP router is active will be responsible for establishing the tunnel.

Do you have a question about IPSec redundancy? Ask the Net-Pro Expert. Send your question to packet-netpro@cisco.com, with the subject line “Implementing and Troubleshooting IPSec Redundancy.” ■



JAZIB FRAHIM is a senior network security engineer for worldwide security services practices in Cisco’s Advanced Services for Network Security group. He can be reached at jfracim@cisco.com.

Intel, Continued from page 45

Bar-El compares the Layer 3 fast secure roaming capabilities in LWAPP to a post-office change-of-address service. “Mail addressed to your old address [analogous to the original access point the client associated with] isn’t actually sent there,” he says. “Instead, the post office [analogous to the controller] automatically forwards the mail to the most recent address on record.” Thus, all traffic is directed to the Cisco Aironet access point to which the client is currently associated and on to the Cisco wireless LAN controller.

Measuring Success

The Intel wireless team is currently performing packet level studies of wired and primary wireless network performance to measure and validate that wireless performance is comparable to wired. A report is due out early 2006 and will provide performance proof points

Network performance reporting will continue throughout 2006 and primary wireless ROI analysis will be added to strike a balance of technical and business proof points. Primary wireless ROI estimated reductions: network capital costs are expected to drop by 40 to 50 percent, and operational costs by 20 to 30 percent. Another factor is an estimated US\$25 yearly savings per employee for moves, adds, and changes for both voice and data services. “We’re also introducing an improved methodology to maintain an integrated network, which will further reduce support requirements across the entire environment,” says Stump.

Intel expects even more productivity gains when it introduces voice over wireless this year. “Employees will be able to establish a virtual office anywhere on the campus, resulting in more spontaneous collaboration, faster decision-making and action, and increased productivity,” says Jim Johnson, vice president and general manager of Intel’s Handheld Platform Group.

Intelligent Networks Rely Upon Intelligent Clients

Wireless manageability and usability depends upon intelligent networks that connect to intelligent clients. Intel IT has found that primary wireless requires intelligence in the clients to deliver end-to-end performance. “Employees expect their

notebooks, PDAs, and smartphones to easily connect to the mobile infrastructure and onto the service,” says Johnson.

More Consistent Than the LAN

The Intel Jones Farm deployment demonstrates that wireless networking can be pervasively deployed to support business-critical applications and advanced wireless services such as voice. The proof is in the uptake. “Originally, employees preferred the LAN because the quality and stability of the WLAN was inconsistent,” says Stump. “Now they prefer the WLAN because the user experience has been made consistent, simple—and with mobility—ubiquitous.” ■

EIGRP, Continued from page 17

B examines its local tables and finds the only path it has to 10.1.1.0/24 through A directly. Because D is advertising itself as a stub router, B has no reason to query D for an alternate path to 10.1.1.0/24.

Router A receives this reply from B, marks this route as unreachable, and removes it from the local routing table.

If the steps look familiar, they should; multiple router sites configured as stubs are treated the same as a remote site with a single router configured as a stub.

◆ ◆ ◆

If you’re counting neighbors to get to sleep, and you’ve designed the network correctly, you can count high enough to overcome almost any insomnia issues. Not only will the neighbor count be high enough to lull you to sleep, you’ll rest easier knowing your network is designed to withstand just about anything. The key is to limit the information that EIGRP advertises to each remote site, and to configure the remote sites as EIGRP stubs. In the future, newer EIGRP features will be introduced to increase scaling even further than the neighbor counts discussed in this article. ■

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.Aladdin.com/Cisco	IFC
Boson Software	www.boson.com/p16	A
Cisco Marketplace	www.cisco.com/go/marketplace/packetdvd	4
Cisco Press	www.ciscopress.com	B
Cisco Systems	www.cisco.com/poweredby	42
Citrix	www.citrix.com/cisco	52
Colt	www.colt.net	32
eIQnetworks	www.eiqnetworks.com/cisco	20
Empirix	www.empirix.com/cisco	12
Energis	www.energis.com	64
Funk Software	www.funk.com/cisco	70
Hong Kong Broadband Network	www.hkbn.net	26
IPcelerate	www.ipcelerate.com	10/56
Network General	www.networkgeneral.com/cisco4	60
OPNET Technologies	www.opnet.com	18
Panduit	www.panduit.com/dp38	IBC
Solsoft	www.solsoft.com/packet2	14
Spanlink Communications	www.spanlink.com	6
Spirent Communications	www.spirentcom.com/go/securitytest	50
Statseeker	www.statseeker.com	F
Trend Micro	www.trendmicro.com/cisco	46/47
Websense	www.websense.com/security	OBC

CACHE FILE

Snippets of Wisdom from Out on the Net

CYBER QUOTE

“Where a calculator on the ENIAC is equipped with 18,000 vacuum tubes and weighs 30 tons, computers in the future may have only 1,000 vacuum tubes and weigh 1.5 tons.”

—*Popular Mechanics*, March 1949

China’s Skyrocketing Broadband Usage

While the US currently has the highest number of broadband subscribers in the world, with 46.9 million subscribers, a Computer Industry Almanac report says China could surpass the US in broadband users in the next few years. Subscribers to broadband services worldwide are projected to exceed 500 million by 2010. South Korea leads in broadband subscribers per capita.

E-Paper’s Killer Application?

Electronics maker Siemens is readying a paper-thin electronic-display technology so cheap it could replace conventional labels on disposable packaging. In less than two years, Siemens says, the technology could transform consumer-goods packaging from the fixed, ink-printed images of today to a digital medium of flashing graphics and text that displays prices, special offers, or alluring photos, all blinking on miniature flat screens. [wired.com]

Small Shops Get Up to Speed but Still Like Phones

Switching from dialup to broadband Internet access improves productivity and efficiency in small organizations, but the telephone is the dominant business tool, according to a joint report from Covad Communications and Sprint, and conducted by Equation Research. The survey of nearly 500 representatives of US companies with fewer than 100 employees found that respondents spent more time online than they did on the phone, yet more than half chose the telephone as the tool their business could not function without. Telephones are the primary communication tool for small businesses, while the Internet is viewed as an information resource. [clickz.com]

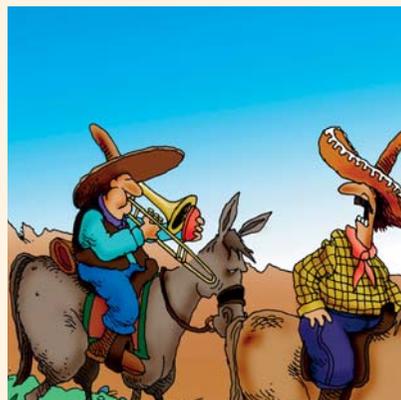
Net Lingo

Alpha Geek—The most knowledgeable, technically proficient person in an office or work group. [whatis.com]

Consumers Privacy Fears Continue to Escalate

Personalization remains something most consumers want, though their privacy fears continue to escalate. According to the second annual personalization study conducted by personalization vendor Choicestream, 80 percent of consumers in a 2005 survey were interested in receiving personalized content. Despite the fact users want more personalization and would buy more if they could get more personalized content, they’re not willing to share as much personal information as they once were. Respondents indicated decreasing willingness to share preference (59 percent in 2005 compared to 65 percent in 2004) and demographic information (46 percent in 2005 compared to 57 percent in 2004) to receive personalized content. [clickz.com]

THE 5TH WAVE



“Why can’t you just bring your iPod like everyone else?”

©The 5th Wave, www.the5thwave.com