

PACKET

CISCO SYSTEMS USERS MAGAZINE

THIRD QUARTER 2005

SOFTWARE INNOVATION

Expanding the Horizons
of Cisco IOS 32

The Application-Aware Network 59

Networks in Space 19

SPECIAL REPORT:
The Evolving Data Center 47



CISCO SYSTEMS

CISCO.COM/PACKET



PACKET

CISCO SYSTEMS USERS MAGAZINE

THIRD QUARTER 2005
VOLUME 17, NO. 3



32

ON THE COVER

Software Innovation

32

A spotlight on Cisco's commitment to driving technology leadership and meeting customers' evolving needs through software innovation.

IOS Modularity Debuts in the Enterprise

35

New Catalyst 6500 with Cisco IOS Software Modularity enables modular Cisco IOS subsystems to run in independent, self-healing processes—boosting operational efficiency and minimizing downtime for enterprises.

Managing the Network as a System

38

A new perspective on managing networks and software tools help enterprises and service providers mold the network to fit their business.

IOS XR: Scaling New Heights

43

Specifically designed and optimized for platforms that can scale and distribute processing as well as perform distributed forwarding, Cisco IOS XR Software directly addresses the requirements that IP next-generation networks demand.




38

SPECIAL REPORT

The Evolving Data Center

47


The next-generation data center must align business priorities to protect and optimize operations, increase productivity, and allow for growth. Find out about emerging architectural changes in the data center in this report.



19



65



71

IN EVERY ISSUE

Mail	3
Acquisitions	7
Calendar	6
Tech Tips	14
Advertiser Index	89
Cache File	90
The 5th Wave	90

TECHNOLOGY

ROUTING: Internet to Orbit	19
-----------------------------------	-----------

Exploring the potential of IP networking technologies in space.

ETHERNET: 10 GbE and Its X Factors	25
---	-----------

Demystifying 10 Gigabit Ethernet port types and pluggable modules on Cisco routers and switches.

ROUTING: Detecting Network Failures	29
--	-----------

New Bidirectional Forwarding Detection capability in the Cisco IOS Software helps speed failure detection and recovery.

ENTERPRISE SOLUTIONS

180,000 IP Phones	55
--------------------------	-----------

Bank of America begins a massive rollout of Cisco IP telephony to 6,000 of its branches.

Getting the Message	59
----------------------------	-----------

An overview of Cisco Application-Oriented Networking (AON) and the application-aware network.

Voice Goes Wireless	65
----------------------------	-----------

Delivering mobile voice over IP, data, and location services to diverse Wi-Fi environments.

SERVICE PROVIDER SOLUTIONS

IPTV/Video over Broadband	71
----------------------------------	-----------

Providing video in digital form over a broadband IP network is becoming a "must do" for many service providers to succeed in the convergence of consumer services.

Ridding Networks of DDoS Attacks	75
---	-----------

An integrated solution enables providers to offer new services to security-conscious customers.

SMALL AND MIDSIZED BUSINESSES

Smart Connections	81
--------------------------	-----------

Red Hat has room to grow with Cisco Catalyst 4500 Series Switch.

DEPARTMENTS

From the Editor	1	Technically Speaking	84
------------------------	----------	-----------------------------	-----------

Digital Edition Unveiled

Cisco's Russ White on mobility and pervasive networks.

User Connection	5
------------------------	----------

World Expo, Japan • Cisco Interactive Games • Cisco Network on Wheels

New Product Dispatches	85
-------------------------------	-----------

What's new from Cisco over the past quarter.

Tech Tips & Training	9
---------------------------------	----------

Cisco IT@Work Best Practices • Voice over IP Monitoring and Management • Reader Tips

NetPro Expert	88
----------------------	-----------

Advice from Cisco's Nick Chong on using Cisco Clean Access for Network Admission Control.

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

Charly Franklin/Getty
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@external.cisco.com.

Aironet, Catalyst, CCDA, CCIE, CCNA, Cisco, Cisco IOS, Cisco Networking Academy, Cisco Press, the Cisco Powered Network logo, the Cisco Systems logo, Cisco Unity, IOS, 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 © 2005 by Cisco Systems, Inc. All rights reserved. Printed in the USA.

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

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

This magazine is printed on recycled paper.



10%
TOTAL RECOVERED FIBER

FROM THE EDITOR

Digital Edition Unveiled

Many of the network services that have distinguished Cisco routers over the last 20 years can be credited directly to the innovative, intelligent features of Cisco IOS Software. In this issue, we explore what Cisco is doing to evolve IOS through technology innovation and capabilities that address the ever-changing requirements of the many markets IOS serves.

Speaking of change, I'd like to tell you about *Packet's* own software innovation: *Packet Digital Edition*.

Through new digital publishing technology, we're now able to deliver all the content, look, and feel of the print edition directly to your PC—before the print edition even rolls off the press. For readers abroad, who often endure lagging mail delivery, this new format can mean receiving *Packet* content weeks, sometimes months, before their print edition arrives.

Unlike many digital-format publications, *Packet Digital Edition* does not require a separate reader. You can view the magazine with any browser or platform with Flash Player 6 or above installed. Although *Packet Digital Edition* is best experienced while connected to the Internet—all URLs, e-mail addresses, and Further Reading links are clickable—you can also download your issue to read anytime, or anywhere you take your laptop.

Since we launched the digital edition in June 2005, we've had more than 30,000 unique visitors. Nearly 6,000 of you have already subscribed. Initial findings from our online survey show that most readers are excited about the new format. Ninety-three percent of all respondents reported a very positive or positive overall impression, while 86 percent found the format very convenient. If forced to choose between the print and digital versions, an impressive 61 percent of respondents said they would subscribe to the digital edition.

The most requested feature is the ability to keyword search an entire issue. This feature already exists in the digital format, but a recent upgrade to the software will soon allow you to search our entire archive collectively. (We hope to have all 2004/2005 *Packet* issues digitized and posted to cisco.com/packet by the time most of you are reading this letter.)

In the near future, we will add interactive content. We have plenty of ideas. Imagine diagrams that turn into narrated network animations, or photographs that offer up a short video on demand with the click of a mouse. We're also looking into the feasibility of embedding live chats, network simulations, and links to live labs directly in the pages of the digital magazine.

These are just some of our ideas. As always, we'd like to hear yours.

To check out the digital edition, visit cisco.com/packet/digital. You can fill out the survey online or send us an e-mail at packet-editor@cisco.com to let us know what you think.

David A. Ball

David Ball
Editor in Chief
daball@cisco.com



Rob Brodman

MAIL

Tip for Newbies

In your “Reader Tips” section (Second Quarter 2005), you published a tip on handling mistyped commands, but the recommended solution might be difficult for people who are new to routers and switches. An alternative solution is to use the `no ip domain-lookup` command in configuration mode, which avoids the name translation of mistyped commands.
—Antonio Manuel Santos, Verizon Dominicana, Santo Domingo, Dominican Republic



Missing Bandwidth

I tried following the “Reader Tip” submitted by A.G Teslicko (Second Quarter 2005) but did not find the bandwidth mentioned in my output. Please let me know if I need to check any more dependencies. The following output is from my router:

```
2621 Router#show frame-relay map
Serial0/0.1 (up): point-to-point dlci,
dlci 21(0x15,0x450), broadcast
status defined, active
Serial0/1.1 (up): point-to-point dlci,
dlci 22(0x16,0x460), broadcast
status defined, active
Serial0/1.2 (up): point-to-point dlci,
dlci 23(0x17,0x470), broadcast,
CISCO
status defined, active,
```

I am using a Cisco 2621 XM router running Cisco IOS Software Release 12.3-10 IPPlus image with two Frame Relay links.
—Girish Arora, Copycat Ltd., Nairobi, Kenya

The following is a response from the Cisco technical support engineering team:

When you use the show frame-relay map command, if the bandwidth reported by the switch is zero (0), the display will not show the bandwidth. The nature of the Gof4 LMI protocol is that the CIR field in the LMI response from the switch is

Correction

In the Second Quarter 2005 issue, page 7, we incorrectly stated that 8,852 network industry professionals have achieved CCIE certification.

As of June 2005, CCIE expert-level certifications totaled more than 12,000 worldwide. We apologize for the error.—Editors

not filled in by many switches. In this case, the switch sets the value to zero. Unfortunately, the router cannot distinguish between a CIR of zero, and an unknown CIR. Because of this, the zero value is treated as unknown (the safer option), and is not displayed.

Integrated Services Routers and Encryption

We are currently using link encryptors for data encryption. How can we use the Cisco 2800 and 3800 Series Integrated Services Routers for encryption? Can you recommend some links on how to configure encryption on these routers?
—Rajesh Ojha, Habib Bank Ltd, Karachi, Pakistan

You can find information on configuring encryption with Cisco Integrated Services Routers at cisco.com/go/routersecurity. The “Network Security Features on the Cisco Integrated Services Routers” data sheet is a good place to start. You may also want to review the white paper “SAFE: VPN IPsec Virtual Private Networks in Depth” at cisco.com/packet/173_2a1.—Editor

Longest Uptime for a Router

I have always enjoyed reading *Packet* and have worked with Cisco routers from the days of the AGS and IOS 9.0. Recently, I was cleaning my laptop and came across this copy of a `show ver` on an old Cisco 2500 Series Router:

```
hostname>show ver
Cisco Internetwork Operating System
Software IOS (tm) 3000 Software
(IGS-J-L), Version 11.1(17), RELEASE
SOFTWARE (fc1) Copyright (c) 1986-1998
by cisco Systems, Inc.
```

Compiled Tue 27-Jan-98 12:14 by phester
Image text-base: 0x030391F0, data-base: 0x00001000

ROM: System Bootstrap, Version 5.2(8a),
RELEASE SOFTWARE
ROM: 3000 Bootstrap Software (IGS-
RXBOOT), Version 10.2(8a), RELEASE
SOFTWARE (fc1)

Uptime is 4 years, 4 weeks, 5 days, 23
hours, 28 minutes System
restarted by reload at 06:57:55 UTC Sun
Mar 22 1998 System image file is
“flash:c2500-j-l_111-17.exe”, booted via
flash

cisco 2500 (68030) processor (revision M)
with 6144K/2048K bytes of memory.
Processor board ID 05344395, with hard-
ware revision 00000000 Bridging
software.
SuperLAT software copyright 1990 by
Meridian Technology Corp).
X.25 software, Version 2.0, NET2, BFE and
GOSIP compliant.
TN3270 Emulation software (copyright 1994
by TGV Inc).
1 Ethernet/IEEE 802.3 interface.
1 Token Ring/IEEE 802.5 interface.
2 Serial network interfaces.
32K bytes of non-volatile configuration
memory.
8192K bytes of processor board System
flash (Read ONLY)

Configuration register is 0x2102

Now that’s what I call uptime!

—Brad Wells, IBM Global Services,
London, Ontario, Canada

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.

First 21st Century World Expo Showcases Cisco Technology

Billed as the first worldwide showcase of “cutting-edge technologies” in the 21st century, an estimated 15 million visitors, including business and political leaders, have descended upon Aichi, Japan, this year to visit EXPO 2005. The World Exposition features spectacular pavilions from 125 participating countries.

Using Cisco Technology

All the networks in the US Pavilion—whose theme is “The Franklin Spirit,” in honor of US “founding father” Benjamin Franklin—incorporate Cisco IP networking technologies, which provide converged communication services, according to Masayuki Onishi, a senior engineer in Global Systems Engineering at Cisco Japan.

“Part of our success derived from the detailed discussion at the beginning of planning on how to best meet the needs of the Pavilion in actual situations such as structure of the building, transmission levels, and surrounding conditions,” says Onishi. “We suggested a converged network for information transmission and then developed the total design,” he continues.

The converged IP network provides access to broadband services, which are used for data exchange, image transfer, equipment



EXPO 2005 The “Mars Exploration Rover” exhibition in the US Pavilion displays a replica of the Rover and projects daily live downloads of real-time images from Mars. The Pavilion is the only institution outside of the US that is allowed to receive these images.

maintenance and management, phone services, voice mail, mobility, and security.

Hiroaki Mori, manager of the Open Service Center in the Tokai Customer Service Department at

Cisco partner PFU Systems adds, “We started to build the networks while constructing the Pavilion. It was a challenge to understand the radio wave propagation characteristics within the steel structure, and the environment was constantly changing as partitioning and displays were added. We had to repeat simulations and site surveys to tune to the most appropriate wireless conditions.”

EXPO 2005 was held March 25–September 25, 2005. For more information, visit cisco.com/packet/173_3a1.

Jumping into the Internet

Cisco’s Packetville (cisco.com/go/packetville) introduces young people to the marvels of the Internet as it relates to math, science, real-world solutions, education, and careers. The Packetville portal uses engaging animation and arcade-style game play to teach about networking. Students can join characters Peter and Penny Packet as they “jump into the Internet” and participate in a variety of activities to explore technology and its important role in society.

“Gaming is an exciting learning tool, and can be highly effective, depending on the audience and content,” says Marcia Sitcoske, senior director of the Creative Learning Studios at Cisco. “The educational and interactive games in Packetville engage students and provide an entertaining learning experience.”

Cisco also offers programs to teach educators and schools how the latest technology can improve the quality of their educational materials and reduce costs:

- **Content delivery networks and IP-based video solutions** help schools to connect remote students to on-campus learning resources, and enable classroom-to-classroom collaboration.

Continued on page 6



Bern King

PLAYING WITH PETER AND PENNY PACKET Learning about networking technology is fun for children through interactive online games with music and action-packed, imaginative scenarios peopled with colorful characters.

USER CONNECTION

Internet Games, Continued from page 5

- **Gigabit Ethernet, optical networking, and storage technology** help schools build advanced network infrastructures, enabling educators to take advantage of new education technologies.
- **Laptop computers and wireless networking** enable schools to make the Internet accessible anywhere on campus.
- **IP communications** help schools integrate data, voice, and video systems.
- **Virtual private networks, firewalls, intrusion detection systems, and video networking** help schools protect their records, information, and property

For more information about Cisco e-learning tools and other educational programs, visit cisco.com/packet/173_3c1, or contact Marcia Sitoske at msitcosk@cisco.com. ■

CISCO WORLDWIDE EVENTS

Sept. 19–22, 2005	Networkers Australia, Gold Coast, Australia
Oct. 26–28, 2005	Networkers Japan, Tokyo, Japan
Nov. 1–3, 2005	Networkers Korea, Seoul, Korea
Nov. 28–Dec. 1, 2005	CIPTUG Annual Users Conference, Las Vegas, Nevada, USA
Dec. 12–15, 2005	Networkers France, Cannes, France
Dec. 19–21, 2005	Networkers China, Beijing, China
cisco.com/warp/public/688/events.html	

1/2 Pg Ad

Network on Wheels Demonstrates Latest Cisco Technologies

Working with Cisco business partners, the Cisco Network on Wheels (NOW) program reaches out to small and midsized (SMB) companies in remote locations throughout the US by delivering hands-on opportunities to view the latest Cisco business solutions. A technology showcase on wheels, the network includes four 25-foot vans that are equipped with the latest Cisco voice, video, networking, wireless, and security technologies, and provides partners with a convenient and effective demo platform and selling resource for working with SMB customers.

“People are busy, and at times companies miss opportunities to investigate new business solutions. Each van provides an environment that includes the very latest Cisco technologies—and companies don’t have to travel far to try them,” says Bill Szmyd, director of market management for Cisco’s Commercial Marketing Group, West Field Operations. The mobile network has completed stops in more than 245 cities, and has covered more than 32,000 miles in 2005.



CISCO NOW With the Network on Wheels, remote SMBs have an opportunity for hands-on interaction with the latest Cisco technologies.

For more information about NOW schedules and activities, contact your area field marketing manager. ■

Recently Announced Cisco Acquisitions

Acquired		Employees	Location
FineGround	Provider of network appliances that accelerate, secure, and monitor application delivery while minimizing bandwidth usage and maximizing infrastructure capacity in the data center. FineGround employees will become part of Cisco’s Security Technology Group.	42	Campbell, California, USA
KiSS Technology A/S	Provider of networked entertainment devices. KiSS employees will join Cisco’s Linksys group, the leading provider of wireless and networking hardware for home, small office/home office (SOHO), and small business environments.	65	Hørsholm, Denmark
M.I. Secure Corporation	Provider of advanced features and functionality for security and VPN solutions. M.I. Secure employees will join Cisco’s Security Technology Group.	11	San Jose, California, USA
NetSift	Provider of deep packet processing solutions. NetSift will become part of Cisco’s Internet Systems Business Unit.	15	San Diego, California, USA
Sheer Networks	Developer of software that enables operational and billing systems to use more accurate data. Technology from Sheer will be used to develop Cisco’s next-generation management products for its service provider and enterprise products. Sheer employees will join the Cisco Network Management Technology Group.	100	Mountain View, California, USA and Bangalore, India
Sipura Technology	Provider of consumer voice over internet protocol technology. Sipura employees will join Cisco’s Linksys group, the leading provider of wireless and networking hardware for home, SOHO, and small business environments.	12	San Jose, California, USA and Petach Tikva, Israel
Vihana, Inc.	Provider of semiconductor solutions for the computer and communications industry. Vihana technology will be integrated across multiple platforms across Cisco technology groups.	27	Sunnyvale, California, USA

The Cisco Connected Workplace

Implementing advanced technologies makes for a leading-edge work environment.

By Rich Gore

With one of the largest and most complex enterprise environments in the world, Cisco faces the ongoing challenge of fostering collaboration within a diverse and distributed employee population. But recently, employees at Cisco gained a more innovative, flexible, and effective place to work, thanks to a partnership between Cisco IT and the Cisco Workplace Resources (WPR) group, which implemented Cisco wireless and IP communications technologies in the workplace.

This article describes how others can draw on Cisco IT's real-world experience with the Cisco Connected Workplace to help support similar enterprise needs within their own organizations.

Setting the Goals

At Cisco, the primary goals for the enhanced workplace environment are to increase employee productivity and satisfaction, promote and support collaboration, accommodate different employee work requirements and styles, reduce real estate costs, and reduce technology operational expenses.

Before implementing the new workplace environment, it became clear that Cisco employees were working more collaboratively than individually. And as business and technological issues become increasingly complex, the need for collaboration with team members in the same building or at various sites worldwide has become an essential component in achieving success. It was evident that the traditional workspace environment did not fully support the needs of employees, and that some workspaces were not being fully utilized; studies showed an average vacancy of 65 percent.

"Nobody would consider building a manufacturing facility that they intended to use just one-third of the time. And yet that's what we routinely do with work space. We realized that assigning resources based on utilization would significantly reduce costs," says Mark Golan, vice president, Worldwide Real Estate and Workplace Resources at Cisco.

Flexible Work Areas

The new workplace environment supports a range of work styles and culture by offering a variety of workplace options that provide space for private to collaborative and informal to formal work activities. In addition, employees now have designated spaces for both planned and spontaneous meetings, and specific areas are designated for concentrated work. The new environment allows employees to choose any workspace or technology option for as long as they need it—for a short time or for the entire day.

"Increasingly, we see that even employees who work onsite are likely to be mobile . . ."

—Christine Ross, Manager, Workplace Effectiveness Team, Cisco

Incorporating Advanced Technologies

Along with the enhanced work spaces, wireless technology and the mobility features of Cisco IP communications solutions support this new employee mobility.

"The work environment we've been building [in the past] is not necessarily what employees need now," says Christine Ross, manager in the Workplace Effectiveness Team in the WPR group. "Increasingly, we see that even Cisco employees who work onsite are likely to be mobile within the building," adds Ross.

After the cross-functional team defined the business requirements, Cisco IT identified the technological innovations that would help to support the new environment and achieve the technology cost-reduction goals.

The central technology design principle was to achieve a predominantly wireless environment, which would require 60 percent less cabling (cabling is one of the most labor-intensive network infrastructure components). The resulting reduction in switch ports would free up space for other switch resources, such as blades for security, content service, or additional infrastructure applications.

To support mobility in the workplace, Cisco IT built a wireless infrastructure that supports a denser-than-usual user population. In a typical Cisco building, Cisco IT provides wireless data connectivity at a density of about 25 employees per access point. Each quarter of a typical building has two or three wireless access points, whereas the new workplace building has 11 access points.

Supporting voice over wireless LAN also required special design attention.

“Wired networks set up for voice can provide a measure of call admission control, whereas wireless access points currently can’t differentiate between a voice and data stream,” says Dave Castaneda, a member of the technical staff in Cisco IT.

The need for increased access points arose because more than twice as many employees can work in the new workplace environment than in a typical Cisco building; employees cluster in small areas, which further increases density; and Cisco IP Communicator uses the same wireless connection that is used for data, adding to wireless bandwidth requirements.

Technology Solution

Cisco IT’s overall technology solution includes:

- Individual workstations, each with a Cisco IP Phone 7960 enabled with Extension Mobility, which allows any Cisco employee to log onto a phone and personalize it, and Cisco VT Advantage, which provides video telephony capability
- Cisco MeetingPlace for audio, video, and Web conferencing
- Cisco IP Communicator with wired or wireless headsets, to turn any laptop anywhere into a corporate IP phone
- Portable videoconferencing units in several meeting rooms
- Interactive white boards, instant messaging, e-mail, and voice mail



RICH GORE is an IT program manager with the Cisco IT@Work group. He can be reached at rgore@cisco.com.



CISCO CONNECTED WORKPLACE A new workplace environment provides space for both private and collaborative work activities.

By moving to the wireless infrastructure, the resulting capital spending reductions have included a 40 percent reduction in the number of switches and switch ports, a 60 percent reduction in the number of wired IP cables required per workstation, and a 50 percent reduction in equipment room space due to racking fewer switches.

Because the enhanced workplace environment has the flexibility to support any employee working anywhere in the workplace, more than 140 employees can work in a space that is normally allocated to 88 employees in traditional cubes, reducing overall costs while increasing employee satisfaction and productivity.

For more information about the Cisco Connected Workplace and other Cisco IT technology efforts and solutions, visit Cisco IT@Work at cisco.com/go/ciscoitwork. ■

FURTHER READING

- Cisco Voice and IP communications
cisco.com/packet/173_4c1
- Cisco Wireless Solutions
cisco.com/packet/173_4c2

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 Locking Down an Access List by Turning the Router into a Packet Sniffer

Many applications need access through an access control list (ACL) on a router, but you might not know what ports or protocols to allow in the ACL rules. I use an access list and the router's logging function to "sniff" and report what it sees. I then take a transparent approach to locking the ACL down to only what is needed. The process takes some time, but is effective and can be used on virtually any interface ACL. Follow these steps:

1. Discover: Use the router's ability to log ACL matches in its own buffer in order to catalog the traffic that crosses it. The configuration is the following:

Router (config)# logging buffered 15000 (this creates a large enough buffer to look at locally on the router, or you can configure the router to log the ACL matches to a Syslog server).

```
Router (config)# access-list 101 permit tcp any gt 0
any gt 0 log
Router (config)# access-list 101 permit udp any gt 0
any gt 0 log
Router (config)# access-list 101 permit icmp any any
Router (config)# access-list 101 permit ip any any log
(this entry is a "catch-all")
Router (config)# interface interface
Router (config-if)# ip access-group 101 in
```

Look at the log by using the **show log** command from the exec prompt. You should see IP addresses (source and destination), along with the used TCP or UDP ports (in parentheses):

```
Mar 18 20:05:10.628: %SEC-6-IPACCESSLOGP: list 101
permitted tcp 192.168.19.137(50051) ->
10.2.9.30(15648), 1 packet
Mar 18 20:05:20.697: %SEC-6-IPACCESSLOGP: list 101
permitted tcp 192.168.19.137(50054) ->
10.2.9.30(15648), 1 packet
Mar 18 20:05:30.757: %SEC-6-IPACCESSLOGP: list 101
permitted tcp 192.168.19.137(50057) ->
10.2.9.30(15648), 1 packet
Mar 18 20:05:40.854: %SEC-6-IPACCESSLOGP: list 101
```

```
permitted tcp 192.168.19.137(50060) ->
10.2.9.30(15648), 1 packet
Mar 18 20:05:51.006: %SEC-6-IPACCESSLOGP: list 101
permitted tcp 192.168.19.137(50063) ->
10.2.9.30(15648), 1 packet
Mar 18 20:06:01.115: %SEC-6-IPACCESSLOGP: list 101
permitted tcp 192.168.19.137(50115) ->
10.2.9.30(15648), 1 packet
Mar 18 20:06:10.354: %SEC-6-IPACCESSLOGP: list 101
permitted tcp 192.168.19.137(50118) ->
10.2.9.30(15648), 1 packet
Mar 18 20:06:20.423: %SEC-6-IPACCESSLOGP: list 101
permitted tcp 192.168.19.137(50121) ->
10.2.9.30(15648), 1 packet
```

When using earlier IOS versions you can specify **permit ip any any log** or **permit tcp any any log**. This shows you the port numbers in the **show log** command. However, with later IOS versions, using these permit statements produces a port zero (0) in the **show log** command, which is why I use statements such as **permit TCP any goto any goto log**. This produces the TCP port numbers. The same goes for UDP.

In this example, two IP addresses need to communicate, and the destination TCP port is 15648. The source port changes, so I cannot create a static ACL for that port. I can create a TCP ACL that states:

```
access-list 101 permit tcp host 192.168.19.137 host
10.2.9.30 eq 15648
```

This should be sufficient for a tight access list. Using this information, I can create a new access list.

2. Notify: Even though we do not expect any connectivity issues, notify those who could be affected by an error or an unexpected condition during this process.

3. Implement: We do not want to interrupt a production environment to make changes. I used this method:

```
Router (config)# interface interface
Router (config-if)# no ip access-group 101 in
(Removes the ACL from active service)
Router (config-if)# exit (I could use Ctrl+Z here as well)
Router (config)# exit
Router# show running-config
```

(I do this every time, because the router will not error out if you mistype the **no ip access-group 101** in line. Verify that it is removed!)


```

Router# ena
Router (config)# no access-list 101 (Now go and
delete the ACL.)
If you use "named" access lists, you can edit out
individual access-list entries. I use Microsoft
Notepad to edit the access list from the configuration
by pasting it into the Telnet/SSH session. Next, I
add the new access list:
Router (config)# access-list 101 permit tcp host
192.168.19.137 host 10.2.9.30 eq 15648
Router (config)# access-list 101 permit tcp any gt 0
any gt 0 log
Router (config)# access-list 101 permit udp any gt 0
any gt 0 log
Router (config)# access-list 101 permit icmp any any
Router (config)# access-list 101 permit ip any any log
Router (config)# interface interface
Router (config-if)# ip access-group 101 in

```

4. Verify: Use the `show access-list` command and the `show log` command to verify that the access list is working.

As the more specific access-list entries are added, the more general ones are pushed to the bottom of the list, and eventually the hit counters on the general ACL entries will be zero. Then you can delete those general ACL entries, leaving only the specific rules.

The sequence of commands I used is important in retaining a transparent change to a production environment.

—Brian Murphree, CCNP, Nashville, Tennessee, USA

Editor's Note: This tip recommends using an access list with the log keyword to discover which traffic goes through. Cisco recommends using NetFlow instead.

TIP Getting a MAC Address from an IP Address and Vice-Versa

MAC from IP: MAC is Layer 2 information. We could get this from a Layer 2 device; for example, Cisco Catalyst switches such as the 1900, 2900, 3500, 4000, and 6000 series, etc. The command for this is:

```

#show cam dynamic port 3/2 (In Catalysts running
CatOS)
#show mac-address-table dynamic interface fa3/2
(In Catalysts running IOS)

```

The above helps bind the MAC address to the port. Using this, we retrieve the IP address. IP is Layer 3 information. We could get this from a Layer 3 device; for example, Cisco routers such as the 1700, 2600, 3700, etc., or from the Route Switch Module (RSM)

of a Layer 3 Switch (4000, 6000, etc.). The command for this is:

```
#show arp | include <mac-address>
```

The `<mac-address>` should be in the format `<xxxx.xxxx.xxxx>`. Using this, we get the IP address related to the MAC address.

IP from MAC: First log in to the RSM or the router:

```
#show arp | include <ip-address>
```

The MAC address related to the IP address will be listed. Select the appropriate MAC address related to the IP address. Log in to the Catalyst switch:

```
#show cam dynamic <xx-xx-xx-xx-xx-xx> (In Catalysts
running CatOS)
```

```
#show mac-address-table dynamic address
<xxxx.xxxx.xxxx>
```

In Catalysts running IOS the above would list the port and the MAC address that is bound to that port.

—Gireesh Hariharasubramani, US Technology Resources LLC, Trivandrum, India

TIP Show Controller Utilization

The `show controllers utilization` command on the Cisco Catalyst 2950 (other IOS switches probably support this) running Cisco IOS Software Release 12.1.22+ shows a summary of port utilization on all ports, including percentage and backplane/fabric utilization:

```

2950#show controllers utilization
Port      Receive Utilization  Transmit Utilization
Fa0/1      0                    0
Fa0/2      0                    0
Fa0/3      0                    0
Fa0/4      0                    0
< -- truncated -- >
Total Ports : 50
Switch Receive Bandwidth Percentage Utilization : 0
Switch Transmit Bandwidth Percentage Utilization : 0
Switch Fabric Percentage Utilization : 0

```

This can also be done on a per-port basis using `sh controllers FastEthernet 0/1 utilization`.

—Aurelio DeSimone, CCIE No. 10267, Refco Group, Inc., Chicago, Illinois, USA

Troubleshooting

TIP Using Cisco BERT to Monitor Link Quality

Cisco bit-error-rate testing (BERT) is a very helpful utility for monitoring link quality; just loop the suspected media at the far end, configure a BERT profile, and run a test for the needed duration. The user profiles are stored as part of the configuration in the NVRAM. We can define a maximum of 15 profiles on the system. This provides the results in terms of pass or fail. You can use the tool can to remotely verify your media from the command-line interface (CLI) and quickly isolate media problems.

```
!  
<Define BERT Profile>  
bert profile 1 pattern 211-0.152 threshold 10^-2  
error-injection none duration 60  
!  
<Run The Test>  
Router#bert controller e1 0 profile 1  
<Output>  
Router#show controllers e1 bert  
Controller E1 0 Profile default : Test Never Ran  
Controller E1 0 Profile 1 : Test passed with BER of  
10^(-2) --> [wow Media is OK till the looped end]
```

```
Controller E1 1 Profile default : Test Never Ran  
Controller E1 1 Profile 1 : Test Never Ran  
Controller E1 2 Profile default : Test Never Ran  
Controller E1 2 Profile 1 : Test Never Ran  
Controller E1 3 Profile default : Test Never Ran  
Controller E1 3 Profile 1 : Test Never Ran  
Controller E1 4 Profile default : Test Never Ran  
Controller E1 4 Profile 1 : Test Never Ran  
Controller E1 5 Profile default : Test Never Ran  
Controller E1 5 Profile 1 : Test Never Ran  
Controller E1 6 Profile default : Test Never Ran  
Controller E1 6 Profile 1 : Test Never Ran  
Controller E1 7 Profile default : Test Never Ran  
Controller E1 7 Profile 1 : Test Never Ran
```

—Sheeraz Ahmed, Supernet Limited, Karachi, Pakistan

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.

Tech Tips

Configure route-maps for IP routing protocol redistribution. Learn how to use commands for configuring route-maps that are applied with the **redistribute** command of dynamic routing protocols. Get help with applying the route-map function to different tasks, such as policy-based routing (PBR) and Border Gateway Protocol (BGP) neighbor update modification. cisco.com/packet/173_4e1

Optimize security for Cisco ONS 15454 Release 5.0. This document describes security considerations for Cisco ONS 15454 Release 5.0 in secure operating mode. Together with timing, communications, and a Control Version Two Plus (TCC2P) card, you can provision the two management LAN ports either with independent IP/MAC addresses for additional network security and segregation, or with a single IP/MAC for simplicity. cisco.com/packet/173_4e2

Use CiscoWorks VMS to authenticate with Cisco Secure ACS. This document provides a sample configuration for enabling CiscoWorks VPN/Security Management Solution (VMS) to authenticate with Cisco Secure Access Control Server (ACS). cisco.com/packet/173_4e3

Configure VPN with Cisco PIX Security Appliance 7.0. Get help configuring LAN-to-LAN sessions between PIX firewalls. This document provides a configuration for static and dynamic LAN-to-LAN tunnels with spoke-to-spoke connectivity through the hub PIX firewall. cisco.com/packet/173_4e4

Upgrade Cisco PIX Security Appliance 7.0 and ASDM Software. Learn about upgrading the PIX appliance from version 6.3 or 6.2 to version 7.0. This document also covers the installation of Cisco Adaptive Security Device Manager (ASDM) version 5.0. cisco.com/packet/173_4e5

Statistics Versus Dashboards

A Practical Guide to VoIP Monitoring and Management

By Dave Chapman

Everything moves in cycles, especially the technology industry, which experiences swings between boom and bust. Technology cycles between distributed networks and centralized computing, databases and storage-area networks (SANs), and wireline and wireless. At the same time, new technologies are constantly emerging and receding, and as technologists know, history repeats itself.

In few technology areas is this phenomenon more apparent than within network management. When data networks took off in the late 1980s, it was not long before data network management tools were a necessity in the enterprise. When SANs took off in the late 1990s, SAN management tools also took off.

With this kind of recent history, it is surprising that many IT executives and CIOs are still coming to terms with the idea that voice over IP (VoIP) phone systems and IP communications networks also need to be managed, and that many organizations are taking the wrong approach to managing their new systems.

IT Investment in EMS

IT managers have invested billions of dollars in enterprise management systems (EMS) to manage their networks, servers, applications, and other infrastructure. Products such as Tivoli, Openview, BMC Patrol, CA Unicenter, CiscoWorks, and many others fill this critical need.

Yet while some EMS implementations pay big dividends, others are a colossal waste of time and money, even when they use the same software. How does an IT team tasked with implementing management tools avoid making the same mistakes with IP communications networks? What will make the difference between huge return on investment (ROI) and a management money pit?

Begin at the Whiteboard

The difference between success and failure begins early—not after the network is deployed, but at its inception—when it's being architected on a whiteboard, long before equipment is specified or ordered. The more critical the enterprise system, the more careful the architecture, design, and implementation must be. And few applications are more critical than a company's phone system.

Yet time and again, EMS's are deployed only after operations organizations realize that they cannot maintain what they have implemented. Managers

rushed to buy whatever solution they could get their hands on, often under pressure. "Requirements," if they were used, were often very technical in nature and not focused on business objectives.

Even more amazing is that most IT managers understand the need to implement management early, yet still end up buying only when they face a crisis.

Management Requirements

In general, enterprise systems, such as e-mail, enterprise resource planning (ERP), payroll, websites, networks, and others are most successful when they are implemented with a careful, top-down, requirements-based approach.

But whose requirements? Ask network engineers what they want to know and you might be tempted to base management requirements on technical questions such as CPU utilization of the CallManager at 5-minute polling intervals, buffer utilization of the core routers, or packet loss on all WAN circuits.

Ask your handset users "*What are the most important aspects of a VoIP phone system?*" and you are likely to hear something like, "The phone must work when I pick it up," "Customers have to be able to reach me," "It has to sound good—certainly better than my cell phone," etc.

Ask the CIO and you will probably hear still different answers:

- It must have high uptime, approaching five nines, just like the old system.
- It must cost the same or less to operate than the old one.
- My staff has to be able to troubleshoot it quickly and easily.
- The users have to like how it sounds.
- It must support growth over the next 10 years.
- The call center must be equipped to handle the holiday rush.

So who is right? Actually, they all are. But ultimately the users and the CIO matter most, not the technical engineer.

All the requirements are important, but when you examine them in depth, all have complex implications. Thus the management tool must take complex statistics and turn them into easy-to-understand, actionable information.

The above requirements tend to revolve around four main issues: *functionality*, *reliability*, *quality*, and *cost*—the primary tenets of any enterprise application. Functionality is not something that an EMS can improve; it typically determines what solution a company purchases. The last three however are the mainstays of EMS. If your EMS can help improve reliability and quality while lowering cost, you are on the right track.

And what you *don't* hear as requirements in the above list is as important to understanding the state of VoIP monitoring and management as what you *do* hear.



DAVE CHAPMAN is president of Chapman Consultants, LLC, a technology consultancy based in Maryland. He is a former director of product management at Cisco partner Qovia, a voice monitoring and management company. He can be reached at dave@chapmanconsultants.com.

Users aren't asking for statistics such as "*packet loss*," "*CPU utilization in the call server*," or "*call completion rate*." Yet that is how many VoIP management systems and their implementers approach the problem of hundreds of performance statistics that have no business context. Just like the early versions of today's EMS's.

Action Heroes

Information that is not actionable is useless. It is just data. But information that works for you is priceless. So management tools need to address requirements by providing the right information at the right time to the right team members—so that they can keep a network running at its best.

In other words, the IT team must be able to measure success and failure, as well as be able to act on it when something goes amiss. And the team needs to know immediately if something is happening. "Hey, if I can't tell quality is dropping now, I can't fix it now!"

Decimal Places Versus Dashboards

One way to think about a management tool is to consider the dashboard of your car. The average driver uses the speedometer, tachometer, and dashboard lights to assess how a car is running. They don't need a complex instrument panel showing details like

Ad

exhaust gas mixture; fuel/air ratios; fan speed; and water temperature to five decimal places. They need a simple amalgam of information based on these statistics. Save the overwhelming details for the mechanic.

The same applies to your VoIP monitoring and management system. The system needs to tell you in plain, simple language that you are doing fine, or that you have a problem. In fact, in some areas, VoIP management systems need to do one better because as it turns out, idiot lights do not begin to help you fix the problem. Therefore, you have to strike a balance between the complex statistics that help you troubleshoot and the idiot lights that make things simple and can send you instant alerts. Like a car dashboard, the management system must notify you the instant something goes wrong.

Doing one better requires you to extract business-level impact from the technical metrics that are used to characterize behavior. You must create a mapping from each user and business requirement down to the features and technical metrics of the monitoring and management system. For features and metrics that cannot map to a business requirement, ask yourself why you have them. In fact, if you start with product features and work your way up, you will also fail, as was the case with many an EMS implementation for data networks.

For a concrete example, let's examine the user requirement that it *"has to sound good,"* a cousin to the CIO requirement that *"users better be happy with the way it sounds."* This is a subjective requirement that changes from one person to the next. For decades phone companies and vendors have relied on the Mean Opinion Score (MOS) to determine if a user would be happy with the way a call sounds. Dozens of algorithms, patents, and standards are dedicated to measuring this one metric both in the VoIP arena as well as the traditional time-division multiplexing (TDM) phone systems. Many of these measure MOS out to several decimal places.

So, you get a system to track MOS scores, and it sends an alert when the score misses a MOS score target. *"Average MOS of 3.3473 reported on a phone call."* Realistically though, users don't even know what a MOS score is, much less care about 4 decimal places of accuracy. In fact, it is likely that the CIO might not care either. What is important is that the operations organization can tell when users are unhappy so they can fix it.

Unfortunately, even if the previous alert is accurate, it is not all that helpful to the operations organization either. First, what does 3.3473 mean? Is it good? If it is poor, why is it poor, and what should I do about it?

In reality, the alert generates more questions than it answers, and after enough of them, the network operations center (NOC) will ignore them, and rightly so. We need more business context to help prioritize, and more detail for the NOC to take action.

So, go back to the user requirements. There was nothing in their requirements about a MOS of a certain score; it just has to *"sound good."* What may be more effective is detecting a change in the relative MOS between calls, or even within a call. For example, if the average MOS is 4.0 and it suddenly drops down to 3.1, it is likely that users will be upset about quality. A good analogy is your expectation about your cell phone. You might not care much if it does not sound as good as your land line, but it is very annoying if quality drops as you drive down the road.

Next, who was affected by the change? Was it a call between two people? Was it all calls to a site? Was a customer involved? This kind of information helps the NOC prioritize what problem to work on next.

To take action, the NOC must get to the cause of the problem. In an IP network, several metrics can have a direct impact on a MOS score. It is helpful for the NOC to know which of those contributed to the low score.

In the end, you might expect a message such as *"Call Center agent 37 experienced five minutes of poor call quality for a call inbound from 800 555 1212."*

The operations organization knows there is a problem because the message clearly says the quality was poor. They know it is important because it is a call center agent and they were talking to a person who called the catalog number.

Next on the NOC's list is the troubleshooting. Like the mechanic, they will drill down within the system, looking for the other details that tell them why the "check engine" light is on. For example, the actual MOS was 3.1 when the average for the other calls is 4.0, packet loss was 10 percent, when it is normally "0," and the call came in on PSTN gateway 27, on the first T1, Channel 6. These are the detailed statistics that are appropriate to the problem. The NOC will also want to look at other calls to determine if the problem was isolated or not.

By recognizing that management tools implementations should be designed using a top-down requirements analysis, and that they must strike a balance between statistics and a dashboard, provide actionable information for the IT team, and be designed into a network during the design and pilot phase, history won't repeat itself for VoIP networks—at least as far as management is concerned. ■

Internet to Orbit

Exploring the Potential of IP Networking Technologies in Space

By Daniel Floreani and Lloyd Wood


On October 4, 2004, Burt Rutan's SpaceShipOne rocketed into history, becoming the world's first private manned spacecraft to reach the edge of space by flying to an altitude of over 100 km (climbing to just over 70 miles up). The builders of SpaceShipOne, backed by Microsoft cofounder Paul G. Allen, were awarded the US\$10 million Ansari X Prize by the X Prize Foundation (xprizefoundation.com), a not-for-profit organization that uses competitions to encourage innovative breakthroughs in private space flight. The X Prize Foundation intends to jumpstart the personal spaceflight industry through competition between the most talented rocketry experts and entrepreneurs in the world.

Once a top-down venture driven by governments and based on highly customized, costly components and systems, space exploration is finally becoming a free-market endeavor. With governments no longer able to afford or justify unrestricted budgets, and with a new generation of space enthusiasts mastering the principles of space technology and rocket propulsion, the commercial market has stepped up its involvement in space activity.

This new generation of deep-pocketed space entrepreneurs includes Richard Branson of the Virgin Group, who has teamed with the prize-winning flyers of SpaceShipOne, Paul Allen's Mojave Aerospace Ventures, to create Virgin Galactic, which will build five spaceliners to bring ordinary (though wealthy) citizens into space. Elon Musk, a cofounder of PayPal, the electronic payment system, has established SpaceX to develop a family of launch vehicles intended to reduce the cost and increase the reliability of access to space. The first two launch vehicle designs from SpaceX, Falcon I and Falcon V, are rockets capable of placing approximately 670 kg and 6020 kg, respectively, into low Earth orbit (LEO). Falcon I will be priced at US\$5.9 million per launch, plus payload-specific costs and range-related fees, making this vehicle a cost-effective ride to orbit.

To help lower the costs of his launch rockets, SpaceX's Musk has turned to an old standby in the world of standards-based networking. In a 2003 *Wired News* article, Musk said, "In launch vehicles, for communications you typically have bundles of serial cables that are as thick as someone's arm. We thought that made no sense, so we put in an Ethernet system."

It's no coincidence that Musk is on the board of the X Prize Foundation.



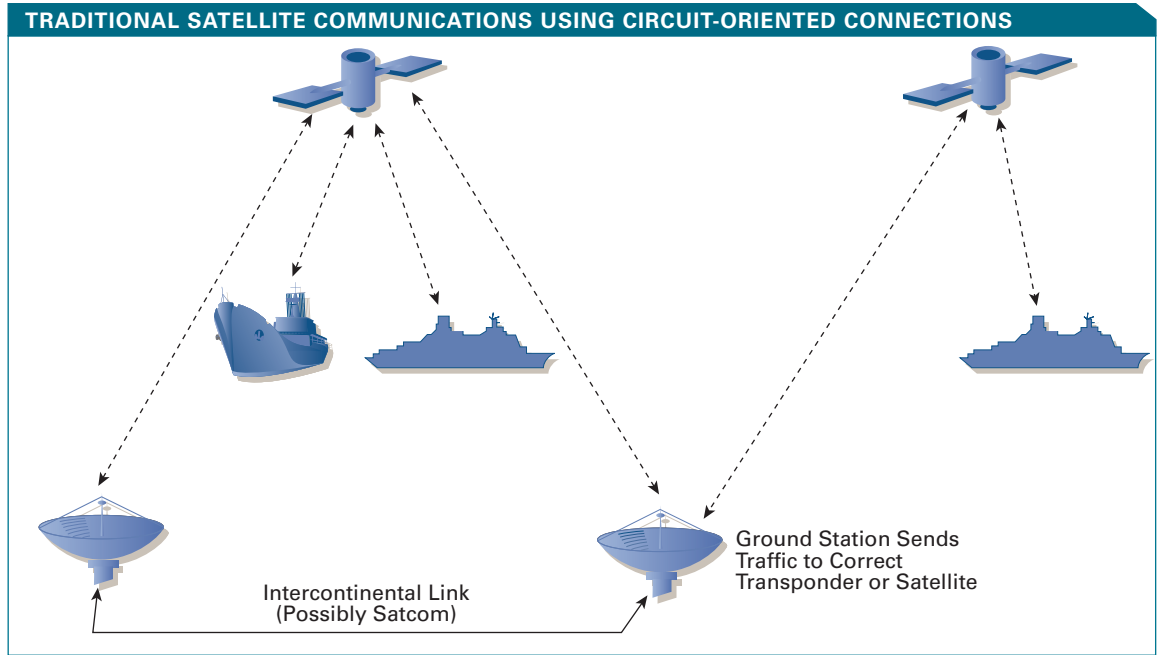
Surrey Satellite Technology Ltd.

UK-DMC The UK-DMC and NigeriaSAT-1 Disaster Monitoring Consortium satellites enter vacuum chamber testing at Rutherford Appleton Laboratory.

Cisco and Space

Cisco has been eyeing space for some time, and in September 2003 a Cisco 3251 Mobile Access Router was launched into low Earth orbit as a secondary experimental payload onboard the UK-DMC (Disaster Monitoring Consortium) satellite built and operated by Surrey Satellite Technology Ltd. (SSTL) of Guildford, England. The UK-DMC is part of a "constellation" of five LEO satellites capable of providing large multispectral images of the Earth's environment for international disaster monitoring and a variety of civil and commercial uses. The satellites and

FIGURE 1 In the “bent-pipe” architecture, banks of transponders receive an RF signal, amplify or regenerate it, then send a lower-frequency signal back to Earth.



ground stations together form a worldwide IP network that extends into space.


This project marked the first time that commercial Cisco hardware was taken to and tested in space. According to Rick Sanford, director of the Cisco Global Space Initiatives group, the venture’s ultimate aim is to help lower the cost of building satellites while improving their communications capabilities with already-established data networks on earth, particularly those using Internet-based communications.

“Typically, satellites are a ground-up effort,” says Sanford. “We want to see if we can reduce the cost of satellite communications by applying the same open standards and commercially-produced equipment that have been used to build the Internet to satellites.”

In order for this new generation of satellite manufacturers and launch companies to drive down costs, shrink project timetables, and increase the flexibility and capabilities of their hardware, it is necessary to adopt open standards of increasing capability and intelligence. Many of these new satellites will use best-of-breed commercial-off-the-shelf (COTS) technologies

DANIEL FLOREANI is a network architect for the space team in Cisco’s Global Defense, Space and Security group. He holds a PhD in modeling tactical packet radio networking from the Institute for Telecommunications Research at the University of South Australia and can be reached at danielf@cisco.com.

LLOYD WOOD is a space initiatives manager in the Global Defense, Space and Security group, where he is responsible for Cisco’s first router in orbit. He holds a PhD in internetworking satellite constellations from the Centre for Communication Systems Research at the University of Surrey in the UK. Lloyd Wood can be reached at lwood@cisco.com.



Surrey Satellite Technology Ltd.

CLEO The Cisco router in Low Earth Orbit assembly sits in the left half of its payload tray.

and subsystems; this will help enable more rapid on-demand deployment of satellite infrastructure.

Standards-Based Communications

Standards-based communications is an important element in the drive to expand the market for satellite services. With IP as a standard protocol for communication onboard satellites and between satellites and ground stations, space-based communications will be able to migrate towards the characteristics of ground-based communications. Unlike today, where satellite service providers maintain their own services and links, and each link can be dedicated to a single service—paging, landline phone, portable phone, or

television—IP enables convergence. Service providers will be able to use one IP-based connection to carry multiple services and support many different applications. For example, a single space-based Internet connection could simultaneously support audio streaming, telephony, and data services. This could replace the need for separate dedicated hardware for satellite radio, TV, and telephone services.

An IP-based network infrastructure would also simplify the ability to use the service, dramatically expanding the market. Today, choosing a satellite service provider is a lengthy, complex process. The user, typically a large company or a government or military organization, must consider the following in its selection process:

- **Coverage area**—This can range from an area as small as a state in the US (or smaller) based on spot beams, to a coverage area that can extend to almost one-third of the Earth. Typically, the larger the coverage area, the smaller the available shared link capacity per user.
- **Service requirements**—Services are typically fixed or mobile. Mobile services require portable or mobile dishes often less than 1 meter in size. Fixed ground stations can be larger, with dishes from 2.5 meters to 8-meter dishes, and often cost millions to install and operate. Smaller dishes have been known for decades as very small aperture terminals (VSATs). The other service consideration is the radio frequency used, including C, X, Ka, Ku, and S-bands, to name but a few. Each band has different physical characteristics, such as susceptibility to rain fade and ability to penetrate clouds and foliage.
- **Available bandwidth**—This is determined by terminal sizes and the radio frequencies they use. Truly mobile terminals (less than 30 cm in size) use the newest technology and can typically provide 64-kbit/s bidirectional speeds. Portable VSAT terminals typically provide up to 512 kbit/s, but more often provide between 64 and 256 kbit/s for cost reasons. Trailer-based terminals (2.5m dishes) can provide 1 to 2 Mbit/s, and larger static dishes can provide 8 to 60 Mbit/s and higher.
- **Radio frequencies and medium-access mechanisms**—Today, medium-access mechanisms and frequency-reuse options that get the most from a shared satellite link are largely proprietary, with little ability to reuse equipment between two providers, or have equipment from multiple providers share a satellite link. As users of satellite services are increasingly requiring symmetrical communications to handle returning data or video, as opposed to asymmetric broadcast services today, more sophisticated allocation mechanisms for capacity reuse are required to support these demands and get the most out of the satellite link.

Frequency Reuse and Multicast with Spot Beams

Satellites are used extensively today for television or for satellite radio, via XM or Sirius Radio. These broadcast services are generally engineered to cover large footprints to reach many users. The challenge is that if one large footprint is also used to receive return channel traffic, as becomes increasingly likely with more users transmitting data, it creates contention for shared channel access, limiting the performance of VSAT satellite networks using many terminals.

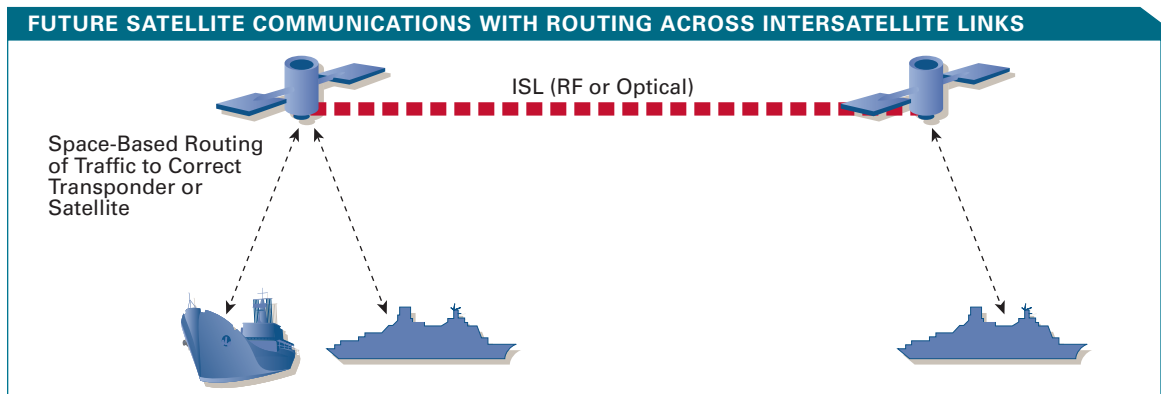
One solution to this problem that engineers have created is called using spot beams. Similar in concept to cellular systems on Earth, spot beams are smaller areas (up to the size of a US state or smaller) that can reuse available frequencies also used in other spotbeams for distant areas and thus support more users and more return traffic per beam. Introducing spotbeams changes the nature of broadcast, and moves broadcast from the satellite to each broadcast beam.

Using beam capacity efficiently for multicast becomes challenging. With IP multicast, multicast should only be sent to the spot beams that have active clients present who have registered for a specific IP multicast group. With IP routing functionality onboard the satellite, it becomes possible to replicate packets onboard the satellite for individual spotbeams as required, and switch traffic directly between spotbeams so that clients sharing a group can communicate. This is preferable to replicating packets on the ground to send up the feeder link.

Routing functionality enabling multicast onboard can enable a service providers to make more revenue from finite link capacity, because packet replication for the spotbeams can be done on the satellite, rather than on the ground. This saves feeder uplink capacity and also decreases the delay in many-to-many multicast communications, because there is no need to return to the ground to switch traffic between group members. However, enabling this flexibility for IP multicast requires integration of IP and Layer 2 switching onboard the satellite, combining the routing and switching functionality with the spot beams.

Setting up a satellite link is clearly not straightforward. Compare this process with that of using Internet-based connectivity, where the user chooses a suitable connection speed from a local Internet service provider (ISP). Key challenges must be overcome before satellite communications will be as easy to use as networking on Earth. The most significant barrier is the lack of common standards for medium access and frequencies. Companies must work together to build standards-based technologies that allow multiple systems to work together seamlessly. Once this standardization has occurred, satellite users should not have to worry about how a ground terminal at the other end of the link has been configured for access, or where or how their traffic will come back to Earth. IP packets could even be routed in space, as well as on the ground, to reach their correct destinations.

FIGURE 2 With routing functionality on every satellite and interconnections between satellites, each ground node needs only one last-hop communications link for connectivity to the network “cloud.”



Just as the terrestrial Internet is migrating from a best-effort IP network to a network supporting distinct classes of service, based on packet quality-of-service (QoS) markings, satellite-based networks also require QoS to provide the levels of quality that different services need to perform as desired. For this to happen, satellite networks must share a common understanding of QoS between satellite terminals and other networking equipment on the ground. Satellite modems must be QoS-aware and condition traffic (with preferential queuing and other QoS techniques) so that they can map QoS onto medium access functions before sending out the frames, enabling them to prioritize traffic based on the QoS marking of individual packets. The separation of modem and routing functionality and use of encryption makes this shared understanding and mapping of QoS hard to achieve.

Reduced Frequency, Faster Transmission

Satellite service providers have seen their ground-based equivalents migrate to IP-based converged networks, and they are beginning to express a strong interest in using IP networking for space to realize the same significant bandwidth, management, and operations savings that have been seen many times on the ground. A satellite Internet ISP wants to use the same mature network management tools that a cable ISP uses.

Traditionally, communications satellites have been support platforms for banks of transponders that receive a radio-frequency (RF) signal, amplify or regenerate it, then send a lower-frequency signal back to Earth. This has produced the familiar “bent-pipe” architecture (see Figure 1 on page 20).

In this example, two navy ships on shared operations, but from separate countries, are communicating by phone. But because each ship uses its own country’s proprietary satellite (or transponder) for its communications, connectivity can only occur indirectly, between the two gateway ground stations. Rather

than a signal traveling directly from ship to satellite to satellite to ship, the signal travels a long, latency-inducing, capacity-consuming trip: ship to satellite, down to ground station, across an intercontinental link, to ground station, up to satellite, and finally down to the other ship. Similarly, traffic between different transponders on the same satellite would need to be switched on the ground.

If routing and switching functionality was deployed onboard satellites with interfaces to standardized intersatellite links (ISLs), then this voice-over-IP (VoIP) traffic could feasibly be transmitted directly from one satellite to the other in space, removing the delays and other overheads associated with the extra hop and decreasing reliance on terrestrial infrastructure (Figure 2).

Newer satellites are already being built with switch fabrics onboard that allow data frames to be sent directly from one transponder to another without returning to the ground for switching. Examples include the AmerHis satellite from Hispasat, a leading supplier of space segment capacity in Europe and the Americas. The Intelsat 9 series of satellites have a reconfigurable switching matrix between transponders to increase overall flexibility. The Japan Aerospace Exploration Agency (JAXA), is building the Wideband InterNetworking engineering test and Demonstration Satellite (WINDS), with an onboard cell switch that can operate at 155 Mbit/s.

These switch fabrics, however, are still managed by ground-based controllers. When routing functionality is included onboard satellites it can extend switching to true Layer 3 IP routing onboard and even between satellites, allowing the satellites to actively recognize and act on the individual IP traffic that their transponders carry, rather than just cross-connect transponders.

Ultimately, with satellites using more onboard communications technology, new architectural approaches to satellite communication can evolve based on smaller, lower-cost satellites flying in single geostationary cluster formations using ISLs. In time, one or more “core” satellites in each cluster could provide connectivity to other clusters. These “core” satellites should be able to connect clusters together to form a geostationary “backbone” ring of high-speed, free-space optical links, similar to fiber-optic networks in the terrestrial Internet.

New Potential with IP and Clusters

The potential benefits of interconnected clusters of satellites, based around IP standards and using inter-satellite links, include allowing the routing decision for packet contents to be guided by onboard functionality that is driven by the following:

- Knowledge of available connectivity to transponders, local payloads, other interconnected satellites and ground stations
- Topology of the space-based network and the packet markings
- Policy and administrative decisions covering link costs, delay tolerance, security policies, and service-level agreements (SLAs).

When this network-level consensus point is reached, onboard satellites, interoperability between different families of payloads and transponders intended for different purposes becomes possible. The functionality and flexibility of all payloads and transponders becomes available across all interconnected payloads in the geostationary satellites holding position in their orbital slots. This would be subject to agreed interoperability and peering administrative agreements at the policy level.

One World, One Network

Developing a standards-based infrastructure in space will not only allow users and service providers to change the way that terrestrial devices communicate using satellites, but will also revolutionize how data is stored, disseminated, and used. A satellite-based IP communications network will help extend services globally to users in any region of the world. This is important to governments that are interested in using satellite technology to provide universal service to

their rural citizens, who currently do not have access to the broadband Internet. This will improve rural quality of life, just as government programs providing countrywide postal and electricity services do today.

In time, users will get cheaper and ubiquitous access to space services. In fact, new merged space-ground architectures will be created where users should not need to know or care how their packets are routed around the Earth.

Our vision is that one day, each and every manned and unmanned spacecraft or airframe will be a network carrying active nodes on the Internet. One world, one network, with terrestrial and space-based communications together forming a seamless universal service. ■

FURTHER READING

- Cisco Space Networking
cisco.com/go/space
- Cisco Global Defense, Space and Security (GDSS)
cisco.com/go/gdss
- Cisco Router in Low Earth Orbit
[ftp://ftp-eng.cisco.com/lwood/cleo/README.html](http://ftp-eng.cisco.com/lwood/cleo/README.html)
- Suneel Ratan, “Net Maverick Sets Sights on Space,” *Wired News*, April 22, 2003.
wired.com/news/business/0,1367,58493,00.html
- Will Ivancic et al., “Secure, Network-Centric Operations of a Space-Based Asset: Cisco Router in Low-Earth Orbit (CLEO) and Virtual Mission Operations Center (VMOC),” NASA Technical Memorandum 2005-213556, May 2005.
- Keith Hogue et al., “Using Standard Internet Protocols and Applications in Space,” *Computer Networks* 47 (2005) 603-650, Elsevier.
- Lloyd Wood et al., “Adopting Internet Standards for Orbital Use,” paper SSC05-IV-03, 19th Annual AIAA/USU Conference on Small Satellites, Logan, Utah, August 2005.
- Lloyd Wood et al., “Slot Clouds: Getting More from Orbital Slots with Networking,” paper IAC-03-U.4.07, 54th International Astronautical Congress, Bremen, Germany, October 2003.
- Joab Jackson, “The Interplanetary Internet,” *IEEE Spectrum*, August 2005.
www.spectrum.ieee.org/WEBONLY/publicfeature/aug05/0805inte.html
- *Packet* article on Rate-Based Satellite Control Protocol (First Quarter 2005)
cisco.com/packet/173_5a1

10 GbE and Its X Factors

Demystifying 10 Gigabit Ethernet Port Types and Pluggable Modules on Cisco Routers and Switches

By Alessandro Barbieri

You have read so much about 10 Gigabit Ethernet (hereafter referred to as 10 GbE), and you are finally convinced that this is the next big thing for your data center. Now you need to get down to the *nitty-gritty* of the technology. Or perhaps you are researching options for your next wiring closet upgrade, or considering 10 GbE to consolidate next-generation metropolitan area network (MAN) requirements. Then again, you might just be a technology enthusiast who likes to stay ahead of the curve.

If any of these or similar scenarios applies to you, this article is meant to help you knowledgeably navigate the labyrinth of 10 GbE technologies and standard port types, as well as grasp the differences, similarities, and applications of the many 10 GbE pluggable modules on Cisco switches and routers.

The IEEE 802.3ae Task Force ratified 10 GbE as a standard in 2002. Although three years in the Internet age could qualify this as an “old” standard, 10 GbE technologies are still in a profound phase of maturation, and the 10 GbE market is still in its infancy.

Before diving into the details of pluggable optics, let’s start with a review of the 10 GbE standards and overview of the most relevant applications for the various 10 GbE standard port types.

10 GbE over Fiber Optics: 802.3ae at a Glance

First and foremost, 10 GbE is still Ethernet, only much faster. Besides raising the speed bar to 10,000 Gbit/s, the main objectives of the IEEE 802.3ae 10 GbE standard were to preserve the Ethernet frame

format, maintain the maximum and minimum frame size of the 802.3 standard and support only full-duplex operation (therefore dropping the requirement for the CSMA/CD protocol). A big portion of the IEEE 802.3ae standard has focused on defining the 10 GbE physical layer.

As a result of the standardization effort, four new optical interface types have been defined to operate at various distances on both single-mode and multi-mode fibers. In IEEE jargon, these are known as physical medium dependent sublayers or PMDs. In addition to these four PMDs, the standard introduces two families of physical layer specifications (PHYs) to support LAN as well as WAN applications. In general, the properties of the PHY are defined in the physical coding sublayer (PCS) responsible for the encoding and decoding functions. Overall there are seven possible 10 GbE port types (see Figure 1).

IEEE 802.3ae PMD Sublayers

The PMD sublayers promoted by the 10 GbE standard can appear confusing. Each PMD has different technical characteristics to support different fiber media and operating distances. In defining the PMD sublayers, the IEEE intended to offer the least expensive optical technology for a particular application:

PMDs for single mode fiber only

- 10GBASE-L operates in the 1300 nm band to support distances to 10 km (6.2 miles).
- 10GBASE-E operates in the 1550 nm band to reach distances to 40 km (24.8 miles).

IEEE 802.3ae 10 GIGABIT ETHERNET PORT TYPES

		PCS		
		10GBASE-R	10GBASE-X	10GBASE-W
PMD	10GBASE-E	10GBASE-ER		10GBASE-EW
	10GBASE-L	10GBASE-LR		10GBASE-LW
	10GBASE-S	10GBASE-SR		10GBASE-SW
	10GBASE-L4		10GBASE-LX4	
		LAN PHY		WAN PHY

FIGURE 1 The IEEE 802.3ae standard introduces four PMDs and three PCS layers for either LAN or WAN applications, producing seven possible 10 GbE port types.

PMD for multimode fiber only

- **10GBASE-S** leverages low cost 850 nm laser technologies and covers distances of 26 to 82 meters on legacy multimode fiber. With laser optimized multimode fibers (known as OM3 fibers), 10GBASE-S operates distances to 300 meters.

PMD for multimode and single-mode fiber

- **10GBASE-LX4** adopts an array of four lasers, each transmitting at 3.125 Gbit/s, and four receivers arranged in wavelength-division multiplexing (WDM) fashion. Working in the 1300 nm region, this PMD supports link lengths of 300 meters on legacy FDDI-grade multimode fiber and distances of 10 km on single-mode fiber.

For more detail on the various fiber types and distances supported, refer to the Cisco data sheets at cisco.com/packet/173_5b1.

IEEE 802.3ae PHY Families

Two new physical layer specifications are part of the 10 GbE standard framework: LAN PHY and WAN PHY. There are also three types of PCS sublayers: 10GBASE-X, 10GBASE-R, and 10GBASE-W. The first two are in the LAN PHY family, and the latter in WAN PHY.

LAN PHY and WAN PHY differ in the type of framing and interface speed. Serial LAN PHY (10GBASE-R) adopts Ethernet framing with a data rate of 10.3125 Gbit/s (the MAC runs at 10,000 Gbit/s; by adding the coding overhead of 64B/66B, the effective line rate becomes $10,000 * 66/64 = 10,3125$ Gbit/s). On the other hand, WAN PHY wraps the 64B/66B

encoded payload into a SONET concatenated STS-192c frame for a data rate of 9.953 Gbit/s.

So Why Do We Need WAN PHY?

SONET/SDH are the dominant technologies deployed in optical transport networks, and thus the traditional optical transport infrastructure is based on SONET/SDH protocols that operate at 9.953 Gbit/s. However, the 10.3125-Gbit/s line rate of LAN PHY does not match the speed of SONET/SDH, and, as such, cannot be transported over WANs based on SONET/SDH. WAN PHY is the IEEE's answer to adapt 10 GbE data rate to SONET/SDH speed.

WAN PHY renders 10 GbE compatible with SONET STS-192c format and data rate, as defined by ANSI, and with the SDH VC-4-64c container specified by ITU. WAN PHY is not strictly SONET compliant. Rather, it is more aptly described as a SONET-friendly variant of 10 GbE. The optical specifications and the timing and jitter requirements remain substantially different from the SONET/SDH protocols.

10 GbE over Copper: 802.3ak at a Glance

In 2004 the family of 10 GbE standard port types welcomed a new member: 10GBASE-CX4. Ratified by the IEEE 802.3ak Task Force, this is the first 10 GbE specification based on a copper interface. CX4 addresses the market demand for very low cost 10 GbE links in applications that do not need the reach of fiber-optics media. 10GBASE-CX4 runs over four pairs of twin-axial copper wiring to 15 meters and adopts the IBX4 connector standardized by the Infiniband Trade Association.

FIGURE 2 The various 10 GbE port types address a variety of enterprise and service provider applications, with the aim of delivering the most cost effective technical solution.

IEEE 802.3ae 10 GIGABIT ETHERNET PORT TYPES MAPPING								
Application Best Fit	Dominant Transmission Media	Typical Application Reach	10GBASE-CX4	10GBASE-SR	10GBASE-LX4	10GBASE-LR	10GBASE-ER	10GBASE-xW
In-Building Horizontal	Copper - Structured Cabling	<= 100 m						
In-Building Vertical (Building Backbone)	Multimode Fiber	<= 300 m			Best Match			
Campus Backbone (Between Buildings)	Multimode Single Mode	<= 2 km			Best Match	Best Match	Best Match	
Data Center/ Server Farm	Multimode/ Copper IBX4 Patch Cables	<= 50 m	Best Match	Best Match				
Metro	Single Mode	<= 80 km			Best Match	Best Match	Best Match	
Transport Over Traditional SONET/SDH Infrastructure	Single Mode	<= 10 km (to the Transport Equipment)				Best Match		Best Match

10GBASE-CX4 will not be the only standard for 10 GbE over copper for too long; in fact, the IEEE Task Force 802.3an is currently working on the 10GBASE-T standard, which will support distances of 55 to 100 meters on Cat6 UTP cabling and 100 meters on Cat7 and “augmented Cat6” or Cat6e cabling. This standard is slated to be ratified in July 2006.

IEEE 10 GbE Applications

As previously noted, the number of port types in the 10 GbE standards is designed to address a variety of applications cost effectively. Each physical layer technology addresses specific market requirements and, within that context, deliver the least expensive technical solution. Figure 2 presents a general mapping of the most common 10 GbE applications, categorized by reach and availability of transmission media, to specific 10 GbE port types. Indeed, port types are not necessarily restricted to the applications shown in this chart.

10GBASE-SR and 10GBASE-CX4 are mostly used in data center or server farm applications, when the limited distance requirements and the flexible choice of transmission media available (copper/fiber) make intra-rack/inter-rack server-to-switch and switch-to-switch interconnects the ideal target for these technologies. Note that CX4 requires IBX4 patch cables and does not support the structured cable paradigm.

The flexibility of 10GBASE-LX4 in terms of supported fiber media and reach makes it the most relevant technology in enterprise campus networks. Although not cost optimized for the data center, LX4 can be used in this multimode, fiber-rich space as well. For these reasons, LX4 has been deemed the “Swiss Army knife” 10 GbE technology for enterprise networks.

With 10GBASE-LR and 10GBASE-ER, we step into the range of applications that require extended distances and single-mode fiber only: campus backbone and metro applications. These two port types are most relevant to service providers for deployments such as metro Ethernet service aggregation, DSL backhaul, and inter-POP connectivity.

Service providers consider WAN PHY for inter-POP or interoffice applications when interface with an optical transport infrastructure based on SONET/SDH for long distances is required.

10GBASE-W is often the only option for enterprises that want to lease from a carrier an OC-192 TDM



ALESSANDRO BARBIERI is a product manager in the Transceiver Module Group at Cisco. He has been at Cisco for more than seven years in various engineering and marketing positions, and can be reached at abarbier@cisco.com.

“X” FORM FACTOR SUPPORTED FEATURES

Supported Features	XENPAK	X2	XFP
All IEEE Port Types	Yes	Yes	No LX4
Non-IEEE Port Types	Yes (80 km/DWDM)	No	No
Size (mm)	126x36x17	100x36x12	78x18x10
Connector Type	SC	SC	LC
Protocols Other Than 10 Gigabit Ethernet	No	No	Yes (OC-192 9.95 Gbit/s, G.709 10.709 Gbit/s)

circuit or an OC-192 lambda for WAN applications spanning countries or even continents.

Let’s now take a look at the three types of Cisco 10 GbE optical transceivers, which represent the real-world incarnation of the port types examined thus far.

Cisco 10 GbE Pluggable Interfaces

Nearly every 10 GbE port on Cisco switches and routers today is sold separately from the optical or copper physical interface. This allows end users to customize the 10 GbE port according to its application requirements by selecting different hot pluggable transceivers.

Cisco 10 GbE transceivers are based on industry-wide standards (known as multisource agreements or MSAs). Cisco standardized on three flavors of 10 GbE transceivers with distinct electrical and mechanical characteristics:

- Xenpak (www.xenpak.org)
- X2 (www.x2msa.org)
- XFP (www.xfpmsa.org)

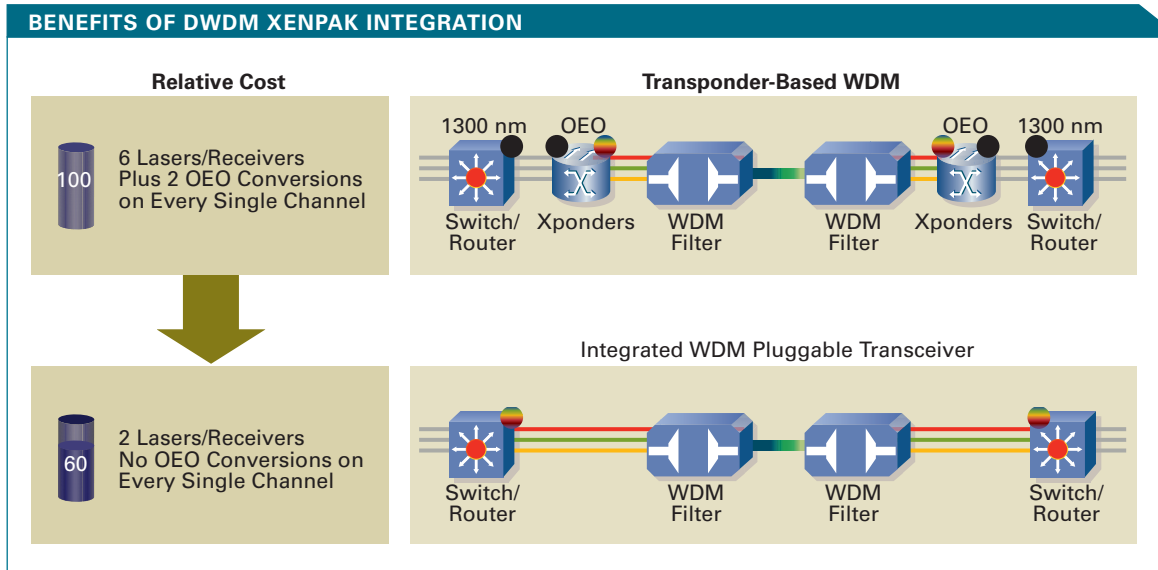
From the end user’s perspective, the most relevant differences among Xenpak, X2, and XFP are found in the mechanical dimensions, and breadth of supported port types and 10G protocols besides 10 GbE (see Figure 3; note that Fibre Channel can technically be supported in all three X modules, but no commercial implementation is yet available).

Xenpak is the largest form factor of the three. But with size comes also a versatility advantage: more room on the module to integrate optical and electrical components as well as higher thermal efficiency, which in turn relaxes the requirement on power dissipation. Consequently, Xenpak supports the deepest range of port types, sporting a power-hungry DWDM interface as well as an 80 km PMD. As the

FIGURE 3 The primary differences among the “X” form factors—Xenpak, X2, and XFP—are lie in their size and supported port types and 10G protocols.

For the full text of Alessandro Barbieri’s paper, *10 Gigabit Ethernet and Its “X” Factors*, see cisco.com/packet/173_5b2.

FIGURE 4 Integrating WDM optics into a switch or router eliminates the transponder from the network.



technology matures and evolves, the existing gap in supported port types is bound to disappear.

Electrically compatible with Xenpak, X2 allows the sharing of board designs and components with Xenpak, and simultaneously small switches with size and thermal constraints to support 10 GbE.

XFP is the smallest of the X form factor pack, partly because it removes much of the electronics from the module, transferring the cost on the host line card. Interestingly, by design XFP supports the data rate of telecom protocols, making it an attractive form factor for routing platforms using Ethernet as well as packet over SONET/SDH or Resilient Packet Ring (RPR) interfaces.

Beyond 10 GbE Port Types: Xenpaks

Cisco sports two non-IEEE Xenpak port types in its 10 GbE product line: DWDM (shipping since July 2004) and ZR (slated for release in the fourth calendar quarter of 2005). DWDM implements a PMD to transmit 32 different channels for a whopping 320 Gbit/s over a single strand of fiber. With the aid of optical amplifiers, the DWDM signals can be carried up to 200 km. ZR is capable of reaching in the 80 km range, doubling the distance specified by the 802.3ae 10GBASE-E standard.

DWDM Xenpak: the Quest for Unlimited Bandwidth

DWDM Xenpak allows businesses to significantly reduce the amount of transport equipment in their network by eliminating the need to perform wavelength conversion using dedicated DWDM transponders (see Figure 4). But what is really unique about DWDM Xenpak is that it enables a seamless integration of the optical layer with switching and routing. With DWDM Xenpak, 10 GbE routers and switches combine in the same platform and architecture the OSI functions of Layer 1 through Layer 3 and above.

In doing so, DWDM Xenpak offers a new level of integration between transport and Ethernet/IP functions on high-end switches and routers.

For service providers building next-generation triple play networks, consolidating the DWDM function and Layer 2 and Layer 3 forwarding on the same platform yields significant cost savings in capital expenditures (no need for transponders, i.e., less optics in the network) and operating expenditures (management and operational costs of transponder systems are eliminated).

For enterprises, this 10 GbE DWDM technology largely removes the need for dedicated transport equipment, which is sometimes too complex and expensive for an enterprise to operate. With DWDM Xenpaks, the transport layer can be limited to simple passive multiplexer devices that don't require any management.

Enterprises can employ DWDM Xenpak even for initial deployments when only one channel is required: the value of DWDM Xenpak lies in its ability to gracefully upgrade bandwidth on the same fiber infrastructure by plugging new DWDM Xenpaks into the core switches. This process is accomplished by managing the optical layer from the familiar IOS command-line interface (CLI)—simply. The DWDM Xenpak is, and will most likely be for quite some time, the most technology-rich pluggable module on the market.



Cisco is committed to support all three 10 GbE pluggable form factors (Xenpak, X2, and XFP) across its 10 GbE product lines, and to deliver all the port types defined in the IEEE 802.3 standard. What's more, Cisco will continue to look beyond the standards to provide the widest range of technologies and to keep expanding the breadth of 10 GbE applications for its customers. ■

Detecting Network Failures

Bidirectional Forwarding Detection protocol in IOS helps increase the speed of failure detection and recovery.

By Asad Faruqi

Bidirectional Forwarding Detection (BFD) is a protocol used to detect faults in the bidirectional path between two forwarding engines. Currently, detection can be fairly quick in certain circumstances when data link hardware comes into play (such as SONET alarms). However, some media, such as Ethernet, do not provide this kind of signaling and might not detect certain kinds of failures in the path, for example, failing interfaces or forwarding engine components.

BFD is supported on Gigabit Ethernet LAN interfaces and higher only, beginning with Cisco IOS Software Release 12.0S. Currently this feature is supported on the Cisco 12000 (12.0S), 7600 (12.2S), 7300 (12.4T), and 7200 (12.2S, 12.4T and 12.0S) series platforms.

Routing protocols use slow hello mechanisms to detect failures when there is no hardware signaling to assist. The detection times available in the existing protocols are no faster than a second, which is far too long for some applications, and represents a great deal of lost data at gigabit rates. BFD tries to bridge this gap by providing fast fault detection of neighbors and furnishing this information to the routing protocols as soon as a link goes down so they can start the convergence process. An additional advantage of BFD is that it provides a single mechanism that can be used for liveness detection over any media, at any protocol layer, with a wide range of detection times and overhead.

BFD Protocol Description

BFD detects communication failures with a data plane next hop. The client protocols that currently support BFD are Open Shortest Path First (OSPF), Intermediate System-to-Intermediate System (IS-IS), Enhanced Interior Gateway Routing Protocol (EIGRP), and Border Gateway Protocol (BGP). Refer to "Cisco IOS Software Support" at the end of this article to learn which routing protocols are supported for BFD in which releases of the Cisco IOS Software.



ASAD FARUQI CCNP, CCNA, is a software engineer for the Core IP Routing Group in the Internet Technologies Division at Cisco. He can be reached at afaruqi@cisco.com.

After BFD starts, the client protocol tells BFD to establish adjacencies to the client's neighbors. The BFD process creates an adjacency structure for the neighbor and attempts to establish a session. The initial hello interval is one second. When the session is established and the state transitions to Up, then the configured transmit and receive interval values are used. To guarantee timely transmission of BFD control packets, BFD control packet transmission and link down detection is done by a pseudo-preemptive BFD process. In most cases, BFD control packets are sent and received from the Cisco Express Forwarding switching path to avoid the delay incurred by queuing up the packet at process level. The BFD packets are unicast between two directly connected neighbors.

BFD operates in two modes: Asynchronous mode and Demand mode, with an Echo adjunct mode. The Cisco IOS Software currently supports only the Asynchronous mode, which is the primary operating mode. In Asynchronous mode, systems send BFD control packets to each other. If a prenegotiated number of these control packets is not received in a row by the other system, the session is taken down. The BFD protocol relies on sending a 24-byte control packet that includes the values specified for *localDiscriminator*, *remoteDiscriminator*, *min_TxInterval*, *min_RxInterval* and *DetectMultiplier*. These values are defined as follows:

- *localDiscriminator*—The local discriminator for this BFD session, used to uniquely identify it. It must be unique on this system and nonzero.
- *remoteDiscriminator*—The remote discriminator for this BFD session. This is the discriminator chosen by the remote system and is completely transparent to the local system.
- *Min_TxInterval*—The minimum interval, in microseconds, between transmitted BFD control packets that this system wants to use. The actual interval is negotiated between the two systems.
- *Min_RxInterval*—The minimum interval, in microseconds, between received BFD control packets that this system requires.

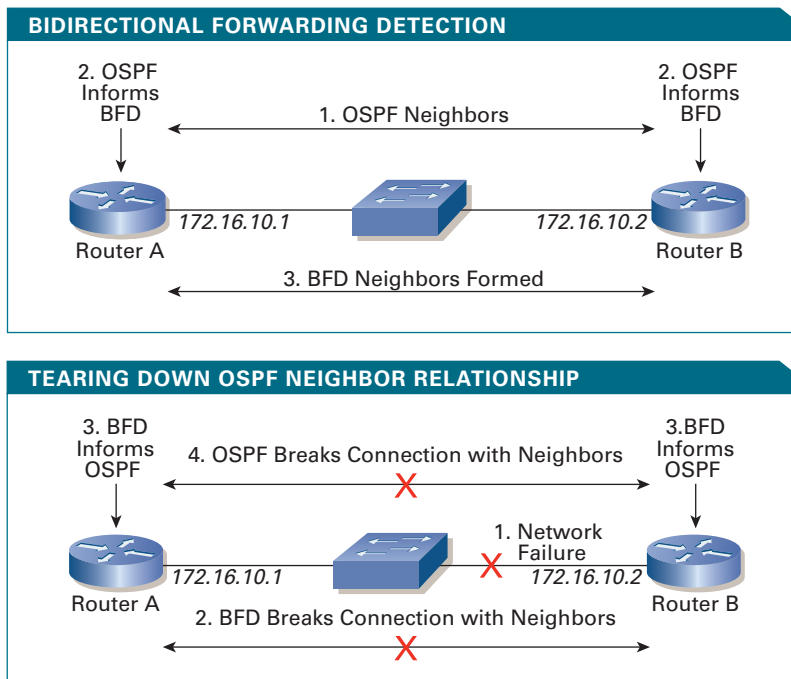


FIGURE 1 When OSPF discovers a neighbor (step 1) it sends a request to the local BFD process to initiate a BFD neighbor session with the OSPF neighbor router (step 2). The BFD neighbor session with the OSPF neighbor router is established (step 3).

FIGURE 2 When a failure occurs in the network (step 1), the BFD neighbor session with the OSPF neighbor router is torn down (step 2). BFD notifies the local OSPF process that the BFD neighbor is no longer reachable (step 3). The local OSPF process tears down the OSPF neighbor relationship (step 4).

- DetectMultiplier**—The desired detect time multiplier for BFD control packets. The negotiated control packet transmission interval, multiplied by this variable, will be the detection time for this session (as seen by the remote system).

The Detection Time (the period of time without receiving BFD packets after which the session is determined to have failed) is not carried explicitly in the protocol. Rather, it is calculated independently in each direction by the receiving system based on the negotiated transmit interval and the detection multiplier. In Asynchronous mode, there may be different detection times in each direction. In Asynchronous mode, the Detection Time calculated in the local system is equal to the value of *Detect Mult* received from the remote system, multiplied by the agreed transmit interval (the greater of *RequiredMinRxInterval*, and the last received *Desired Min TX Interval*), in other words the system reporting the slower rate determines the transmission rate. The *Detect Mult* value is (roughly speaking, due to jitter) the number of packets that must be missed in a row in order to declare the session to be down.

Example

Figure 1 shows a simple network with two routers running OSPF and BFD.

Figure 2 shows what happens when a failure occurs in the network. If an alternative path is available, the routers immediately start converging on it.

Now let's look at how routing protocols use BFD and how to configure BFD for different protocols.

BFD for OSPF

OSPF registers all neighbors that it is interested in tracking via BFD with the BFD process. After a neighbor is registered, BFD initiates a session with the neighbor (if a session does not already exist). When BFD detects that a neighbor has gone down, it sends a notification. If BFD for OSPF is enabled on this interface, it forces the neighbor down. This generates a neighbor down event that causes new versions of needed link-state advertisements (LSAs) to be generated. This triggers Shortest Path First (SPF) calculation, causing a reroute of traffic to other viable paths.

You can configure BFD for OSPF on the router mode and enable it on all interfaces where this protocol instance runs and BFD is enabled. From the command-line interface (CLI), enter the following:

```
router ospf 1
[no] bfd all-interfaces
```

You can also enable BFD on a particular interface, in which case OSPF will register all of its neighbors with BFD for that interface. Interface-level CLI overrides any router mode CLI. Enable or disable BFD using the following Interface mode command:

```
[no] ip ospf bfd [disable]
```

The following **show** commands indicate BFD status:

```
RouterA#sh ip ospf int e2/0
Ethernet2/0 is up, line protocol is up
  Internet Address 172.16.10.1/16, Area 0
  Process ID 1, Router ID 172.16.10.1, Network Type
  BROADCAST, Cost: 10
  Transmit Delay is 1 sec, State BDR, Priority 1,
BFD enabled
  Designated Router (ID) 172.16.10.2, Interface
  address 172.16.10.2
  Backup Designated router (ID) 172.16.10.1,
  Interface address 172.16.10.1
  Timer intervals configured, Hello 10, Dead 40,
  Wait 40, Retransmit 5
```

show ip ospf neighbor detail tells you whether the neighbor is being monitored by BFD.

```
RouterA#sh ip ospf nei 172.16.10.2 det
Neighbor 172.16.10.2, interface address
172.16.10.2
  In the area 0 via interface Ethernet2/0, BFD
  enabled
  Neighbor priority is 1, State is FULL, 6 state
  changes
```

show ip ospf will also include a change to incorporate the case of router-level CLI.

```
RouterA#sh ip ospf
Routing Process "ospf 1" with ID 172.16.10.1
```

Supports only single TOS(TOSO) routes

 Number of areas transit capable is 0
 External flood list length 0
BFD is enabled

I Hear You bit: 1 - Demand bit: 0
 Poll bit: 0 - Final bit: 0
 Multiplier: 3 - Length: 24
 My Discr.: 2 - Your Discr.: 2
 Min tx interval: 50000 - Min rx
 interval: 50000
 Min Echo interval: 0

Sample Configuration Using OSPF

In the following example, the simple OSPF network consists of Router A and Router B. Ethernet interface 2/0 on Router A is connected to the same network as Ethernet interface 1/0 in Router B. The example, starting in global configuration mode, shows the configuration of BFD. For both Routers A and B, BFD is configured globally for all interfaces associated with the OSPF process.

Configuration for Router A

```
interface Ethernet2/0
 ip address 172.16.10.1 255.255.0.0
 bfd interval 50 min_rx 50 multiplier 3
!
router ospf 1
 log-adjacency-changes detail
 network 172.16.0.0 0.0.255.255 area 0
 bfd all-interfaces
```

Configuration for Router B

```
interface Ethernet1/0
 ip address 172.16.10.2 255.255.0.0
 bfd interval 50 min_rx 50 multiplier 3
!
router ospf 1
 log-adjacency-changes detail
 network 172.16.0.0 0.0.255.255 area 0
 bfd all-interfaces
```

The output from the **show bfd neighbors details** command verifies that a BFD session has been created and that OSPF is registered for BFD support. The relevant command output is shown in bold in the following output.

RouterA#sh bfd neighbors details

```
OurAddr      NeighAddr    LD/RD RH  Holdown(mult)
State      Int
172.16.10.1  172.16.10.2  2/2  1  94  (3 )
Up          Et2/0
Local Diag: 0, Demand mode: 0, Poll bit: 0
MinTxInt: 50000, MinRxInt: 50000, Multiplier: 3
Received MinRxInt: 50000, Received Multiplier: 3
Holdown (hits): 150(0), Hello (hits): 50(8605)
Rx Count: 8607, Rx Interval (ms) min/max/avg:
32/72/49 last: 56 ms ago
Tx Count: 8609, Tx Interval (ms) min/max/avg:
32/72/49 last: 16 ms ago
Registered protocols: OSPF
Uptime: 00:07:08
Last packet: Version: 0 - Diagnostic: 0
```

RouterA#

Output for Router B is similar:

RouterB#sh bfd neighbors details

```
OurAddr      NeighAddr    LD/RD RH  Holdown(mult)
State      Int
172.16.10.2  172.16.10.1  2/2  1  142 (3 )
Up          Et1/0
Local Diag: 0, Demand mode: 0, Poll bit: 0
MinTxInt: 50000, MinRxInt: 50000, Multiplier: 3
Received MinRxInt: 50000, Received Multiplier: 3
Holdown (hits): 150(0), Hello (hits): 50(10131)
Rx Count: 10137, Rx Interval (ms) min/max/avg:
32/64/49 last: 8 ms ago
Tx Count: 10135, Tx Interval (ms) min/max/avg:
36/64/49 last: 36 ms ago
Registered protocols: OSPF
Uptime: 00:08:24
Last packet: Version: 0 - Diagnostic: 0
I Hear You bit: 1 - Demand bit: 0
Poll bit: 0 - Final bit: 0
Multiplier: 3 - Length: 24
My Discr.: 2 - Your Discr.: 2
Min tx interval: 50000 - Min rx
interval: 50000
Min Echo interval: 0
```

RouterB#

Cisco IOS Software Support

The BFD feature was introduced in Cisco IOS Software Release 12.2(18)SXE. This version supports the OSPF, EIGRP, and IS-IS protocols. Cisco IOS Software Release 12.0(31)S supports the BGP, IS-IS, and OSPF protocols. Support for 12.4T will be available in the October 2005 timeframe and will support the BGP, IS-IS, OSPF, and EIGRP protocols.

◆ ◆ ◆

For additional BFD sample configurations for EIGRP, IS-IS, and BGP, visit cisco.com/packet/173_5c1. ■


FURTHER READING

- IETF BFD Working Group
www.ietf.org/html.charters/bfd-charter.html
- Cisco IP Routing
cisco.com/packet/173_5c1

SOFTWARE INNOVATION

EXPANDING THE HORIZONS OF CISCO IOS



A landscape photograph of a desert with mountains in the background under a cloudy sky. The sky is filled with soft, white clouds, and the mountains are a dark, hazy blue. The foreground is a flat, sandy desert floor with sparse, low-lying vegetation.

OPERATING ON MORE than 10 million active systems ranging from small home office setups to the largest service provider networks, the Cisco IOS Software Family is the world's most widely adopted networking software. Its standards-based design integrates innovative technologies, business-critical IP services, and unmatched platform support.

Underscoring this industry-leading position is Cisco's ceaseless commitment to evolving IOS through technology innovation and capabilities that address the ever-changing requirements of the many markets IOS serves. This commitment is demonstrated in the following pages.

"IOS Modularity Debuts in the Enterprise" (page 35) describes how the new Catalyst 6500 with Cisco IOS Software Modularity helps enterprises minimize downtime and boost efficiency. For service providers migrating to an IP next-generation network (IP NGN), the spotlight is on IOS XR. "Scaling New Heights" (page 43) details the scalability, availability, and service flexibility requirements that IP NGNs increasingly demand, and how IOS XR satisfies each one of them.

Complementing Cisco's IOS Software, "Managing the Network as a System" (page 38) lays out a trio of recently launched, automated software tools that help you (enterprises and service providers alike) configure and monitor specific devices and applications in the network, and also defines a fresh perspective on managing the network to fit *your* specific business needs. ■

IOS MODULARITY DEBUTS IN THE ENTERPRISE

CATALYST 6500 WITH CISCO IOS SOFTWARE MODULARITY MINIMIZES DOWNTIME AND BOOSTS OPERATIONAL EFFICIENCY IN THE DATA CENTER AND ENTERPRISE CAMPUS.



CISCO IOS SOFTWARE Modularity on the Catalyst 6500 Series sets new standards for network availability through fault containment and faster fault recovery in places where it is needed the most: single points of failure in the network, ranging from the data center to the enterprise wiring closet.

With this innovation on the Catalyst 6500, network administrators can apply incremental patches to address time-sensitive requirements such as critical security fixes—a capability that not only reduces complexity of the software certification and upgrade process but does so without impacting network availability. Additionally, support for Cisco IOS Software Modularity on the Catalyst 6500 combined with the monitoring and automated response features of Embedded Event Manager (EEM) further simplifies network operations.

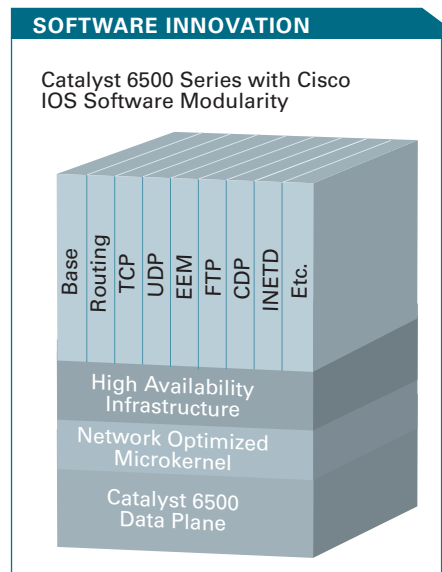
“The stability and uptime of Cisco’s network has always been one of our highest priorities, especially at the access layer where technology solutions such as

MODELS FOR MODULARITY
Cisco IOS Software Modularity on Catalyst 6500 Series switches enables modular Cisco IOS subsystems to run in independent, self-healing processes.

IP telephony depend on constant connectivity,” says Mike Norman, director of intelligent network solutions for Cisco IT. “Catalyst 6500 with Cisco IOS Software Modularity is a vital and welcome component that will significantly enhance that uptime, alleviating potential outages due to failure, and allowing network devices to have security fixes applied without incurring downtime.”

Industry studies show that demands on the network—by advanced technologies such as IP telephony, sophisticated security measures, and converged applications, for example—are growing faster than IT staff, creating an “operational gap.” The impact of downtime in both real dollars and soft costs has grown increasingly higher. Nowhere are these pressures more acute and the need for operational efficiencies greater than in the enterprise data center, core, and access layers.

Now, the Cisco IOS Software Modularity innovation from Cisco promises to go a long way toward alleviating some of the operational and downtime pressures for enterprises. Catalyst 6500 with Cisco IOS Software Modularity will be delivered on the existing Cisco IOS 12.2SX Release for the Supervisor Engine 720 and Supervisor Engine 32 in the fourth calendar quarter of 2005. No additional licensing is required. “As a result, existing and new Catalyst 6500 Series customers can quickly benefit from this high-value technology,” says John Yen, senior manager for enterprise switching



MODULARITY GAINS Catalyst 6500 Series switches with Cisco IOS Software Modularity simplify software updates, minimize unplanned downtime, and enable process-level, automated policy control.

Cisco IOS Software Modularity on the Catalyst 6500 Series yields the following benefits for enterprise network administrators:

- **Minimized unplanned downtime** through self-healing, independent processes
- **Simplified software changes** through subsystem in-service software upgrades (ISSU)
- **Process-level, automated policy control** by integrating EEM

“OUR ABILITY TO DELIVER NEW INNOVATIONS SUCH AS CISCO IOS SOFTWARE MODULARITY ALONG WITH THE UNMATCHED INVESTMENT PROTECTION OF THE CATALYST 6500 PLATFORM IS HIGHLY VALUED BY OUR CUSTOMERS AND IS AN AREA OF CONTINUED INVESTMENT FOR US.” —JAYSHREE ULLAL, SENIOR VICE PRESIDENT FOR SECURITY, DATA CENTER, AND SWITCHING, CISCO

in Cisco’s Product and Technology Marketing Organization.

“Our ability to deliver new innovations such as Cisco IOS Software Modularity along with the unmatched investment protection of the Catalyst 6500 platform is highly valued by our customers and is an area of continued investment for us,” adds Jayshree Ullal, senior vice president of security, data center, and switching at Cisco.

Simplifying Software Changes and Security Fixes

Especially compelling about Catalyst 6500 with Cisco IOS Software Modularity is the ability to simplify the software change process. These new software advancements enable subsystem ISSU so that code changes can be made while

other control and data plane components remain in operation with zero downtime.

Generally, IT departments go through a process of certifying new software first. Then they schedule a downtime window and finally roll the software out to a number of devices. Now with support for Cisco IOS Software Modularity on the Catalyst 6500, code certification of patches and bug fixes can be performed more quickly because these code changes are specific and contained, reducing the risk of cross-software effects. Thus, the number of tests to be run is smaller. And because these code upgrades can be performed in-service, they can be deployed sooner because IT administrators don't have to wait for planned downtime windows.

Based on customer testing, enterprises can reduce the time it takes to certify and deploy new software from months to weeks. Accelerating the rate of change is particularly crucial to mission-critical environments, such as healthcare and trading-floor businesses, as well as to real-time network service environments that support voice and video applications.

Also important in all businesses is hastening the deployment of security patches, such as those issued from the Cisco Product Security Incident Response Team (PSIRT). Some organizations have postponed PSIRT upgrades or have instituted a policy of performing them once a year or so because of the burden and cost of certification and scheduled downtime. By significantly lowering the operational costs associated with security patching, security fixes can be kept much more current, observes Sachin Gupta, Catalyst 6500 product line manager.

"Getting security fixes in place quickly reduces business risks associated with vulnerabilities in software that emerge as hackers become increasingly sophisticated," says Gupta.

Minimizing Unplanned Downtime

With Catalyst 6500 support for Cisco IOS Software Modularity, software processes run in their own protected memory space. So the impact of a software fault is contained, resulting in less disruption in the Catalyst 6500 Series Switch.

In addition, this innovation supports restartable processes and state checkpointing, allowing a faulty process to restart and to return to a previous known state and configuration without a system reboot. This capability accelerates restart recovery time to milliseconds, down from minutes.

"Without these capabilities, the switch would have to rebuild all the forwarding tables," explains Gupta. "For example, previously, if the routing process were to restart, forwarding would stop while the routes were relearned." But now a network administrator could transparently restart a process, such as a routing, TCP, UDP, or FTP process, thereby containing faults. This can also be done automatically, via policy.

"If a modular process enters an undesirable state, it does not impact the entire system," Gupta notes. "You restart that process, not the entire system."

Catalyst 6500 with Cisco IOS Software Modularity supports cold, warm, and hot restarts, depending on the degree of state checkpointing. The more critical a particular process, the more data about that process will be checkpointed.

"The software checkpoints enough data for each process to enable continuous operation," says Gupta. "You really need both restartable processes and stateful checkpointing for full operational transparency. Either one alone is not enough."

The continuous-forwarding capability associated with restartable processes and stateful checkpointing in Catalyst 6500 with Cisco IOS Software Modularity complements the Non-Stop Forwarding/Stateful Switchover (NSF/SSO) high availability capabilities in Cisco IOS Software in a dual-Supervisor Engine configuration. Using the transparent restarting capability in the modular software delivers the software resilience benefits of NSF even with a single Supervisor Engine.

Automated Policy Control

Automated Policy Control offloads time-consuming tasks to the network, enabling IT staff to focus on higher value activities and, in doing so, helping to close the aforementioned operational gap created by demands on the network growing faster than IT staff.

Cisco has long offered a distributed approach to event detection and recovery from directly within a Cisco IOS device in the form of EEM. Cisco EEM has now been integrated on the Catalyst 6500 Series with Cisco IOS Software Modularity to give network administrators more granular, automated control in a policy-based fashion.

Administrators define policies using the Cisco command-line interface (CLI) or tool common language (TCL) scripts. EEM can be used in a variety of ways, such as checking for an updated, certified patch on servers, downloading the patch, and installing it at a pre-specified time.

"If a patch is installed and something undesirable happens, EEM can automatically uninstall it," adds Gupta.

EEM with modularity can be used to detect a faulty process, automatically restart it, and then e-mail administrators with diagnostic information to accelerate troubleshooting. It can also regularly check the health of processes through user-defined tests at regular intervals. If a test doesn't pass, a policy can dictate that the particular modular process should be restarted.

◆ ◆ ◆

Catalyst 6500 with Cisco IOS Software Modularity is the latest in a deep pipeline of software innovations being developed by Cisco's IOS Technology Division. The advancements now delivered on the flagship Catalyst 6500 Series will not only help maintain this platform's technology leadership in the LAN switching market, but, more importantly, they address operational and network issues for enterprise IT administrators. ■

FURTHER READING

- Cisco Catalyst 6500 Series Switches
cisco.com/go/catalyst6500
- Cisco Catalyst 6500 Software Modularity White Paper
cisco.com/packet/173_6c1
- Cisco IOS Software
cisco.com/go/ios

MANAGING THE NETWORK AS A SYSTEM

A NEW PERSPECTIVE ON MANAGING NETWORKS AND SOFTWARE TOOLS HELP ENTERPRISES AND SERVICE PROVIDERS MOLD THE NETWORK TO FIT THEIR BUSINESS. BY JANET KREILING

A NETWORK IS A network, not an aggregation of discrete parts. It must be managed as a network, not as a collection of routers, a collection of servers, a collection of switches or as the northern region, the southern region, the territory (and equipment) acquired from X company, and the territory (and equipment) acquired from Y company. Yet, too often, the network management tools available encompass only classes of devices or devices from a given vendor.

Silos (disconnects) in network management get in the way. And, as Cliff Meltzer, senior vice president of the Network Management Technology Group at Cisco, points out, they create specific problems for each type of user. For service providers, he says, “They hinder time to market with new services and applications and raise operations and management costs.” For enterprises, “They make changes and configuring systems more costly and increase operating expenses.” For small to midsized businesses, “They make setting up a network confusing and time consuming, especially as these businesses want voice, video, and data on a converged network just like larger enterprises.”

“Cisco’s goal is to develop software and technologies to help manage networks, not just individual components,” Meltzer says. “We want to create tools that will help each of these market segments solve its individual problems and reach its own goals. That means using standards-based abstractions for resource management and open application programming interfaces [APIs] for element management and application systems, and using Web services so they can all communicate with overall network management. It also means creating software that automates the configuration of multiple devices to consistently implement a system-wide change or to deploy a service or an application. This includes creating programs that demystify the setup and running of a network for Cisco channel partners or customers themselves.”



Tools for Setup, IP Voice, Locating Wireless Devices

Cisco's ongoing flow of new network management capabilities is represented recently by three tools and the widespread application of the language of e-business:

- The Cisco Configuration Engine enables enterprises to configure and reconfigure tens or thousands of devices throughout their networks from a network operations center without human intervention.
- The CiscoWorks IP Telephony Environment Monitor (ITEM) enables enterprises and midsized businesses to monitor and proactively maintain IP communications, again from a central location.
- A new addition to Cisco's integrated wireless network portfolio, the Cisco 2700 Series Wireless Location Appliance locates any IEEE 802.11 wireless device in a building or campus network.
- Cisco's continuing leadership in bringing the Extensible Markup Language (XML) into networks is helping enable devices and applications throughout the network to communicate with each other and with network management.

Cisco Configuration Engine

The Cisco Configuration Engine is a self-contained, Web-based network management software application that enables "zero touch" deployment and management of Cisco integrated services routers. A service provider or enterprise IT staff

plugs information into one or more templates—e.g., for IP communications or virtual private networks (VPNs)—and when one or ten or a thousand routers identify themselves to the network by serial number or some other attribute and ask to be configured, the engine downloads the correct template and other attributes automatically into each device. After each router has been configured, it announces that the task has been successfully completed.

The system can also be used to reconfigure routers, collect data from them, and feed them security, policy, or other software. Thus, routers are easily changed to accommodate new tasks as the business needs them.

“This is plug and play,” says Brian Junnila, marketing manager for network management in Cisco’s Product and Technology Marketing Organization. “On the smaller end, a teleworker or a branch staffer can unpack the box, plug the router in, and everything else is automatic.”

On a larger scale, he adds, “Look at the work and time involved in configuring, for example, 30 routers across the state of Iowa. Factor in a day per router and travel time. You’re tying up an IT person for a month, and during that month he or she is on the road, not back at home and able to handle other work.”

Given an initial deployment of 800 routers, Junnila says, updating the configuration, updating certain software, and gathering configuration data manually costs more than US\$260,000. “Using the Cisco Configuration Engine to do all those tasks costs about \$50,000, with some input from CiscoWorks, for a savings of over \$200,000. That’s two and a half times the cost of an IT operator’s salary for a year.”

Additionally, Junnila points out, what those savings represent is opportunity cost—what the person can do instead of manually managing routers. That is an important consideration now and becoming disparately more so. The need for IT resources is far outstripping the availability.

Web-Based XML Communications

The Cisco Configuration Engine transfers data in XML format. This represents one example of how Cisco is using—and driving—standardized communications in networks. Cisco has been a leader in pushing open interfaces between devices and network management systems and open APIs (see sidebar, “Standards for the Network as a System”).

Web services, the mechanisms that enable e-business communications among companies, are becoming the basis of communications among network components and systems. They rely on the exchange of structured data via standardized, composable technologies addressing security and reliable messaging and eventing. This allows devices, operations and business support systems, APIs, management systems, and anything else that interacts with the network in total or in part to do so easily and independently of the component or its vendor.

Such open communications are becoming even more important as networks bring together geographically distributed and logically diverse elements, according to Nino Vidovic, vice president and chief technology officer for network management at Cisco. “Functions such as storage and computing are placed remotely and united by networks into virtual storage or distributed processing facilities,” he says. “So, more information can be stored, more people can be given access, and larger applications can be run. But for the networks to support businesses effectively, all pieces (computing, storage, and network) must work together—interface smoothly—and be managed together.”

IP Telephony Environment Monitor

IP voice users get vocal when quality isn’t good. The ITEM software-based tool, part of CiscoWorks, continually monitors the quality of voice communications over an IP network to ensure that network managers learn of problems before their clients do.

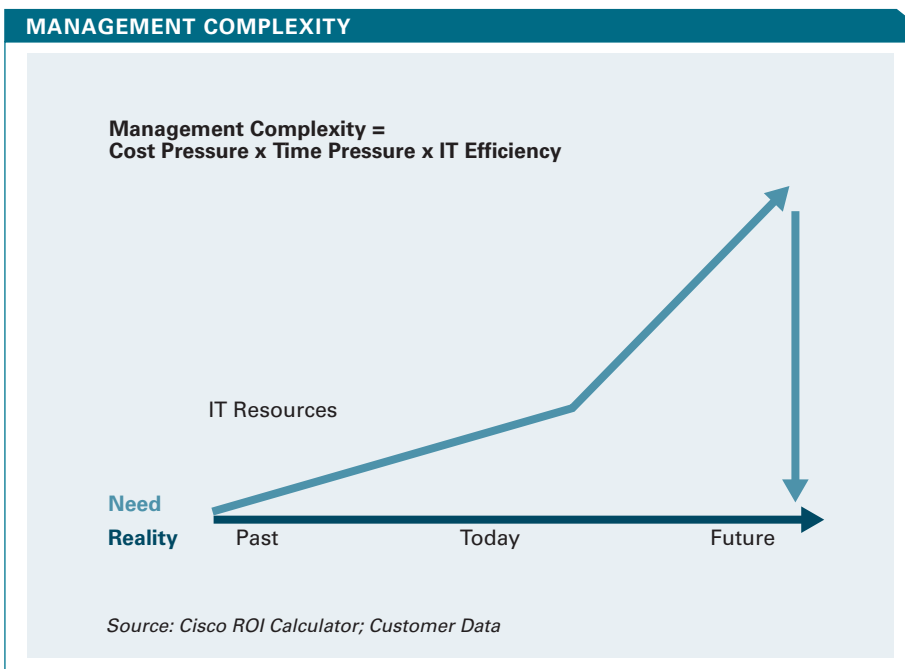
ITEM uses intelligence built into Cisco network hardware and software to send simulated voice packets along various pathways to check for latency, packet loss, and jitter, the basic problems that can bedevil IP voice communications. According to Dave Wetzel, marketing manager, network management, in Cisco’s Product

and Technology Marketing Organization, “ITEM can test access to Cisco CallManagers, voice mail systems, routers, and switches; how long it takes to turn on a message waiting light on the phone and turn it off again; transport links—just about anything that involves voice on an IP network.”

For example, he says, applied in a school district, ITEM found a fan in a switch that was going bad. This enabled network managers to address a developing problem and potentially avoid a catastrophic network outage resulting from a simple equipment failure.

ITEM employs functionality within Cisco IOS Software called IP Service-Level Agreements (IP SLAs), an agent that can make up packets to simulate all types of traffic and route them between any two or more points on the network. ITEM asks IP SLAs to generate voice-like packets and spells out what is to be tested. Tests can be run on a set schedule, on demand, or continuously, enabling companies to find problems before their users get vocal and call to complain, notes Wetzel.

ITEM makes it easier for businesses to converge their voice and data networks, delivering voice cost effectively anywhere their data networks run.



IT PRODUCTIVITY IN DEMAND The need for skilled IT people is growing far faster than the supply. Among Cisco’s recommendations to help alleviate management complexity are a system-level view of the network, a systemic approach to leveraging standard interfaces, and automation of routine tasks.

STANDARDS FOR THE NETWORK AS A SYSTEM

Cisco is changing its view of the network from focusing on point products to emphasizing a systemic approach, says Nino Vidovic, vice president and chief technology officer in the Network Management Group at Cisco, and with that change comes an intense involvement in the standards bodies that deal with intranetwork communications and Web interfaces.

"While Cisco is still involved in organizations such as the Internet Engineering Task Force (IETF), we've been involved now for a number of years in the Distributed Management Task Force (DMTF), International Telecommunications Union (ITU), Organization for the Advancement of Structured Information Standards (OASIS), and World Wide Web Consortium (W3C) as well," he says. "We were one of the founding members of DMTF, which includes many customers and vendors, in 1992, and one of our people, Jim Turner, is now chairman of its board of directors."

One of Cisco's major standards efforts in this area has been the application of XML to intranetwork communications. "We've worked to define protocols and standard data structures governing network management," Vidovic says. "We've been a catalyst in bringing companies together to develop an abstract information model for managing systems, networks, and storage, called the Common Information Model [CIM]." CIM abstracts and standardizes basic management concepts and interfaces such as physical asset and protocol information, or job management and diagnostics services, explains Vidovic. "These interfaces and data work across computing, storage, and networking domains—breaking down the management silos that have sprung up between these areas, and allowing the integration of information in support of a business. Management information can be made available from any device via open APIs and Web services, and data from any device is readily available and integrateable since it is normalized to CIM."

For years, Cisco routers, switches, IP phones, and other devices have shipped with XML interfaces, and Cisco has participated especially in creating a bridge between XML, standardized by W3C, and the CIM, administered by the DMTF. From 2001 to 2004, the DMTF CIM Technical Committee was chaired by Andrea Westerinen, a Cisco employee. Westerinen and other Cisco employees continue to be active in nearly every DMTF working group.

Vidovic expects CIM to be as important to network management in the future as Simple Network Management Protocol (SNMP) has been in the last several decades.

Cisco 2700 Series Wireless Location Appliance

Looking for a wheelchair, a sales manager, a job foreman, a cargo container, a laptop, or an individual toy in a warehouse? As long as he, she, or it has an IEEE 802.11 wireless communication device—a laptop, 802.11 voice over IP handset/ PDA, wireless modem, even a WI-FI radio frequency ID (RFID) tag—the Cisco 2700 Series Wireless Location Appliance can locate the person or thing to within a few meters.

"Through a technique known as RF fingerprinting, the location appliance enables a powerful, Wi-Fi-based solution

for physical asset tracking and security, extending the value of the network beyond traditional data access and electronic communications to real-world applications," explains Mark King, marketing manager, network management, in Cisco's Product and Technology Marketing Organization.

"The technique uses RF prediction to create grids that identify how each portion of a building's floorplan looks to all relevant wireless LAN access points. Real-world information is then gathered by the access points and compared to the

grid to determine a device's location to within a few meters."

Of five suggested applications for the location appliance, four are more reminiscent of business activities than network ones: asset tracking, inventory management, workflow automation, security, and voice over wireless LAN (see related article, page 65).

Asset tracking can locate people or equipment. It's already helping more than one hospital get EKG scanners and wheelchairs to patients faster and also helping prevent them from being stolen. Inventory management is an obvious use; monitoring where people are can help speed them to where they're needed. The appliance can also locate rogue wireless access points and devices so they can be removed quickly. The ability to locate people accurately also helps staff deal with emergencies. In addition, the appliance supports location-based trending for RF capacity management, troubleshooting, and ongoing network modeling.

Forethought

Network management is often an afterthought. But its planning should be front and center with other network and business planning. With forethought in planning, the network can be nimble enough to essentially deliver itself wherever it's needed. For example, many businesses need their networks to direct additional bandwidth to the sales offices at quarter's end so they can file timely reports.

In fact, King adds, "Companies depend so much now on their networks that the network has become an integral part of doing business. Understanding what the network can do becomes crucial. Can it accommodate a move of a manufacturing plant from one city to another? Can it support a new customer support application? Network management tools give business management knowledge of what the network can do, visibility into how it is behaving, and control of what it does." ■

FURTHER READING

- Cisco Network Management cisco.com/packet/173_6d1
- White paper: *Cisco Integrated Wireless Network* cisco.com/packet/173_6d2

IOS XR

SCALING NEW HEIGHTS

ADDRESSING THE SCALABILITY, AVAILABILITY, AND SERVICE FLEXIBILITY REQUIREMENTS THAT IP NEXT-GENERATION NETWORKS DEMAND

RECENT DEVELOPMENTS with Cisco IOS XR Software exemplify Cisco's leadership and commitment to evolving IOS through innovation and customer-focused enhancements. Aimed at service providers phasing out their multiple, single-service networks in favor of a single, converged network, IOS XR Software squarely addresses the key requirements for scalability, continuous system operation, and service flexibility that IP next-generation networks (IP NGNs) increasingly demand.

Specifically designed and optimized for platforms that can scale and distribute processing as well as perform distributed forwarding, IOS XR was first introduced a little more than a year ago on the Cisco CRS-1 Carrier Routing System.

"The scale required in the forwarding and control planes for a system like the CRS-1 could only be met through a completely different approach to building an operating system," says Rob Redford, vice president of marketing in Cisco's Product and Technology Marketing Organization.



HEART INSIDE THE HARDWARE

The distributed, highly scalable capabilities of Cisco IOS XR Software have been extended to the Cisco 12000 Series Router line.

Now, this industry's first fully modular, fully distributed OS has been extended to the popular Cisco 12000 Series Router platform.

Key IP NGN Requirements

Service providers migrating toward an IP NGN need an innovative platform that can support *multiservice scale*, *continuous system operation*, and *secure virtualization*.

- **Multiservice scale**—Consolidating multiple services onto a single platform requires scale in both the forwarding plane and control planes. The forwarding plane must support increasing bandwidth, robust quality of service (QoS) that preserves service-level agreements (SLAs), and scalable multicast mechanisms. In parallel, the control plane must support distributed processing intelligence to scale services and customers simultaneously with increases in forwarding performance. This scalable intelligence allows carriers to move up the service value chain from basic transport to content delivery.

Multiservice Scale

Cisco IOS XR Software supports a distributed forwarding architecture that allows the system to scale interfaces and features for multiservice applications. One of the key innovations with IOS XR is moving away from a central CPU control plane model to a fully distributed processing model. This includes the distribution of many processes out to the line cards, such as Layer 2 interface state. Because each line card requires knowledge only of its own interfaces and not those of the entire system, this architecture yields greater scale in the number of customers and services supported.

Processes such as routing and signaling protocols can run on a single processor or be distributed over multiple route processors. Regardless of where these processes are performed and which memory pool they use, they operate as if they are running on the same, single route processor.

An entire route processor can be dedicated to applications such as Border Gateway Protocol (BGP) that require lots

"THE SCALE REQUIRED IN THE FORWARDING AND CONTROL PLANES FOR A SYSTEM LIKE THE CRS-1 COULD ONLY BE MET THROUGH A COMPLETELY DIFFERENT APPROACH TO BUILDING AN OPERATING SYSTEM." —ROB REDFORD, VICE PRESIDENT OF MARKETING, CISCO PRODUCT AND TECHNOLOGY MARKETING ORGANIZATION.

- **Continuous system operation**—A step beyond the concept of 99.999 percent availability, continuous system operation aims to minimize service outages during periods of both unplanned downtime and scheduled downtime when scaling capacity or adding new services or features.
- **Secure virtualization**—The ability to consolidate multiple networks and platforms into a single system drives costs down but raises organizational challenges among groups that traditionally do not share network elements. Secure virtualization enables a single system to be logically partitioned along service and administrative boundaries. Services and customers are isolated from one another for maximum security and enhanced operational and management efficiency.

of CPU and memory resources. This unique software distribution capability enables control plane scalability and ensures that no CPU or memory usage can become a resource bottleneck.

On the Cisco XR 12000 Series, for example, IOS XR Software distributes independent instances of critical infrastructure services, such as TCP/IP and interprocess communication, on every system card. This gives autonomy to each card in the system, improving performance and resource management.

Interprocess communication is the mechanism that enables software processes to communicate to share data or events with one another. Without efficient interprocess communication, a system's performance

Cisco IOS XR Software satisfies each of these requirements.

and reliability can degrade significantly. Cisco IOS XR Software supports both reliable unicast and reliable multicast interprocess communication.

Reliable unicast allows point-to-point communication between two processes; reliable multicast enables point-to-multi-point communication between multiple processes in the system. With reliable multicast, if the hardware replicates packets, the application hands a single copy of the message to the hardware, and the hardware does the rest.

Each interprocess-communication message translates into two process context switches. If only direct process-to-process communication is required, reliable unicast interprocess communication is enough. However when one process needs to communicate with hundreds of other processes, such as when a new Forwarding Information Base (FIB) table is computed and must be populated simultaneously across line cards, the performance impact is significant without reliable multicast.

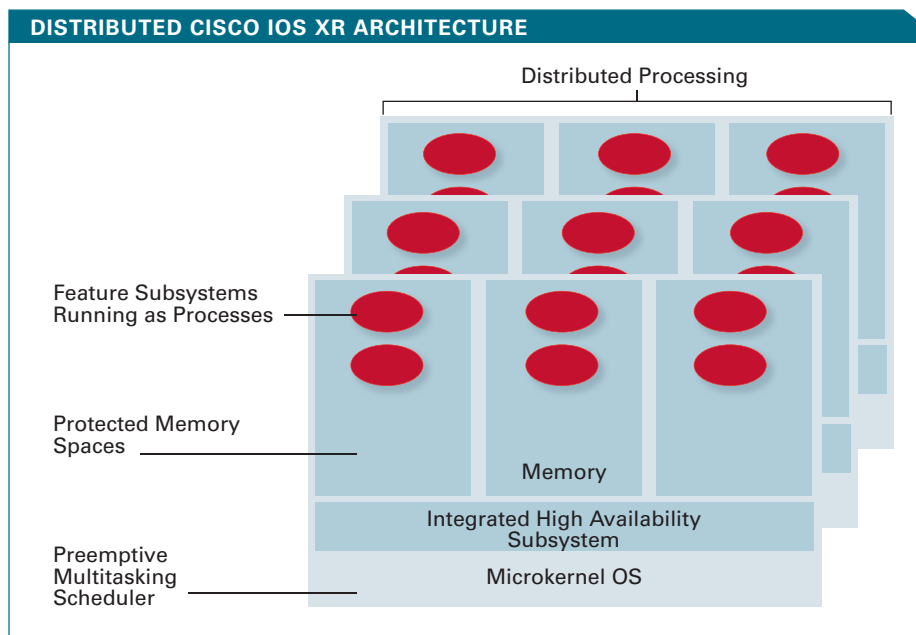
Continuous System Operation

An IP NGN routing OS must proactively prepare and provide protection mechanisms that minimize service outages due to both planned and unplanned downtime. Cisco IOS XR Software minimizes downtime in both instances.

Planned downtime can now be virtually eliminated. Maintenance, upgrades, or replacing software and hardware components can be done with no service loss. Unplanned downtime—resulting from a system failure or security violation, for example—is greatly minimized by the modular, microkernel-based architecture of IOS XR.

The granular modularity of IOS XR ensures that only required processes are restarted upon process failure or during software upgrades. This enables *in-service software upgrades (ISSU)* for software patches and features, and supports Nonstop Forwarding (NSF) and graceful restart extensions of routing and signaling protocols.

A microkernel architecture offers several important advantages. There is fault isolation for all system and kernel processes, and user applications. User applications



DISTRIBUTED EFFICIENCY Processes run independently in protected memory on each microkernel, so faults in one process or its memory do not negatively affect other processes. And because processes are distributed among line cards, failure of one card has minimal effect on others.

do not crash the system, and system processes, such as interprocess communication, do not affect kernel threads.

Like UNIX, the microkernel supports a time-sharing, multitasking scheduler and rapid application development through a POSIX-compliant API. Application and system processes can be added, removed, or upgraded while the system is in service.

The kernel is lightweight and does not include system services such as the file system, device drivers, and network stack. By removing these system services from the kernel, faults can be contained and restarted without having to restart the kernel. Because these services are implemented as independent processes, they can be restarted like any other application.

Each process, such as BGP, Intermediate System-to-Intermediate System (IS-IS), and Open Shortest Path First (OSPF), is further segmented into individual threads and can be distributed to different processing resources. The “isolated” procedure of patching one piece of code relieves network operators from having to certify an entire OS just to change or upgrade one small piece of it.

Keeping the downtime of each router across the network to a minimum also greatly contributes to high network avail-

ability overall and continued network operations. The more granular the software modularity is, the higher the availability of the router OS.

To this end, IOS XR Software supports redundant hardware configurations such as active or standby route processor. Multiple, simultaneous process failures can be handled gracefully within the active route processor with the right protection mechanism invoked, such as switchover to standby route processor or line card link protection.

IOS XR, for example, uses a distributed checkpoint database between multiple pairs of active and standby route processors to speed up switchover. It also supports flexible route processor switchover, allowing applications to be designed with the optimal state replication mode based on their scale, availability, and performance requirements.

“The high degree of granular modularity in IOS XR translates into fewer costs in lost revenue and productivity associated with network downtime,” says John Braskamp, director of marketing for Cisco IOS Software in Cisco’s Product and Technology Marketing Organization.

A Note about IOS XR Packaging

Cisco IOS XR Software is based on a modular release design—an essential ingredient for ISSU. Without modularity, every software release could affect every software component or process. Moreover, this approach minimizes the amount of software users must qualify and install, and decreases operational expenses and Mean Time To Repair (MTTR).

In the release design, IOS XR uses a component-based development model where features, such as BGP and IS-IS, are implemented as *components*. Individual components are aggregated into a *package*. Components inside a package might depend on components in the same or other packages. Example IOS XR Software packages include OS, Admin (shelf management components), Line Card, Multicast, Security, and Manageability. Some packages are mandatory, others optional for the functioning of the router.

Further, packages are bundled into larger packages called *composites*. Packages and composites are built, tested, and released by Cisco based on the features developed and delivered in a particular release, and can be installed on parts of the system or incrementally installed across the entire system.

Secure Virtualization

Cisco IOS XR Software Service Separation Architecture (SSA) allows service providers to consolidate multiple networks and services onto a single platform—while keeping each network and service instance separate and secure. Benefits of this consolidation through secure virtualization include feature transparency, optimal point of presence (POP) design, and efficiencies in operating expenses and capital expenditures resulting from fewer network elements.

The primary building block of Cisco IOS XR SSA is the Service Domain Router (SDR). SDRs are partitioned on a hardware boundary and share only the physical chassis, including the power shelves, switch fabric, and fans. Each SDR has dedicated line cards (including line card memory and CPU) and route processors, along with dedicated Routing Information Base (RIB) tables and control plane.

Using Cisco IOS XR SSA, each domain can be kept secured from every other domain. For example, each domain can have its own route processor and therefore its own instantiation of Control Plane Policing. This capability ensures that each domain is protected separately against reconnaissance and distributed denial of service (DDoS) attacks. An attack on one domain will not affect another domain.

Features, services, and customers can remain partitioned from one another, explains Sanjeev Mervana, product line manager for Cisco IOS XR Software. For example, providers can tailor particular line cards or routing applications to specific services. In addition, they can employ multiple SDRs within one system, each with its own CPU and database.

This approach circumvents the need to redesign a POP with additional equipment as the POP grows, in that processes and resources remain partitioned within a single chassis, Mervana points out.

In May 2005, the European Advanced Networking Test Center (EANTC) AG of Berlin, Germany, conducted extensive tests on the Cisco XR 12416 Router configured as four separate SDRs.

The EANTC confirmed that the system operates with true separation between SDRs with zero overlap or forwarding between them. The EANTC also concluded that software and hardware upgrades to one separated router had no performance effect on the others, and that CPU and memory resources on the different SDRs' route processors proved to be independent of one another.

Proving the Benefits of IOS XR in the Field

To date, more than 100 units of the Cisco CRS-1 have been shipped to customers.

Among the providers using IOS XR are Swisscom, Switzerland's incumbent PTT; SOFTBANK BB Corp., the largest broadband provider in Japan; the Pittsburgh Supercomputing Center; and the National Institute of Informatics, which manages the largest academic research network in Japan.

BellSouth Corporation, a Fortune 100 communications company headquartered in Atlanta, Georgia, and a parent company of Cingular Wireless, is currently evaluating the Cisco XR 12000.

“The modular design of the software in the Cisco XR 12000 could enhance our ability to maintain service levels and performance,” says Mike Duckett, research director, BellSouth Science and Technology. “This redesign, we believe, is critical to reducing the risk associated with new feature introduction and to increasing scalability and reliability.”

The China Education and Research Network (CERNET) has also been testing the Cisco XR 12000 Series. CERNET is the first nationwide education and research computer network in China, and it aims to deliver next-generation Internet applications that support more than 900 education and research institutions worldwide.

“We have asked our router vendors for a fully modular software architecture that enables each routing protocol and other applications to run independently from one another,” says Professor Li, CERNET's chief technology officer. “We have run the Cisco XR 12000 with IOS XR in our live network, and we are pleased with the results demonstrated during this deployment.” ■

FURTHER READING

- Cisco IOS XR Software Introduction
cisco.com/packet/173_6b1
- EANTC XR 12000 Test Results
cisco.com/packet/173_6b2
- Cisco XR 12000 Series
cisco.com/packet/173_6b3
- Cisco XR 12000 Manageability
cisco.com/packet/173_6b4
- Cisco XR 12000 Routing Configuration Guide
cisco.com/packet/173_6b5

The Evolving Data Center

Emerging Trends Impacting the Data Center Network Architecture

By Mauricio Arregoces



Enterprises develop business objectives to control operational costs and increase business agility, and advanced technologies in the data center provide new ways to address these business objectives. The adoption of advanced technologies causes the rapid evolution of enterprise data centers.

This article discusses the effects of emerging trends on the data center network architecture and the architectural changes that are likely to influence next-generation data centers.

Emerging Trends

The need to control operational costs results in consolidation of various technology areas such as server, applications, storage, and even data centers. Other trends derive from achieving higher levels of resiliency, which translates to application and data security, business continuance, and server clusters in distributed environments. For example, mergers and acquisitions help achieve rapid growth in different markets, creating the need to share applications and merge networks while controlling access, authentication, routing, and security policies. Network segmentation and virtualization technologies simplify mergers. Some trends result from growth and expansion, which require capacity planning to deal with environmental issues such as power, cooling, cabling, and airflow, as well as network and compute growth.

Consolidation Trends

Consolidation is undoubtedly the most common trend affecting data centers worldwide today. Key business drivers include cost control, operational efficiency, and effective resource utilization. Figure 1, page 48, illustrates technology consolidation in the data center.

Data Center Consolidation

Data center consolidation implies that the number of distributed server farms is reduced and servers are relocated to existing or new facilities. These consolidated data centers are interconnected and the operational, support, and design best practices are standardized. Consolidation objectives include fewer locations to manage, centralization of the compute infrastructure (servers and applications), as shown in Figure 1, page 48, and using the “know-how” held in the existing data center facilities. From a networking perspective, consolidated environments house significantly more servers and applications, which require higher port density, higher use of network-based services, and improved high availability to match the increased level of exposure. These changes influence oversubscription rates, overall throughput, scalability, and high availability targets.

Server Consolidation

Server consolidation implies the reduction of server hardware platforms and supported operating systems, and standardization of application environments (web and middleware tier). Reduction of supported server hardware platforms permits more efficient time to restoration, alleviates spare parts inventory problems, and eases staff expertise demands and the complexity of performing maintenance. Standardizing server hardware includes choosing the BUS/NIC technology (for example, choosing between PCI, PCI-X or PCI-Express, and between 10/100, 10/100/1000, and 10 Gigabit Ethernet for server connectivity), and choosing between single versus multihoming. Networking considerations include revisiting oversubscription rates between servers and the access layer, between the access and aggregation layers, and between the aggregation and core layers; the port density required for multihoming, and the type of network interface card (NIC) teaming. The oversubscription rates could change substantially using PCI-Express and Gigabit

Ethernet because the average burst rates would increase, and dual homing requires twice the ports, as well as Layer 2 adjacency between different access switches.

Blade-Server Integration

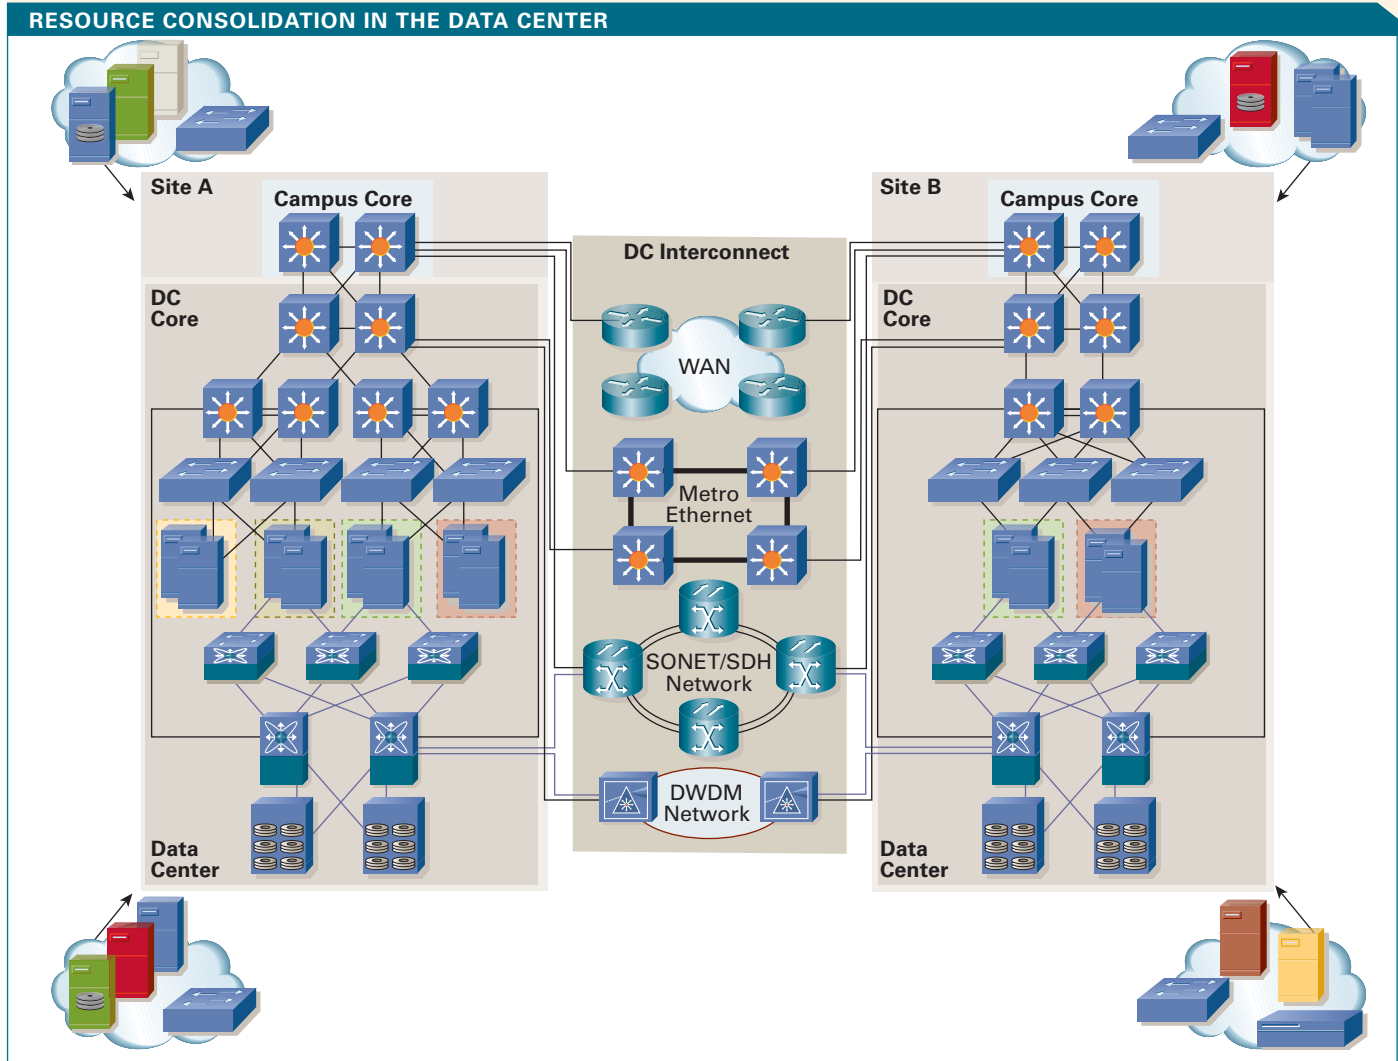
Blade-server technology provides higher compute, memory, and I/O capacity per rack while reducing cabling (network and keyboard, video, and mouse, or KVM). Networking-related aspects of blade-server integration include the I/O and network fabric selection. The I/O choices are pass-through technology and integrated network fabric. Pass-through technology connects the servers directly to the network, whereas the integrated network fabric is a set of redundant switches inside the chassis where the servers connect. In pass-through environments, the servers require external cable connections. When using the integrated network fabric, the

choices include Ethernet, Fibre Channel and Infini-band, and each server is prewired to the slots where the switches are. The integrated switches provide several uplinks to connect to the existing data center infrastructure.

Storage Consolidation

Storage consolidation simplifies storage capacity management through centralization and improves storage capacity utilization. By migrating from Direct Attached Storage Disks (DASD) to storage arrays, storage utilization is more efficient and better managed from centralized storage arrays. Current storage-area network (SAN) fabric technology that supports high Fibre Channel port density and logical SAN separation (through virtual SANs) allows you to further centralize storage and consolidate isolated storage islands while reducing disparate storage fabrics. Networking considerations include higher Fibre Channel port density at

FIGURE 1 Consolidation of data center resources results in fewer distributed server farms and relocation of servers to existing or new facilities.



the edge layer and core layers for server connectivity and storage array and tape subsystem connectivity, respectively. Additional considerations include a high level of redundancy and performance to maintain reasonable levels of oversubscription.

Multi-Tier Application Integration

Multi-tier application integration includes the migration to web-based applications where the presentation (web), application (middleware), and database functions are distinctly decoupled in software and are physically on different servers. Application integration also includes the use of web services, standardization of middleware environments, and support for Service Oriented Architecture (SOA). Network-based services provide the means to scale, secure, and manipulate messages, and enable caching of static and dynamic content, as well as network file system files. These network-based functions are adaptable to physical or logical application tiers per tier, per application, or per group of multi-tier applications. The flexibility in which network-based services are applied comes from their virtualization capabilities, where the underlying hardware supports independent use of logical instances.

Data and Application Security

Data and application security are controlled in the data center as well as the rest of the enterprise network. Data center security addresses the integrity, confidentiality, and accessibility of the information kept in server farms, and it is achieved by mitigating the effects of attacks against servers, applications, storage, and the network infrastructure that supports them.

Server and Application Security

Server and application security include segmentation of server farm application tiers, denial-of-service (DoS) and distributed DoS (DDoS) protection, intrusion/detection protection, and protection against worms. To address the range of requirements, network-based services use agents that run on the servers to specifically protect the application environment and service devices that monitor and analyze traffic including correlating information from multiple devices to reduce false positives. Network-based services include intrusion detection/protection devices, DoS protection devices, firewalls, traffic analysis, and correlation engines.

Data Security

Data security involves protecting information kept in storage arrays, whether connected to the IP or SAN clouds, or kept on the servers themselves (DASD). Two primary goals are maintaining data integrity by preventing tampering and preventing information theft.

When the data is kept in the IP network (NAS heads or DASD), the same protection mechanisms used in server and application security apply. When the information is kept in SAN-connected storage arrays, the security protection mechanisms are provided by the network infrastructure in the form of zoning, device authentication, and port-based security. In cases where the data is kept in the SAN but transported over IP (FCIP or iSCSI) such as in data replication or remote backup, security measures, such as encryption and virtual private networks (VPNs), apply to segment (isolate) and encrypt the data while in transit.

Infrastructure Security

Infrastructure security involves protecting network devices and links that support traffic to/from the server farms. Network device protection includes basic mechanisms such as router authentication, rate limiters to prevent the control plane from being saturated, and DoS/DDoS protection devices that filter unwanted traffic to avoid link saturation. The network infrastructure is hardened using several features such as Control Plane Policing, special purpose rate limiters, router authentication, and port security in addition to other features that address endpoint security problems such as private VLANs.

Virtualization and Segmentation Trends

Virtualization and segmentation affect data center environments and other network places such as the WAN, campus, branch office environments, and the Internet edge. Virtualization and segmentation provide functions required by environments where control of resources such as network devices and data is done based on user roles or functions. User access is controlled at the access points and enforced throughout the network. Functions provided by network devices require logical instances (virtual) from a single physical element to allow its concurrent use in different application environments or by different user groups. Data access is controlled based on the application environment's specific restrictions, which are enforced by instances of network elements. Figure 2, page 51, shows user traffic transported over isolated paths to the server farms where virtual firewalls, virtual load balancers, etc., control access to applications, which in turn might be running on virtual servers. The different application environments are segmented using VLANs, VSANs, and policies on the firewall and load balancers applied through the logical instance of the service.

Virtualization

Virtualization technology covers a wide range of capabilities applied to server and application environments—to network-based services such as firewalls

and load balancers, or to infrastructure components such as routers and switches. Virtualization of servers and applications refers to the use of logical instances of these resources in a manner that is somewhat independent from the hardware platform in which they run. The application service (operating system and application software) is, in essence, a logical instance in a server that is easily moved to any available server when required. This requires pre-existing images of the application service and can require a separation between the I/O functions and the server platform to allow easy manipulation of the logical instances.

Two primary types of server virtualization techniques exist. In one, a single physical server is divided into multiple logical entities, or virtual servers, each running its own operating system and application environment. In the other, multiple servers are logically grouped to appear as a single server with a single operating system in which the CPU capacity could be increased or decreased by adding or removing servers, which simplifies the management of multiple distinct software and hardware configurations.

Network-based service virtualization applies to multiuser group environments to isolate groups from one another, and in multi-tier application environments to isolate each tier or entire multi-tier environment from each other. Within a single application environment, each tier has its own security and scalability requirements, and these requirements are preserved across distinct application environments. Network-based functions apply to each tier and each application environment independently through logical instances. The capability to define and apply logical instances of a network-based function is provided by virtualizing the function, referred to as a *virtual context*. An example of a virtual context is a virtual firewall applied to four different application environments from a set of redundant firewalls, where each virtual context is configured independently per application environment. The firewall function is virtual yet the overall policy is centralized while the specific policy instances are managed independently.



MAURICIO ARREGOCES, CCIE No. 3285, is a manager of an engineering group focused on network design and architecture at Cisco. He holds bachelors and masters degrees in Computer Science and has written and spoken extensively on data center networking. He can be reached at marregoc@cisco.com.

Infrastructure virtualization provides flexibility using infrastructure functions such as VLANs and VSANs to allow logical separation of compute resource groups. These groups extend through the network topology as needed, preserving the group boundary. Other infrastructure virtual services include virtual routers and switches. In virtual routing, routing instances in the same physical switch provide routing capabilities to each application environment independently and collectively from the server farm to the enterprise network. In virtual switching, two physical switches are treated as one, thus simplifying code maintenance and configuration management, but more importantly providing physical redundancy by supporting port channeling and state across different physical switches.

Segmentation

Segmentation is a network-wide trend in which technology aids in the logical separation of user groups and their access to equally segmented application environments. Three areas to consider when designing a segmented environment are user access control, secure transport technology, and segmented application environments. User access control (for internal and external users) is first enforced at the access points, which are typically branch, campus, and Internet edge environments. At these access points users are authenticated, authorized, and once allowed, placed in the group they belong controlled by its associated policy. Once on the network, the transport that allows the users to reach the application environments provides logically isolated paths, and thus a segmented environment. The segmented logical paths provide a secure media for client traffic over a shared infrastructure that terminates near the traffic destination. When the traffic is destined to server farms, the access control is enforced at the data center entry point.

Segmentation technology is useful in mergers and acquisitions, when companies must join networks, share physical data center space, and transport networks between clients and applications.

A network that offers segmentation services requires several technologies used throughout the entire network, some of which rely on virtualization capabilities. At the access points, enforcement mechanisms use IEEE 802.1X and authentication, authorization, and accounting (AAA) to control user access, group assignment, and group policy. The transport of client traffic uses technologies such as Generic Routing Encapsulation (GRE) or Layer 2 Tunneling Protocol Version 3 (L2TPv3) tunnels, Multiprotocol Label Switching virtual private networks (MPLS VPNs), or VPN routing/forwarding Lite (VRF Lite) to offer a secure

VIRTUALIZATION IN THE ENTERPRISE NETWORK

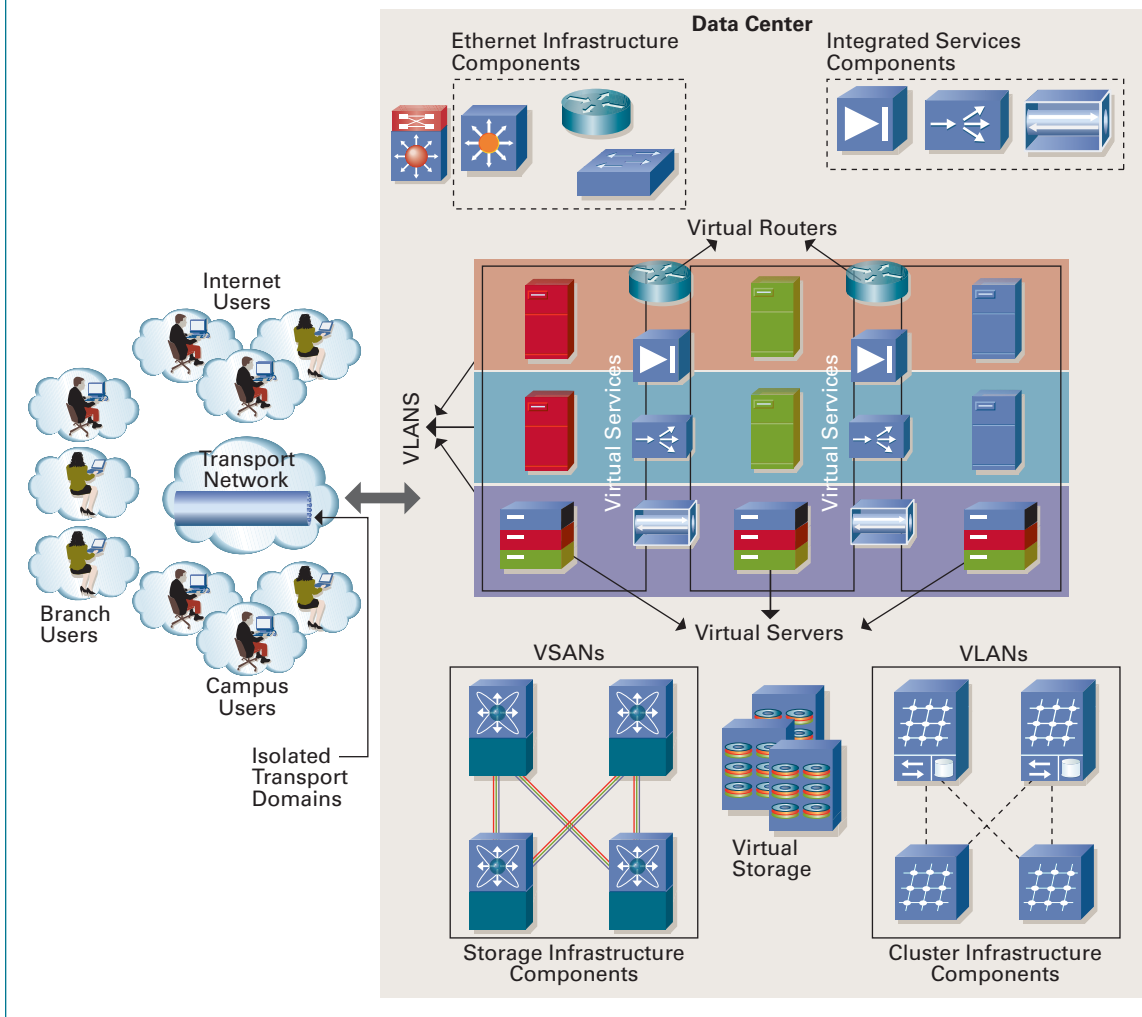


FIGURE 2 With virtualization and segmentation it is possible to create multiple topologies on a single physical infrastructure, thus optimizing the use of network and server resources.

transport using multiple isolated paths over shared infrastructure. At the data center, mapping of the transport path to the application environment transport uses virtual routers, virtual firewalls, and any additional predefined virtual services.

Business Continuation

Business continuation plans address the need to prevent disruption to the business operation, and in many countries worldwide this has translated into government regulations over different industries to prevent economic unrest. To achieve the highest possible availability levels, redundancy in primary data center facilities and across distributed data center environments is required. The network design must consider a resilient primary data

center and the use of multiple data centers: a primary, a backup in close proximity for synchronous replication, a remote for asynchronous replication, and multiple data centers in active-active mode. Traditional business continuation areas have been SAN extension and data replication between distributed data centers. Additional functions include the selection of the application instance closest to the user, known as *site selection*, and the selection of the transport technology supporting storage-to-storage, client-to-server, and server-to-server traffic. In server-to-server environments the transport can also include support for Layer 2 adjacency between server clusters, which is known as *stretched clusters*.

Physical Data Center Environment

Physical environment changes result from consolidation of data centers, servers, applications, and storage;

from the need to provide business continuity; and from the natural evolution of the application environment: higher performance, better response time, and higher scalability. Critical environmental issues include power, cooling, cabling, and floor and rack space. It is often necessary to revisit design assumptions due to new server density per rack and related power, cooling, and cabling requirements. Data center planning also includes reasonable growth mapped down to cooling units, power units, compute capacity per rack-unit, and rack space up to the number or ports required per rack.

New Data Center Network Architecture

These recent trends present many challenges, which taken collectively, require the development of a robust data center network architecture that eases technology adoption. The following criteria are important in developing the data center architecture: selecting the underlying network infrastructure (network fabric), determining data center services, and identifying distributed environment objectives.

Data Center Fabric

The data center fabric is the overall switching infrastructure that supports the server farms (see Figure 3, page 53). The switching infrastructure provides key functions and capabilities to make the communication exchange to, from, and between data center devices more efficient. The communication exchange falls in one of the following types: client-to-server, server-to-server, server-to-storage, and storage-to-storage. Client-to-server traffic belongs to transactional applications where users interact with applications. Server-to-server traffic is either an indirect result of client-to-server interaction (application servers exchanging state information, application servers exchanging information with database servers, and database servers querying each other's databases) or the need for servers to exchange information that belongs to a single job that has been divided into smaller tasks (computational analysis). Server-to-storage traffic consists of hosts accessing their target disk in a storage array or accessing a tape subsystem in traditional block access. For storage-to-storage, the exchange is between storage arrays typically using a protocol for synchronous or asynchronous communication in data replication scenarios.

Connectivity options include Ethernet and Fibre Channel, and emerging technology includes Infiniband. Ethernet continues to evolve by offering higher performance at lower costs, and provides a proven fabric for client-to-server applications. Fibre Channel remains the connectivity fabric for SANs because it provides critical capabilities in the server-to-storage and storage-to-storage communication exchange.

Infiniband emerges by supporting high throughput, low latency, low cost, and critical capabilities in server cluster environments. Each fabric type offers a distinct connectivity alternative, and each provides a set of services that support the various traffic types and their specific requirements.

Data Center Services

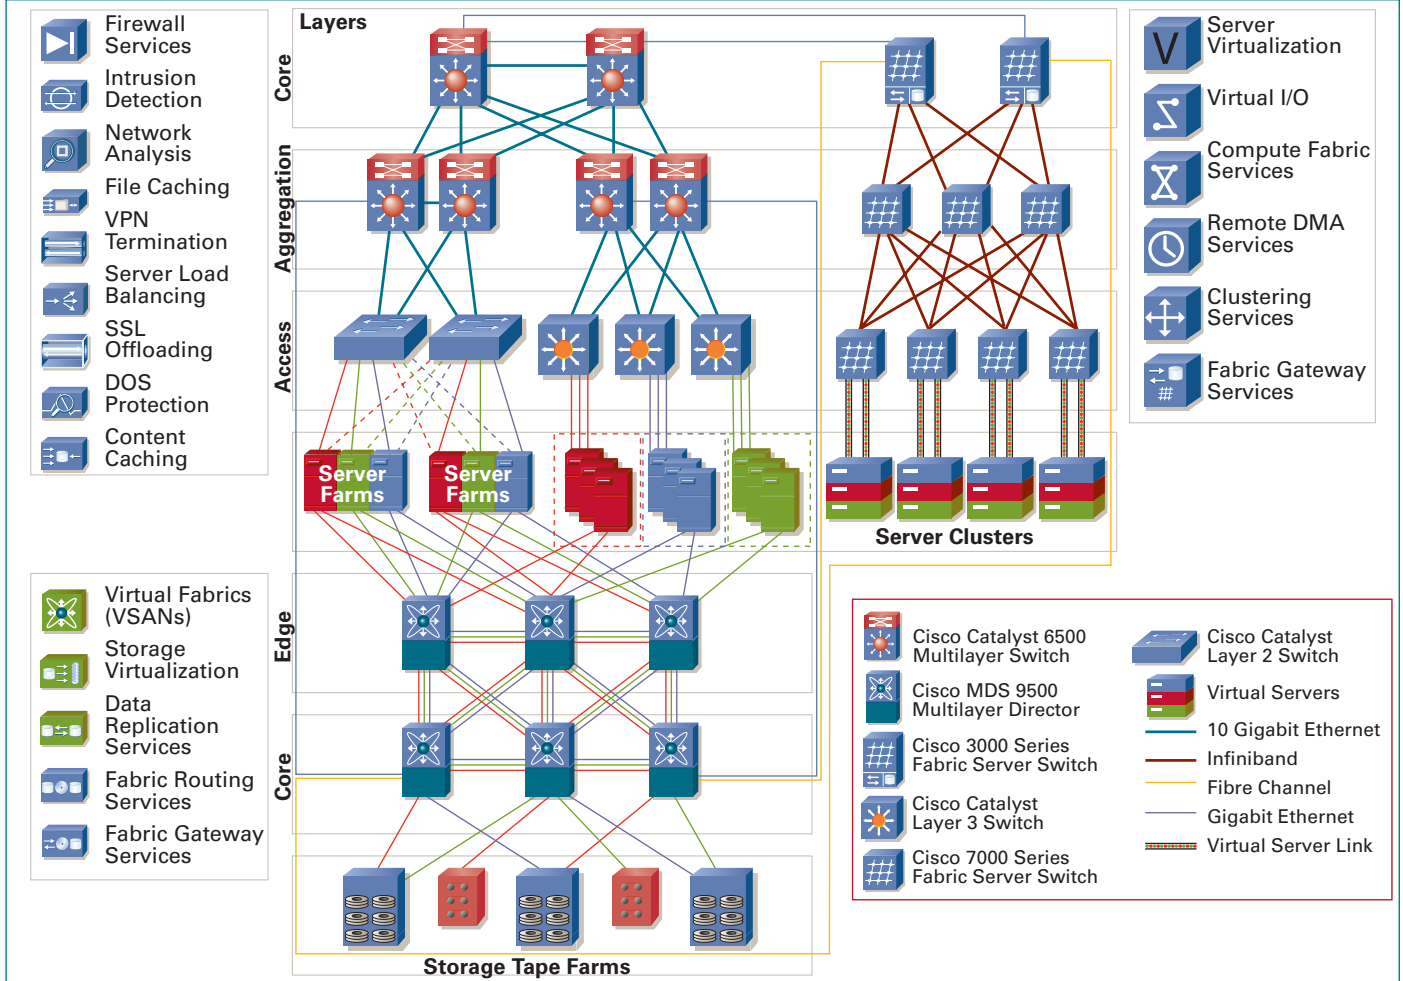
Ethernet fabric services primarily apply to client-to-server traffic and some transactional application server-to-server traffic. These services include security, application optimization, and network management. Security services include firewall, intrusion detection, IP Security (IPSec), Secure Sockets Layer (SSL) VPN, and DoS and DDoS mitigation, as well as basic network-based services such as broadcast suppression, ARP inspection, and PVLANS. Application optimization services include server load balancing, SSL offloading, and two new service types: *file caching* and *message manipulation*. File caching services are used in server consolidation by allowing the centralization of files in the data center while maintaining similar performance to the one experienced by the user prior to the consolidation. Message manipulation services apply to the information exchange between different applications and strive to simplify integration of enterprise applications. Network management services include traditional management of network elements and provisioning functions, but newer trends focus on end-to-end provisioning of virtual services, including compute resources (servers and applications) as well as the network and its services. Additional network management functions use monitoring capabilities to create correlation information used for security analysis, capacity planning, and application optimization.

The storage network fabric provides virtualization services as well; VSANs to separate SAN islands using the same fabric, and storage virtualization by making multiple storage devices appear as a single device. Additional services focused on data replication enhance the replication process between distributed storage arrays and fabric routing services that allow devices on a SAN to communicate with devices on a different SAN (for example, a remote tape subsystem

FURTHER READING

- Cisco data center design best practices
cisco.com/go/datacenter
- Cisco Press book on data centers
Data Center Fundamentals, by Mauricio Arregoces and Maurizio Portolani
(Cisco Press - ISBN: 1587050234)

DATA CENTER FUNCTIONAL LAYERS AND NETWORK-BASED SERVICES



used to back up multiple storage arrays located on different SANs in distributed data centers). Because communication to the Ethernet and IP clouds is also required, the SAN fabric also provides gateway services that support iSCSI and FCIP.

The Infiniband network fabric provides several services for optimizing the operation of server clusters. The Infiniband fabric allows building of very high-speed, low latency, and low oversubscribed server clusters. Additional benefits to Infiniband include optimized interprocess communication via Remote Direct Memory Access (RDMA) and access to the LAN or SAN fabric. The Infiniband fabric also provides server virtualization services by seamlessly provisioning diskless servers (through boot services) using any combination of operating system, application, and storage.

Distributed Environments

Distributed data center environments help achieve high availability beyond single sites. In active-active scenarios the connectivity between data centers is critical, which should also take into account the need for Layer 2 adjacent clusters.

The design criteria for distributed data centers includes how many of them, which one houses the active instance of a specific application, the transport options between them, and the traffic types sharing the transport network. Other design factors include site selection requirements, whether the IP transport infrastructure is also used for data replication traffic, and the quantity and location of Internet access points (in distributed environments, multiple Internet access points are located in the data centers). ■

FIGURE 3 The switching fabric provides the physical media (Ethernet, Fibre Channel, Infiniband) and services for client-to-server, server-to-server, server-to-storage, and storage-to-storage connectivity.

180,000 IP Phones

Bank of America can roll out new voice processes overnight to 6,000 branches with a ubiquitous voice system.



GREATER FLEXIBILITY Migrating to a Cisco optical network gave Bank of America the flexibility to deploy IP communications at a reasonable cost.

By Rhonda Raider

In 2001, Bank of America made a decision about its network that, three years later, would give it the flexibility to change business processes in 6,000 branches across the US literally overnight. That decision was to migrate from a fully meshed ATM network to a privately owned optical network based on the Cisco ONS 15454 Multiservice Transport Platform. The decision to own rather than lease the optical network reflected the economics at the time. “The cost was roughly the same whether we secured our own dark fiber in metro areas or acquired private-line services,” says Craig Hinkley, senior vice president of network architecture and strategic direction. “I tip my hat to the bank executives who had the foresight and vision to recognize that a small incremental cost would buy us greater flexibility.”

That flexibility paid dividends in 2004 when Bank of America was able to use idle backbone capacity to carry voice traffic between metro areas—immediately and with no extra infrastructure cost. The bank signed a deal with Cisco partner EDS to deploy 180,000 Cisco IP phones and expects that every associate will have one within three years.

EDS was already in the process of transforming the bank’s data LAN and WAN to help position the bank to move to VoIP as part of its 10-year network management contract, which was signed in 2002, so it was a natural fit for EDS to then engage Cisco to assist with this project.

“What drove the decision to migrate to VoIP was the fact that we had more than 450 PBXs and 5,200 key [telephone] systems from multiple vendors, 100,000 Centrex lines, and 450 voice-mail systems, only 40 of which were networked so that branch associates could forward voice mails,” says Hinkley. “Bank of America has grown through mergers and acquisitions, so if there was a voice platform ever built, we had one.” While the bank had always budgeted to bring the data networks of acquired financial institutions into compliance, the same attention had not been paid to voice.

Ensuring Consistent Interactions

The trouble was that the different phone systems offered different capabilities, preventing Bank of America from gaining process efficiencies such as implementing a standard customer interaction model in all branches. The bank wanted to enhance customer service by defining models for customer-to-associate, associate-to-associate, and associate-to-customer conversations. The model would define such things as whether customer calls were routed to the branch manager or an associate, and whether the call would be sent to a teller or voice mail if the primary person was not available.

“In the past, we couldn’t even consider standardized processes because we didn’t have ubiquitous capabilities across all branches,” says Hinkley. “This lack of manageability had brought us to a crossroads, and we had to make a strategic decision to bring our voice platform to a current state.”

Faster Time to Market

Bank of America is today migrating to a ubiquitous voice environment by replacing its disparate PBXs with Cisco CallManager servers and Cisco Unity unified messaging servers in regional data centers.

The Cisco IP communications solution speeds time to market for new telephony features and functions available at branch stores. Suppose the bank decides that incoming customer calls will be routed to an associate with a particular role and rolled over to voicemail if the associate is not available.

“In the past, we either couldn’t accomplish this at all, or had to implement in one market at a time because each telephony system had different capabilities,” says Hinkley. “Now we can make changes to all branches literally overnight.”

In fact, in its pilot, Bank of America remotely disabled speaker phones in the meeting rooms of 50 branches in one night, something that would have taken six months in the previous voice environment with the same staffing levels. When the deployment is complete, Hinkley expects to be able to enable new telephony capabilities in all 6,000 stores overnight.

Working with EDS, Bank of America is today migrating to a ubiquitous voice environment by replacing its disparate PBXs with Cisco CallManager servers and Cisco Unity unified messaging servers in regional data centers.

“We recommended the Cisco platform for VoIP because of its manageability and flexibility,” says Travis Gilligan, EDS delivery manager for VoIP services engineering. “EDS is well underway in transforming

the Bank of America LAN and WAN to an end-to-end Cisco network, and using the same solution for VoIP would greatly simplify management.”

Another important advantage of the Cisco IP communications solution is that it increases business flexibility because it will enable the bank to use the Cisco IP phones on associates’ desks to deliver productivity-enhancing XML applications.

Incremental Innovation

As a Six Sigma company, Bank of America has commissioned its IT group to develop multigenerational plans (MGPs) for “incremental innovation.” VoIP is an enabler for incremental innovation, according to Hinkley. In generation 1, all Bank of America branches will acquire the ubiquitous features and functions that enable a common business telephony operating model and lower operating costs. In generation 2, Bank of America will begin capitalizing the ubiquitous infrastructure to improve business processes in ways that increase revenue, decrease costs, improve interaction capabilities with customers, and increase wallet share.

Enabling a Flexible Workspace

Those business-process improvements have already begun. For example, the Corporate Workplace Group

VoIP as a Risk Mitigator

When companies evaluate the advantages of IP telephony, many consider the potential process improvements for customer service or productivity. Bank of America also considered its role in mitigating the risk in the previous voice environment.

The Bank of America IT department conducted a risk assessment of the current voice environment, including manageability and visibility, fault management, performance management, and capacity planning. One risk was the time and cost needed to comply with legislation, such as a recent requirement that financial institutions must provide caller ID when they make outgoing calls. “The lack of manageability and visibility into our system made it difficult to determine what we had to do to become compliant,” says Hinkley. “Once we figured it out, compliance was time-consuming and required expensive truck rolls. To mitigate risk we wanted a ubiquitous voice platform that would enable us to quickly assess our risk for any new legislation, and then comply quickly—preferably through remote implementation.”

Just How Do You Deploy 180,000 IP Phones?

With 180,000 Cisco IP phones, the Bank of America VoIP deployment is the world's largest to date. Often asked about the logistics, Hinkley emphasizes that deployment includes people and processes as well as technology. "VoIP is not simply a convergence of voice and data technologies," he explains. "Rather it's a convergence of the people and processes that support the technologies." That sentiment is echoed by Jennifer Crum, EDS program director for the Bank of America VoIP program. "This deployment is succeeding because of close teamwork between Bank of America, Cisco, SBC, and many resources within EDS," she says.

To approach the challenge, Hinkley divides IP telephony into two components: VoIP Infrastructure Services (VIS) and VoIP Application Services (VAS). One set of people and processes supports VIS, or the traditional data networking elements that go into a VoIP solution, such as the WAN, LAN, quality of service (QoS), and the IP phones themselves. Another set of people and processes support VAS, or the features, functions, computer-telephony integration, and application-layer compatibility that transform an IP phone into something that provides business value. "You can think of an IP phone as simply a three-port switch wrapped in plastic with a handset," says Hinkley. "By adopting VoIP, we're extending the

demarcation of the network out to the phone, which becomes a network interface for the associate. The integration and interoperability of phone to network edge is critical to deliver voice services.

Hinkley notes that VIS and VAS distinction more clearly reflects roles and responsibilities than the traditional voice and data distinction. "The data group is responsible for VIS, or the infrastructure that supports service requirements for real-time communications. The voice team is responsible for VAS, or using its voice skills set to link the capabilities of the voice system to the business need."

Using the VIS/VAS model, the Bank of America IT group divided the deployment challenge into multiple voice domains: for branches, call center traders, banking centers, the regional enterprise, and the metropolitan optical network. The IT group built a separate business case for each domain. For the branch domain, for example, it determined the number of branches, needed features and functions, and scalability and availability requirements. "Once we figured this out for a building block of 500 branches, we could scale it for 6000 branches by replicating the building block," says Hinkley. By dividing the challenge into domains, Hinkley's group was able to deploy incrementally—that is, deploy it in some domains while still making the business case for others.

wanted its personnel team in Charlotte, North Carolina, to enjoy the freedom of a flexible workspace, in which employees choose a workspace that best suits the day's business needs: no assigned seating.

"Using Cisco IP communications, we met the business need with wireless access and the extension mobility feature that lets employees log in to personalize any phone as their own," says Hinkley. Technology enablers include Cisco Aironet wireless access points, Cisco 7920 wireless IP phones, Cisco 7960 IP phones, and Cisco IP Communicator on laptops.

Office-in-a-Box for Disaster Response

Bank of America IT also used the Cisco IP communications solution to meet the needs of the bank's business continuity team. "At Time of Disaster," an office-in-a-box, contains Cisco Aironet wireless access points and other Cisco networking solutions that enable disaster recovery for both voice and data.

"When our voice circuits were provided by a carrier, our only recourse during an outage was to use cell phones," says Hinkley. "By adopting VoIP, we've gained the same control over business continuity for voice that we have for data."

Reducing Operating Costs

Hinkley notes that VoIP is reducing operating costs even as it enables process change. "We're considering reducing the total cost of ownership of peripherals on the associate's desk by delivering voice to the PC, eliminating the need for a separate phone," says Hinkley. "The end result of VoIP is providing the same or enhanced capabilities as traditional voice systems, at a significantly reduced operating cost."

EDS owns the bank's voice-and-data infrastructure, and Bank of America pays for services consumed, such as LAN ports, WAN site connectivity, and IP phone usage. This pay-as-you-go model yielded immediate cost savings. "We didn't have to wait for the transformation to VoIP to be complete before we began seeing cost savings," Hinkley notes.

The cost of moves, adds, and changes has also dropped. The IT group is still deciding whether moves will be handled by the associates themselves, a desktop service provider, or EDS. "Regardless, the price point is significantly lower because nobody has to go to the back room to punch down blocks," says Hinkley.

Continued on page 89

Getting the “Message”

An Overview of Cisco Application-Oriented Networking

By Gail Meredith Otteson

Message to applications: Cisco speaks your language.

Application integration is a well-known, necessary evil in today’s enterprises. It takes a prodigious effort to develop user-friendly, “flow-through” systems that support efficient business processes. Traditional approaches are slow—typically spanning three to five years—and costly, consuming 40 to 80 percent of IT budgets.

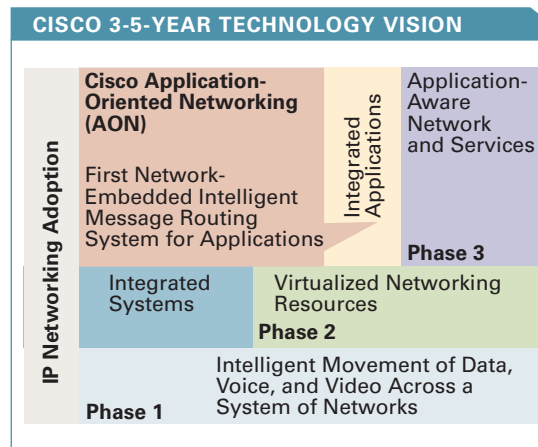
Just as networks used many protocols a decade ago, siloed applications today speak a variety of often proprietary messaging languages. Using custom coding and point products, a multibillion dollar software and service industry links applications together into cohesive business systems. Every project is unique and therefore expensive. Development of middleware and emerging application-messaging standards have begun to address the problem, but they can also exacerbate it, increasing IT complexity with layers and systems that reduce business agility and contribute to soaring costs.

IT managers would rather increase business agility with a radically different application-integration paradigm that reduces IT complexity, takes less time, and frees budgets for strategic purposes. The Cisco Application-Oriented Network (AON) offers that approach. It is the first and only solution that embeds essential application message processing functions into the network fabric, offloading them from servers and middleware for greater performance, consistency, and visibility.

“The network is the only thing that touches every element in the IT infrastructure,” says Rob Redford, vice president of Worldwide Product and Technology Marketing at Cisco. “That’s why it’s the perfect place for functions that promote true collaboration among applications.”

AON is the first Phase 3 deliverable in the Cisco vision of the Intelligent Information Network (IIN), adding application-aware capabilities and services to the network (Figure 1).

For more information on IIN, see the *Packet* special report, “A Smarter Way to Network” (Third Quarter 2004).



From Ad-Hoc to Architecture

What distinguishes AON from standard application integration is an architectural approach that considers the entire IT infrastructure with an eye toward simplifying systems and reducing layers. This differs from ad-hoc approaches that “glue together” point solutions, such as translating a message protocol between application A and application B.

The Cisco AON architecture integrates application messaging overlays with packet networks, achieving a unified system that can operate more efficiently and consistently than previously isolated layers (Figure 2, page 62). “Problems exist because of isolation between the application and network layers,” says Redford. “We’ve already demonstrated the value of integrating layers with our IP communications, storage networking, and security solutions. Now we’re creating an application Internet that addresses the complexities of heterogeneous application environments.”

AON simplifies application message processing, because it requires no *changes to application code*. As a “universal translator” for applications, AON accepts and delivers messages in their native formats, performing necessary policy decisions, field transformations, and message protocol translations in between. (See sidebar on page 63 for a list of application-level standards and protocols that initial versions of AON can interpret and translate.)

For example, AON can route a purchase order to a certain server based on its monetary value and simultaneously transform it into the correct format for the

FIGURE 1 Organizations can take advantage of the convergence of applications on IP to create new applications and uses of applications that were previously impossible. Using the network to help understand and map business policies to application usage, businesses can simplify the application infrastructure.

destination application. With the network doing routine content and context-based translation, transformation, and sorting—activities common to all application integration efforts—end users can better invest their application integration budgets in specialized, advanced functions on servers and middleware.

In addition to translating and transforming messages between applications, AON technology does the following:

- **Enforces consistent security policies** for application access and information exchange
- **Provides visibility of information flows** by monitoring and filtering messages for both business and infrastructure purposes
- **Optimizes application performance** by offloading application-level load balancing, application-level security, and Extensible Markup Language (XML) operations from servers into the network, and offering message-level cache and compression services

AON Components

Application managers can formulate business policies as a series of *bladelets*, or program objects, that are sequenced and passed to the network manager, who downloads policies into AON elements.

A bladelet is an object that represents a specific action, such as “validate digital signature.” Cisco AON Development Studio (ADS) software is a Windows-based

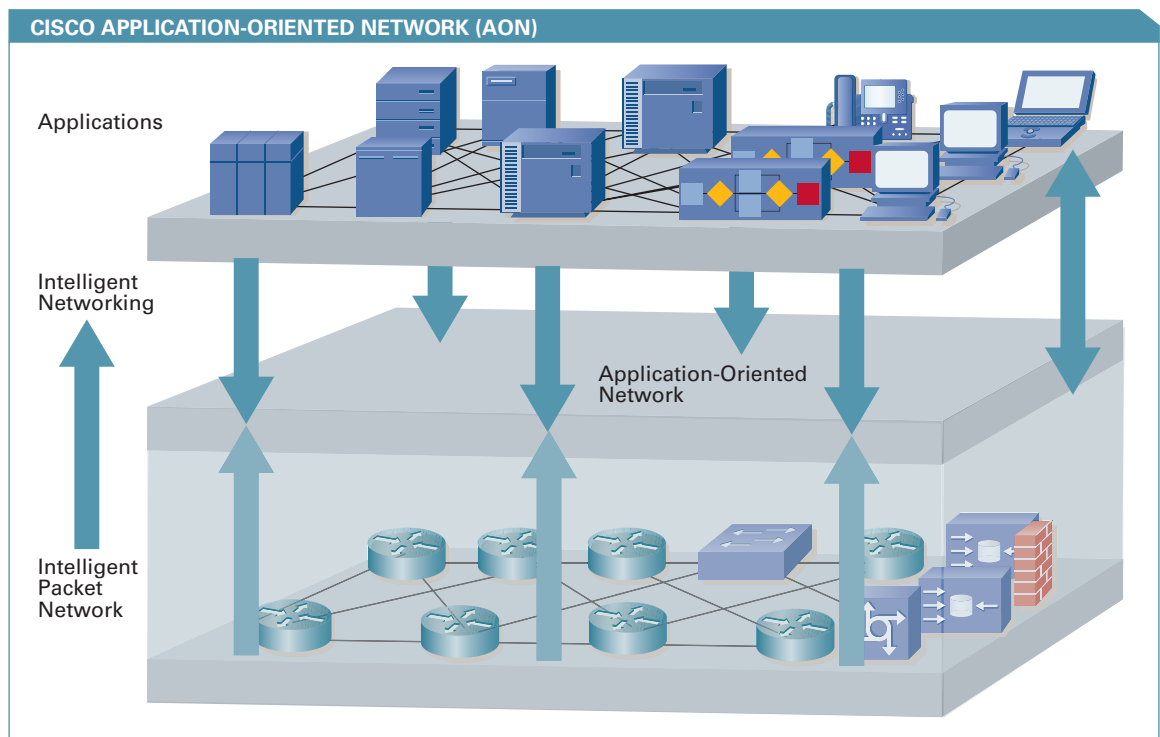
tool for developers to configure bladelet templates. Included is the Bladelet Developer Kit (BDK) for creating custom bladelets. Cisco ADS also includes a custom adapter kit that enables development of adapters that perform unique functions such as metadata interfaces, error recovery and logging, or encryption.

Cisco AON Management Console (AMC) software is a Linux-based Web application for centralized management of the Cisco AON system. Administrators use Cisco AMC to sequence bladelets and adapters into processes that represent business policies, then download them to AON nodes. The console also manages security keys and certificates, monitors node events and logs, and interfaces with each AON element in the network.

Initially, Cisco is offering two AON nodes as blades residing in Cisco Catalyst 6500 Series Switches (in network cores) and Cisco Integrated Services Routers (at network end points and remote sites). Extending AON beyond the data center into branch offices and even customer sites facilitates higher performance levels than a central hub-based solution can achieve. Both AON modules can accept and pass traffic to other service modules in the chassis for additional processing such as firewall services, encryption/decryption, and content switching.

The AON modules have hardware acceleration features for processor-intensive actions. The switch or router detects incoming message traffic and diverts it into the AON blade, where packets are reassembled into messages. The AON engine then applies policies

FIGURE 2 Cisco AON simplifies the IT infrastructure through careful integration of network and application IT layers.



Lingua Franca

The first versions of Cisco AON products and partner add-ons support communication between these application-level protocols:

Protocols	HTTP, HTTPS, TIBCO Enterprise Messaging Service (EMS), Websphere MQ, WebSphere Java Messaging Service (JMS)
Database	Oracle 9I (9.2), Sybase 12.5.1, Java Database Connectivity (JDBC), Lightweight Directory Access Protocol (LDAP)
Security	WS-Security Profile, Secure Sockets Layer (SSL), Public-Key Cryptography Standards (PKCS), Triple Data Encryption Standard (3DES), Advanced Encryption Standard (AES), RSA, Security Assertion Markup Language (SAML), Secure Hash Algorithm (SHA-1), Secure Multipurpose Internet Mail Extensions (S/MIME)
Extensible Markup Language (XML) Processing	XML encryption, XML digital signatures, Schema Validation, Extensible Style Language Transformation (XSLT) 1.0, Simple Object Access Protocol (SOAP) 1.1 and 1.2, Xpath 1.0
Transaction	Financial Information Exchange (FIX)

as defined by bladelet logic, converts messages back into packets, and sends them to their destination. This architecture protects wire-speed performance of normal Layer 2 and Layer 3 traffic.

Greater Security

Cisco AON enables hardware-accelerated, application-level security within the network itself. AON increases the defensive posture of the network. Where middleware and servers can be configured to bypass security appliances, embedded AON enables consistent policy enforcement because all messages must suffer inspection. It can detect and isolate suspicious messages, then activate the security features of the network to block future instances. It can decrypt and encrypt portions of messages (such as credit card numbers) for examination.

For example, AON can detect that a purchase requisition number has the wrong number of digits or that a part number is invalid. The network generates an alert and tells the firewall to block all traffic from the source.

Information Visibility

AON provides real-time event logging, enabling enterprises to track transactions and respond to changing business conditions. "People can track individual transactions, but the distributed nature of applications and processes makes it difficult to get an aggregate picture of business activity," says Stephen Cho, senior director of product management in the newly formed AON Business Unit at Cisco. "AON fixes that with real-time information that allows enterprises to tune their operations to meet service-level agreements [SLAs]."

AON nodes can be configured to act as sensors that can capture, process, and log highly granular information about application messages. For example, on-time delivery is critical to a shipping company's reputation and ongoing business. Before AON, the company could track individual packages, but couldn't view aggregate transaction behavior throughout the day, making it difficult to allocate resources on the fly. AON looks into a message field and builds data that allows the company to track all packages shipped by a particular customer or to a certain location. With this information, the company can adjust resources to assure on-time delivery.

Information visibility can be configured in implicit or explicit mode. Implicit mode is appropriate for protocol translation in transit, such as a bank exchange. Member banks use the Financial Information Exchange (FIX) messaging protocol, but several versions exist. Rather than force all banks to install the same version of FIX, AON automatically translates messages into the version used by the destination bank.

Explicit mode works well in service virtualization. Because Cisco AON understands both the content and context of application messages, it can act as a proxy that provides an abstraction layer for endpoint applications and apply policies without the endpoints being aware of the intermediary. For example, a business-to-business (B2B) transaction for an online travel service may send messages from a corporate travel site to a vendor. AON knows that the customer wants the lowest-price auto rental available. It learns that one vendor offers better rates on the weekend, while another offers a discount during the week. Based on the dates of travel, AON routes the request to the appropriate vendor.

AON White Paper for Network Managers

Interested in learning more about the evolution of application integration and middleware? "The Application Infrastructure Primer for Network Professionals" examines areas of commonality between middleware and network infrastructure and describes how it will soon make economic sense for parts of these common functions to come closer together. Read more at cisco.com/packet/173_7c1.

Application Optimization

AON enhances message-handling performance, but it must do so without adversely affecting network availability or performance of other network traffic. Cisco enables enterprise-class performance with the following features:

- Hardware acceleration for performance-intensive actions such as security and XML operations
- Caching and compression to improve response times and conserve network bandwidth usage
- Placing an AON node in front of an application cluster can assist in high availability and load-balancing services to improve overall IT system performance

Found in Translation

Analysts predict that Cisco AON will gain traction because it offers enterprises an incremental path toward simplifying and speeding application integration. "There are a lot of ramifications to our approach," says Taf Anthias, vice president and general manager of the AON Business Unit at Cisco. "We leave the application infrastructure alone. We're making the network programmable so that our partners and service organizations can extend the platform over time by adding functionality."

Cisco ecosystem partners such as IBM and TIBCO are creating custom adapters and bladelets that perform advanced functions for specific applications. Systems integrators can focus on offering value-added services and expertise beyond essential functionality, while middleware vendors can develop solutions that build upon the power of AON.

"No question—this is a game changer," says Roy Schulte, Gartner application integration analyst. "Everyone's integration strategies will have to take this [AON] concept into account. It's a big change." ■

FURTHER READING

- Cisco Application-Oriented Networking
cisco.com/go/aon
- AON Product Portfolio
cisco.com/packet/173_7c2
- AON White Papers
cisco.com/packet/173_7c3
- Cisco Intelligent Information Networking
cisco.com/go/iin

GL Communications

Voice Goes Wireless

Delivering Mobile VoIP, Data, and Location Services to Diverse Wi-Fi Environments



Carnival Cruise Lines

EARLY ADOPTER Carnival Cruise Lines is the first cruise line to deploy a Cisco wireless LAN with voice over IP onboard a cruise ship.

By Gene Knauer

With today's laptops, handsets, and personal digital assistants (PDAs) able to access many of the same network applications, users need access to applications to be productive and responsive. And these users want high-speed access to voice and data whether mobile or at their offices. Wireless LAN enhancements make support for voice possible, enabling mobile users to access enterprise messaging and applications. Mobile devices are also improving access.

Cisco has been on the front lines of WLAN innovation with products for consumers, enterprises, and service providers. Foremost among the gear are the Cisco Wireless IP Phone 7920 and the newly announced Cisco 2700 Series Wireless Location Appliance.

The IEEE 802.11b-compliant 7920 Series phone provides comprehensive voice communications with Cisco CallManager and Cisco Aironet Series Wi-Fi access points. The new Cisco 2700 Series Wireless Location Appliance, a component of the Cisco Integrated Wireless Network—Cisco Centralized WLAN Solution, lets businesses of all sizes track any mobile device, from wireless laptops and PDAs to devices equipped with IEEE 802.11 radio frequency identification (RFID) tags. The Cisco 2700 Series can be used to track assets, manage inventory, automate workflow, enable enhanced 911 (E911) communications, improve network performance based on location-based usage trending, and tighten WLAN security. Moreover, the 2700 Series has a rich application programming interface (API) that integrates with third-party applications for a variety of location-based applications.

Continued on page 67

Voice Goes Wireless, Continued from page 65

Converged Voice, Data, and Cellular Services

Nearly half of all large and small enterprises in the US have deployed WLANs in some form, according to Michael Ladam, senior research analyst at Stratcast. “The beauty of Wi-Fi is that the enterprise can make its coverage ubiquitous, whereas cellular coverage is spotty in many buildings,” says Ladam. “With ubiquitous access, people using their PDAs and 3G phones for cellular voice can use them for IP data and IP telephony, too. Dual VoIP and cellular phones are going to be widespread within three years.”

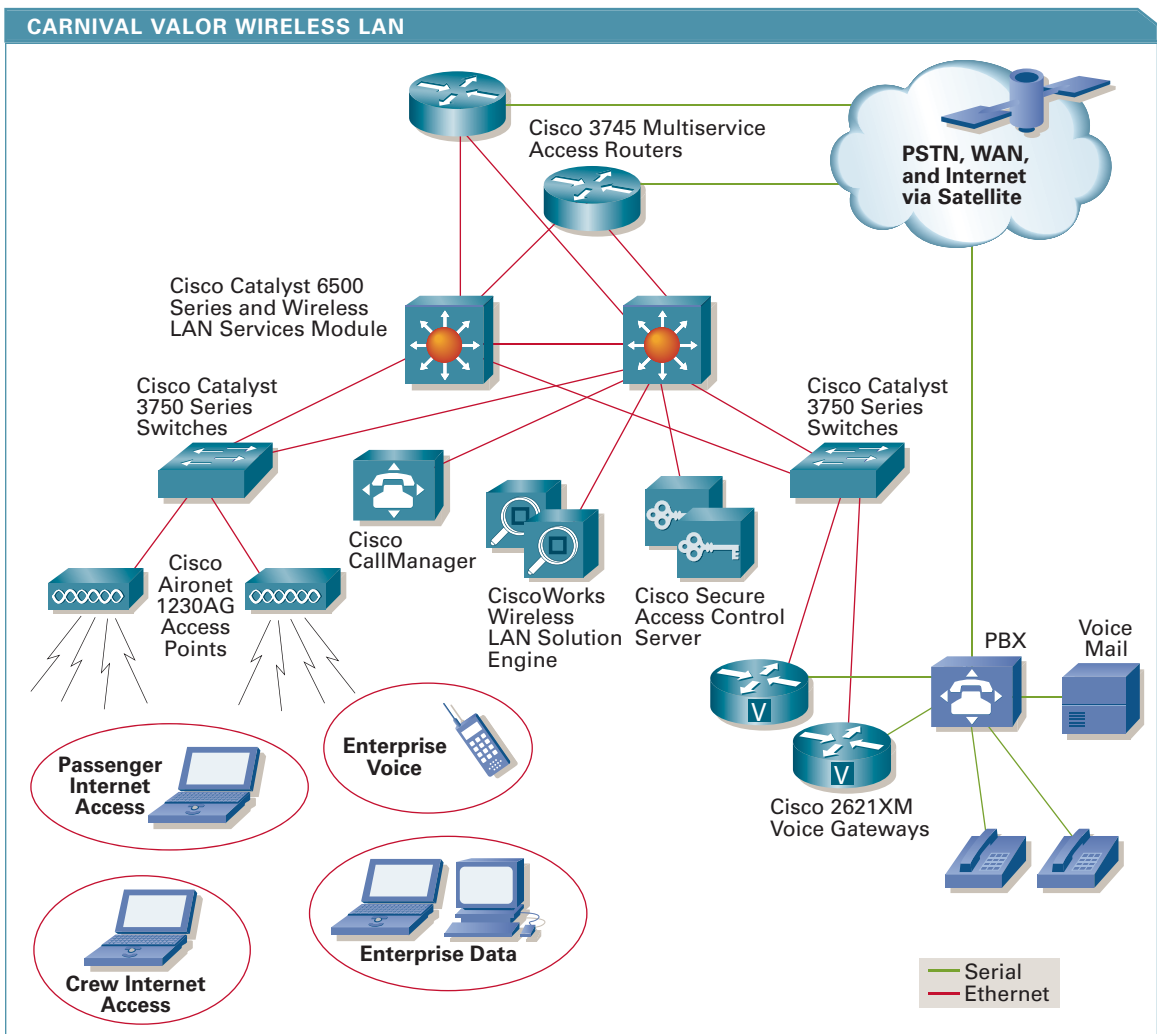
Ladam includes hospitals, factories, hotels, and college campuses in the list of early adopters of VoIP over WLAN because of an array of productivity benefits, from locating a doctor or piece of medical equipment to broadcasting a “Class Canceled” message prior to a college lecture.

“The effect of VoIP together with wireless data applications in a WLAN is like moving from e-mail to instant messaging,” says Ladam. “For some, voice over Wi-Fi will make clear business sense right away, while for others it will be a nice but not crucial feature on a cell phone, like a camera.”

VoIP over WLAN Proves Itself “Ship-Shape”

In December 2004, Miami, Florida-based Carnival Cruise Lines became the first cruise line to deploy a Cisco WLAN with full VoIP capabilities on a newly built cruise ship, the *Carnival Valor*.

“Initially, we had planned to increase the number of workstations in our onboard Internet cafes and to expand Internet access to the staterooms using traditional Cat 5 cabling,” says Tom McCormick, manager of network engineering for Carnival Cruise Lines. “However, using Cisco wireless technology we are able provide wireless data access bow-to-stern and, as an added benefit, we were also able to introduce mobile VoIP on the same infrastructure.”



BOW-TO-STERN WIRELESS All shipboard communications are switched within the WLAN, and a geostationary satellite link enables Internet access.

Adding Location Services to the WLAN

Among the early adopters of WLANs, some environments can derive important benefits from another WLAN component: location-based solutions. The recently introduced Cisco Wireless Location Appliance (2700 Series) a component of the Cisco Integrated Wireless Network—Cisco Centralized WLAN Solution is the first such device to be integrated directly into a WLAN infrastructure. It uses existing Cisco Aironet lightweight access points and Cisco Wireless LAN Controllers, eliminating the need for separate, add-on RFID readers and separate client software.

“Each Cisco 2700 Series appliance can locate up to 1,500 simultaneous Wi-Fi devices, including Wi-Fi active RFID tags, rogue clients, and rogue access points,” says Anita Pandey, product manager for the Cisco 2700 Series. “So it lets you track people and assets for purposes of inventory management, streamlining workflow, and security enforcement,” continues Pandey.

Using the Cisco 2700 Series, hospitals with WLANs can track the location of critical care machines, such as blood and gas monitors and pumps, which move about as-needed on carts. In-patients with Wi-Fi-enabled RFID tags can



REAL-TIME TRACKING The Cisco 2700 Series maps the location of Wi-Fi-enabled devices.

be tracked on the hospital premises. High-value assets in stores, factories, schools, companies, museums, and other locations can be similarly tracked in real time.

Third-party developers can use the Cisco 2700 Series' open application programming interfaces (APIs) to add location based alerts, alarms, and other features, or to integrate with existing enterprise resource planning (ERP) and other applications. The Cisco 2700 Series' patent-pending radio frequency (RF) fingerprinting technology can determine the location of a wireless client to within a few meters. With RF fingerprinting, a grid is created that identifies how every part of the floorplan looks to all access points. Factors such as RF attenuation and reflection are taken into account. Real world information gathered by access points is compared to the grid to determine the device's location to within a few meters.

The Carnival network engineering team of Ying-Yuang Chen and Vicki Ramirez and Cisco Advanced Services engineers Ken Gholston, Michael Kang, and Tony Hong, installed the WLAN infrastructure on the 110,000-ton *Carnival Valor* as the ship was being constructed in Monfalcone, Italy. The installation included 217 wireless access points located throughout the ship, which provide end-to-end voice and data coverage on all passenger decks, crew decks, the engine rooms, and the bridge.

“The placement of the access points and the installation of the wiring closets with the access switches was one of the biggest design challenges because of the unique environment of a large cruise ship,” recalls Chen, of Carnival Cruise Lines. “Thick steel bulkheads mandated by US Coast Guard regulations separate each fire zone. Because the wireless IP phones have to roam between subnets as a caller passes between the different fire zones, we had to carefully position and tune each access point to make sure there

were no dead zones or dropped calls. In addition, many different types of shipboard machinery are in use that can cause radio interference, which must be taken into account during the initial network design.”

Cisco Catalyst 6500 Series switches provide core switching and 16 distribution hub closets with a total of 31 Cisco Catalyst 3750 Series edge switches with power over Ethernet (PoE) connect the access points through subnets to the WLAN (see figure, page 67).

The crew and passengers can surf the Internet wirelessly from any staterooms using laptops or PDAs with any wireless card, or at separate crew and passenger Internet cafes with Wi-Fi-accessible workstations.

Travel agents and travel writers who were treated to the initial voyage on the *Carnival Valor* were able to test the ubiquitous wireless access first hand: they could file their articles and check e-mail throughout the voyage using the WLAN infrastructure.

“Now, for those who won't go on a vacation because of concerns over being disconnected from the

Internet, they can go on a cruise and stay connected; in their room, by the pool, anywhere aboard,” says McCormick.

Of the ship’s 1,000 crew members, 350 use the Cisco Wireless IP phones. This ranges from cabin stewards—who can receive a special request for an extra pillow from a passenger on their VoIP phone or via text message—to the ship’s Captain—who can receive calls, text pages and access voice mail anywhere on the ship via the wireless VoIP phone, and check data in navigational and administrative applications via a wireless laptop. The IP phones have been integrated into the Siemens PBX that was part of the ship’s original design. Cisco CallManager integrates call processing from the IP phones with calls made on the Siemens analog phones used throughout the ship.

“The WLAN on the *Carnival Valor* is useful in serving up both the voice and data environment and by making many applications available to the handset,” says McCormick.

Wireless allowed the *Carnival Valor* to improve the guest experience, as well as improve the crews overall productivity. “Mobile voice and data applications open up so many possibilities for us, including the easy deployment of new systems to different parts of the ship without the cost of new cabling, which can be excessive within a shipboard environment.” says McCormick. “The largest benefit was the increased productivity that is provided by the phones themselves. Our ships are big, and actually locating the person that you need to talk to is often difficult. The wireless phones allow the type of instant communications on the ship that we usually take for granted on shore.”

The ship also installed 14 Cisco Aironet 1300 Series Outdoor Access Point/Bridges to facilitate ship-to-shore communication while in port. These bridges allow the ship to seamlessly connect to the pier side networks so that guest embarkation can take place anywhere the ship goes. This flexibility is important to Carnival because the ship’s itinerary can change at a moments notice, due to weather, tides, or other reasons. For example, a recent high tide situation combined with power lines that traverse over the Mississippi River prevented the Carnival Conquest from being able to reach its usual embarkation port of New Orleans last year. The ship was diverted to a nearby port that had no existing infrastructure for the ship’s embarkation process. The wireless connectivity

provided by the Aironet 1300 Series outdoor access point/bridges allowed the crew to process 2,000 passengers at the temporary location without a problem.

“In this particular case, it made a difference in our ability to keep the impact to our guests to a minimum,” says McCormick.

Evolving Standards

Standards for cellular over VoIP and Wi-Fi are still in progress. “You can’t buy a cellular handset with VoIP over Wi-Fi and use it with different access point vendors and expect seamless reliability right now,” says Michael Ladam. “Common standards are evolving. Also, chip manufacturers have yet to develop chips with dual capabilities that are designed for use on a handset, PDA, or laptop. In three years, you’ll see these capabilities on phones that automatically tap into a Wi-Fi network at home or at the office for voice and data services and, when beyond these WLANs, they will use their cellular network to stay connected,” he continues.

As with other mobile network applications, Ladam predicts that university, healthcare, and hospitality environments like cruise ships and hotels will demonstrate the full value of converged voice and data applications over a WLAN before mass adoption occurs. Because of the IEEE 802.11i specifications for voice security, including encryption, Ladam also believes that WLAN voice and data applications are as secure as conversations over a land line and authentication and encryption can be much more stringent than in most wireline networks. ■

FURTHER READING

- Cisco Wireless IP Phone 7920 cisco.com/packet/173_7b1
- Cisco 2700 Series Wireless Location Appliance cisco.com/packet/173_7b2
- Cisco Integrated Wireless Network cisco.com/go/integratedwireless

IPTV/Video over Broadband

New Opportunity, Revenues, and Loyalty . . . Enabled by IP Next-Generation Networks

By Janet Kreiling

The IPTV video opportunity—providing video in digital form over a broadband IP network—is becoming a “must do” for many service providers to succeed in the convergence of consumer services. For telecom companies, it is the latest source for new revenues as income from fixed voice lines shrink. Cable companies are already finding IP the most efficient path to digital service convergence.

This opportunity, which includes digital broadcast video, video basic and pay per view (PPV) premium channels, video on demand (VOD), network-based personal video recording (nPVR), and the emerging high-definition TV (HDTV), can increase profitability by boosting average revenues per user (ARPU), expanding a provider’s footprint, reducing churn, and taking advantage of a converged infrastructure. VOD already has proven to be a real service differentiator for many providers.

The market is already growing. According to Gerry Kaufhold, principal analyst for multimedia broadband and consumer content at In-Stat, worldwide revenues from digital video offered by telecom companies is expected to rise from US\$1.1 billion in 2005 to \$5.5 billion in 2008. “VOD in particular is demonstrating a take rate of about 20 percent,” says Kaufhold.

Adds Pankaj Gupta, senior manager of broadband marketing at Cisco, “IPTV/video over broadband with strong linkages to the network will be one of the keys to the success of the next-generation, converged network.”

Video Application to Network Linkages

To attain its benefits, service providers must deploy video to scale with a proven infrastructure. Video service will benefit from a core and aggregation Gigabit Ethernet IP infrastructure, so it can be scaled to meet video’s huge demands for bandwidth. It should take advantage of video application to network linkages, such as quality of service (QoS), security, routing protocols, and rapid failover. It must also adapt to any access technology, whether DSL, Ethernet, or coax/HFC.

Cisco has considerable experience in deploying video over IP networks, Gupta points out, having more than 10 million subscribers of VOD already being served by production-scale networks. Furthermore, he says, “Cisco has innovated key network technologies that, when linked to video service, enable providers to deliver it cost effectively and with great customer experience and satisfaction.”

The following innovative network capabilities offered by Cisco, are “crucial to a successful digital video offering,” notes Gupta:



- *Video-optimized asymmetric networking*, which reduces total cost of ownership by enabling providers to configure only the bandwidth needed across the core and distribution networks for downstream and upstream PPV and VOD communications
- *Enhanced dynamic IP Multicast and Broadcast Source Redundancy for live TV, PPV, and VOD*, which minimizes the bandwidth needed to serve customers and supports source redundancy, thus helping ensure a quality experience
- *Source-Specific Multicast (SSM)*, which inhibits denial of service (DoS) attacks from unwanted sources, and is easier to install and manage for video broadcast service

Lots of Bandwidth and Asymmetric Networking

“One of the biggest requirements for interactive and on-demand digital video is much more bandwidth,” says David Benham, senior manager for video technologies and solutions development at Cisco. “Delivering video, particularly VOD and nPVR to homes, will dwarf any other service in bandwidth consumption. For example, supporting only a modest fraction of subscribers with one high-definition movie and standard-definition streams versus two to three standard-definition movie streams at the same time will require an order of magnitude more capacity in the network than just the voice and data services do,” adds Benham (see sidebar, “Example Demand Projections for a Video Service Provider,” page 73).

The emergence of higher DSL speeds with ADSL2+, generally 20 to 24 Mbit/s per DSL subscriber link, and advanced video

codecs have removed the last barriers to supporting multiple, high-quality video streams to each home.

“Bandwidth requirements will accelerate as the on-demand service becomes successful. This is because the video streams associated with on-demand services are unicast while the video streams associated with broadcast services are multicast,” explains Benham. “Thus, the amount of bandwidth required in the aggregation and distribution networks to carry on-demand streams accelerates faster than what is required to add channels to the broadcast service.”

A successful VOD service will consume four to seven times the bandwidth than that of the video broadcast service, adding up to 4 to 8 Gbit/s of downstream traffic to each video central office, on average, and 100 or more Gbit/s to the headend. Yet even this mass of network bandwidth can be provided cost effectively.

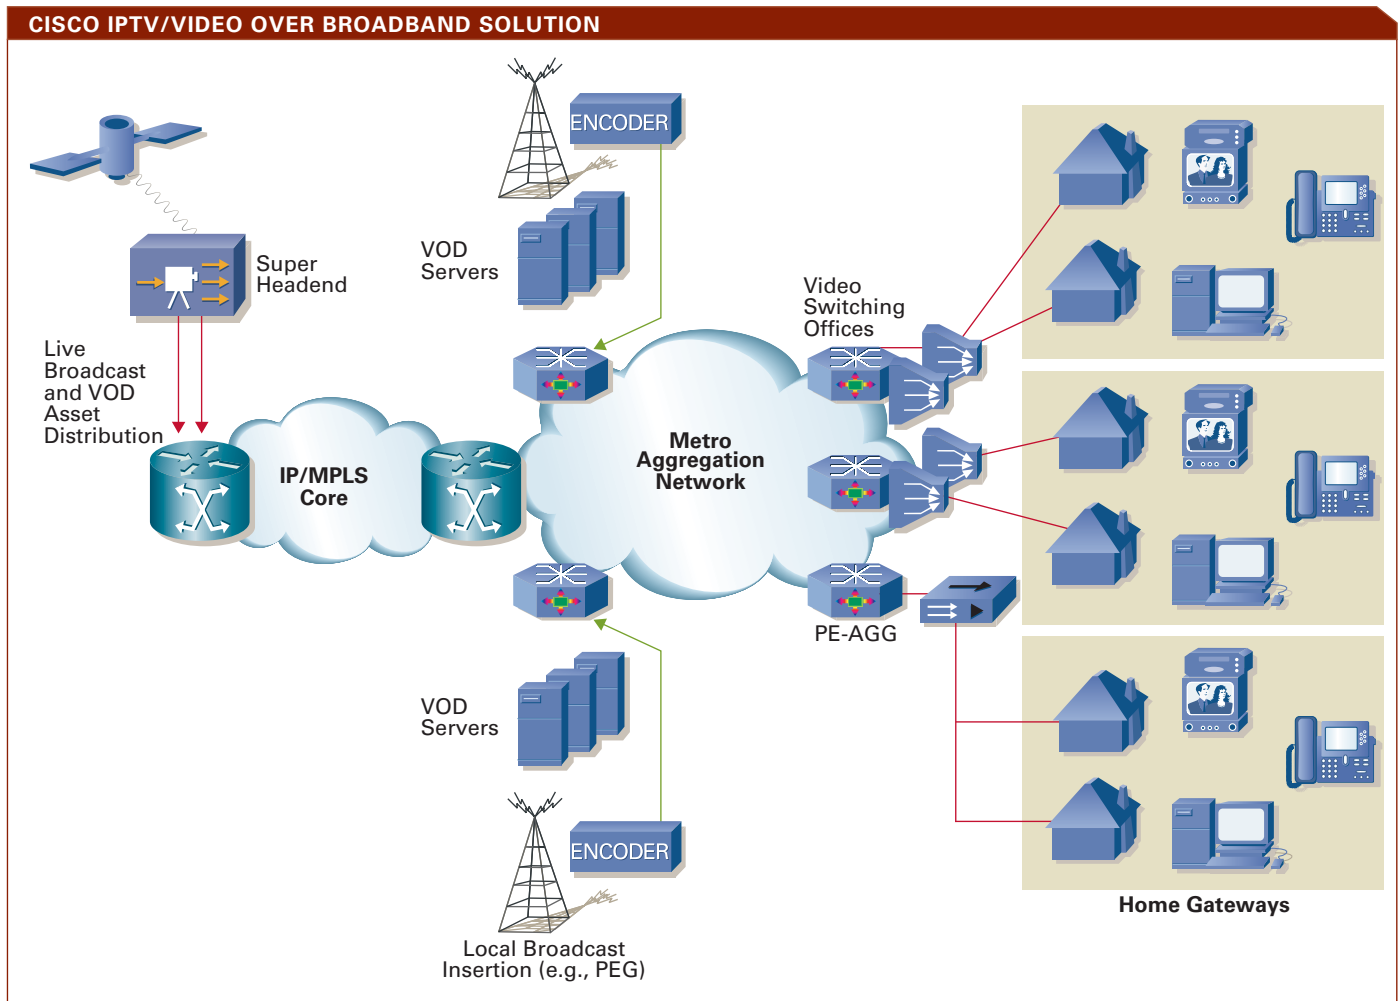
The technology to use, says Benham, is asymmetric networking, making more bandwidth available downstream than up. Video traffic is nearly all in one direction: downstream. Any upstream traffic consists of a few packets requesting channels, titles, or for stream

control. If a service provider deploys symmetrical bandwidth to meet the demand for video, a great deal of upstream capacity will be left expensively idle and result in unnecessary capital expenditures (CapEx). “Asymmetric networking establishes separate paths for up and downstream links that can be dynamically and flexibly configured,” says Benham. “Both up and down signals can travel over any available paths of sufficient bandwidth in an Ethernet IP network.”

Either up or downstream capacities can grow in bandwidth, as needed, and independently for the most efficient use of CapEx. Furthermore, because of the asymmetric networking capability, the cost of the distribution network can be reduced by incorporating less expensive unidirectional links in the transport along with one or more bidirectional links, Benham adds.

Bidirectional and unidirectional connectivity functioning together as one asymmetrical network is made possible through the use of Unidirectional Link Routing (UDLR). It provides bidirectional communication between any two nodes via a combination of unidirectional links on one path and by tunneling rel-

DIGITAL VIDEO NETWORK
Technology already available via linkages between the video layer and the network, such as asymmetric networking and dynamic IP Multicast, reduces capital and operating expenses from the head-end to the home gateway and improves customer satisfaction as well.



atively fewer packets upstream through another path. UDLR enables the cost benefits of asymmetric networking while preserving the operational benefits of a bidirectional network.

Given the enormous amount of bandwidth potentially “stranded” in a symmetric network by VOD especially, asymmetric networking is crucial to delivering the service economically, according to Benham.

IP Multicast and Broadcast Source Redundancy

When a video program is sent only to certain viewers through broadcast channels, for example, a multicast tree configures itself from the headend out through the video-equipped central offices and on to individual viewers’ homes. This selective delivery is enabled by IP Multicast routing. When subscribers choose to watch a subscription program, their set-top box asks the network to add a “leaf” to the multicast tree over the least-cost IP route.

“With dynamic IP Multicast, the formation of trees is dynamic and enables providers to most cost effectively use network bandwidth,” says Benham. “Furthermore, the service provider can also take advantage of ‘anycast’ support in the Cisco IP Multicast solution for load balancing of video encoders/streamers and fast failover of video broadcast sources.”

Cisco IOS Software creates different multicast trees to serve a subset of the subscriber set-top boxes for each of the different broadcast sources located in geographically separate video headends, which are each transmitting the same sets of channels using the same host address. If a broadcast source fails, IOS Software detects the loss of that one broadcast source, automatically switches that source’s multicast trunk and its branches, and merges it over to a remaining multicast tree. Only a subset of the subscribers might experience a momentary outage, but nearly instantaneously they will receive a healthy stream from the second broadcast source dynamically.

Trees and failover linkages can extend over considerable geographic distances. For example, notes Gupta, a server in Boston might failover to one in Atlanta, incorporating Boston’s branches into its own tree with almost imperceptible interruption of service.

Source-Specific Multicast

One of the innovative technologies that helps make dynamic multicasting work is *Protocol Independent Multicast (PIM) SSM*, a feature of Cisco IOS Software that simplifies the creation of an IP Multicast tree for each source by eliminating the need to define a rendezvous point for multicast forwarding. Standard SSM relies on Internet General Management Protocol version 3 (IGMPv3), a signaling protocol not yet supported by most of today’s set-top boxes.

Example Demand Projections for a Video Service Provider

The following projection is based on a real network and reasonable forecasts of video and VOD peak take rates. The Gigabit Ethernet network, converged for triple play, includes a super headend that receives signals from a broadband source and satellite and feeds them out to 400 video-equipped central offices, each of which serves an average of 10,000 homes. The video service take rate is projected to be 40 percent of homes passed, or about 4,000 homes per central office. Each home is expected to have an average of 2.5 TV sets, which can be served independently.

The IPTV/Video over Broadcast Solution uses an advanced video codec and offers 150 channels of standard-definition content, requiring approximately 300 Mbit/s, assuming 2 Mbit/s or less for each channel. Add another 50 HDTV channels, each consuming 8 to 10 Mbit/s, and the total bandwidth required in the network for the broadcast service is still just under 1 Gbit/s to any given video-equipped central office where final multicast leaf replication occurs.

The peak concurrent demand for VOD is forecast to be 20 percent of all set-top boxes, which adds up to a peak 2,000 streams from the servers to a given central office. If 90 percent of those VOD streams are standard-definition video at 2 Mbit/s and 10 percent at 8 Mbit/s each, then the total is a little more than 5 Gbit/s per central office from the on-demand servers at peak.

Note that VOD, as well as nPVR, require much more bandwidth than any of the other services, and quickly become the driving factors in network design.

However, another Cisco innovation, *PIM SSM Mapping*, enables the aggregation network to map IGMPv3 messages into PIM SSM messages that can be read by the network as the tree is established. PIM SSM requires the operator to configure multicast source addresses in the network routers, so only authorized video sources are allowed onto the network. This helps protect against DoS attacks and from others trying to broadcast unauthorized content to subscribers.

These technologies rely on routing, Gigabit Ethernet, and other capabilities embedded in Cisco IOS Software and equipment commonly used in digital video networks: the Cisco CRS-1 Carrier Routing System; Cisco 7600 Series Router; Catalyst 6500, 4500, and

FURTHER READING

- White paper: *Video over Broadband: Taking Video to the Next Scale*
cisco.com/packet/173_8a1
- Gigabit Ethernet-Optimized VOD Solution
cisco.com/packet/173_8a2

3750 Series Gigabit Ethernet switches; Cisco ONS 15454 Reconfigurable Optical Add/Drop Multiplexer (ROADM); Cisco dense wavelength-division multiplexing (DWDM) pluggable integrated optics; and new optical filtering capabilities. Also embedded in these devices are provisions for QoS that help ensure customer satisfaction and for security that protect both the network and revenues.

Where to go from here? Gupta sees a number of challenges—and potential opportunities—ahead. One is figuring out a way to handle oversubscription of video services that gracefully lets customers know the service is not available, and that their request will be met a short time later. The network and video service will need to link together to apply admission control to ensure that QoS is maintained for subscribers already watching a video.

Field Experience: neuf telecom

Challenges notwithstanding, service providers worldwide already offer full-blown video service or are taking steps toward doing so. Among them is one of France's leading telecom carriers, neuf telecom, whose territory includes more than half the country's population. Using the Cisco Video Networking Solution along with the Cisco 7600 Series Router,

Catalyst 3750 Series Switch, and Cisco IOS Software, neuf telecom has deployed a multiservice, multimedia network to more than 3,000 cities that began service in 2004.

In November of that year, the telecom carrier launched neuf TV, a digital service delivered over ADSL with more than 70 channels, and also has begun offering ADSL2+.

“The Cisco-based network gives us the ability to introduce our unique ‘a la carte’ services, whereby customers can select from a clearly priced menu of services,” says Francois Paulus, general manager of neuf telecom's Network Division. “By simplifying the provisioning of TV services, we keep our operating expenses at a minimum and maintain our competitiveness.”

According to Jean-Christophe Dessange, broadband business development manager at Cisco, neuf telecom scaled its residential subscriber base from 100,000 ADSL subscribers at the end of 2003 to 400,000 subscribers for ADSL a la carte services by the end of 2004. What's more, the carrier reported that sales in the first half of 2004 increased 24 percent year over year. ■

Ridding Networks of DDoS Attacks

An integrated solution lets providers offer new services to security-conscious customers.

By Edmund Lam

It is easier than you probably think for a distributed denial of service (DDoS) attack to bring down a corporate network. DDoS attacks are usually created by botnets, networks of compromised individual computers (bots) that can be directed by an attacker to launch a flood of various types of packets at the target. One relatively small botnet might comprise 1,000 machines. If you estimate an average upstream bandwidth of 128 Kbit/s per machine, the botnet can generate more than 100 Mbit/s of incoming data, larger than many WAN connections between service providers and large enterprises.

Worse, DDoS attacks on businesses are growing at a furious pace. The number of bots is exploding, largely because of the rise of home PCs with always-on Internet connections that are often poorly secured and available to hackers. Many attacks are serious criminal operations, threatened to take place during major corporate events.

DDoS attacks can saturate the bandwidth available for incoming, outgoing, and intranetwork communications; overwhelm the capacity of targeted routers, servers, and even firewalls, rendering them unavailable for legitimate traffic; prevent access to specific applications or hosts; attack other network resources such as soft switches, core routers, and Domain Name System (DNS) servers; and cause collateral damage to parts of the network not directly attacked. Not surprisingly, a Gartner Group study found that for most enterprises network security had climbed from tenth place in 2003 to the top of the spending list in 2004.

Compounding this growing problem is the fact that DDoS attacks are difficult to detect, because illegitimate packets are not easily distinguishable from legitimate ones. Typical signature pattern matching, performed by intrusion detection systems (IDS), does not work. Some of the more popular techniques for dealing with DDoS attacks, such as blackholing and router filtering, also fall short in terms of mitigation and ensuring business continuity. Strategies such as overprovisioning do not provide adequate protection against larger attacks, and are far too costly to administer for DDoS prevention.

“Clean Pipes”

The answer to DDoS assaults is a comprehensive protection system that detects attacks when they begin,

diverts and “cleans” the contaminated traffic stream, or pipe, and then returns the legitimate traffic to the network. Cisco now offers such a system called the *Cisco DDoS Protection Solution* that performs all of these tasks (see figure, page 76). It can be deployed by service providers to:

- Manage network DDoS protection for their downstream customers, both over the last mile and within the public-facing infrastructure of their networks
- Protect Web services and e-commerce applications at managed hosting data centers
- Protect links to downstream ISPs from being saturated with DDoS traffic

In addition, the solution enables service providers to defend their own critical network infrastructure elements, such as aggregation, edge, and core routers and switches, DNS servers, and transoceanic links, nearly the instant an attack is detected.

Cisco DDoS Protection offers a comprehensive solution for delivering “clean pipes” capabilities, but service providers are strongly recommended to also implement security measures known as Cisco Network Foundation Protection (see sidebar, page 77). These measures harden the data, control, and management planes of a provider’s infrastructure against security threats and enables better positioning for service delivery. Meanwhile, the functional elements of the Cisco DDoS Protection Solution—*detection* and *mitigation* (includes diversion, cleaning, and injection)—work in concert to protect service providers and their customers’ networks from DDoS attacks.

Detection

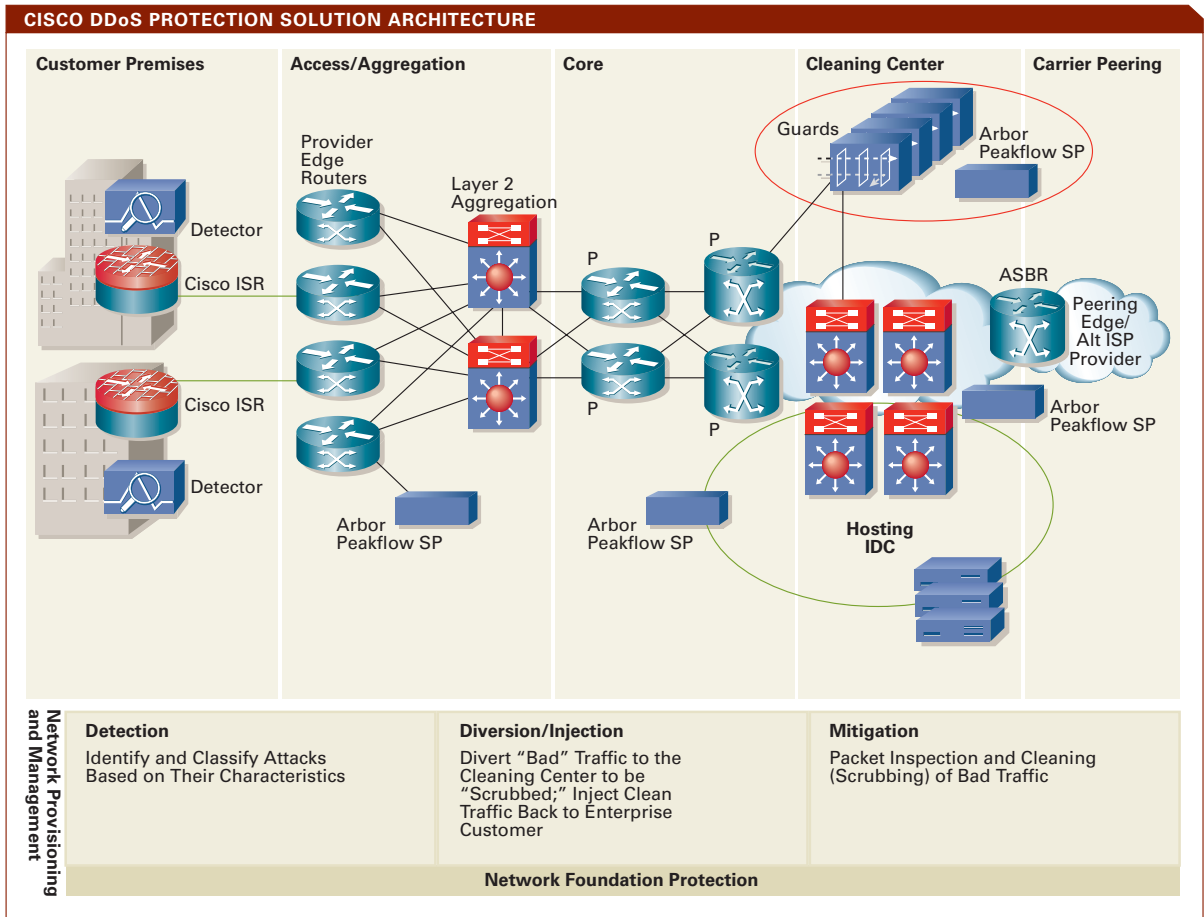
DDoS detection has historically been tough, because illegitimate packets have the same content and header characteristics as legitimate ones. Moreover, many attacks use spoofed source addresses. The defense, therefore, must be predicated on an ability to differentiate between good and bad traffic flows.

Cisco’s DDoS solution employs two systems that dynamically create baseline traffic models and compare flows to them in real time. Any differences above a certain threshold trigger an alarm. The two systems are

For a graphic depiction of the detection and mitigation steps involved in the Cisco DDoS Protection Solution, visit cisco.com/packet/173_8b1.

INTEGRATED PROTECTION

Cisco's DDoS Protection Solution composes detection and mitigation (diversion, cleaning, and injection) functions. Detection takes place at the customer edge and/or within the network, and mitigation near the peering edge.



deployed in different places in the network, depending on the location that best fits the user's needs: the Cisco Traffic Anomaly Detector XT standalone appliance or Arbor Peakflow SP from Arbor Networks, a Cisco Technology Developer Program Partner. Arbor Peakflow SP works in conjunction with NetFlow functionality in Cisco IOS Software to perform detection.

The Cisco Traffic Anomaly Detector XT offers the highest detection capacity because it sits at the customer edge of the network, examining every traffic flow as it is delivered to the customer. It monitors mirrored traffic from the wire. Traffic is copied by a port-based feature such as Switched Port Analyzer (SPAN), virtual LAN (VLAN), or VPN Access Control List (VACL), or by optical splitting, and one stream is fed into the detection device.

The Cisco Traffic Anomaly Detector XT uses the latest behavioral analyses and attack recognition technology



EDMUND LAM, a technical leader in the Cisco's Service Provider Systems Engineering Group, focuses on managed security system design. Prior to his current role, he was a technical marketing engineer specializing in system design for VoIP and MPLS VPN for service providers.

to proactively detect and identify all types of DDoS attacks. If the detector senses an anomalous flow, such as a huge burst of traffic to a specific address or a burst of a certain type of traffic, it immediately and dynamically generates an event in syslog, or it activates a remote Cisco Guard over a secure Secure Shell (SSH) connection.

The Arbor Peakflow SP device resides in an out-of-band network, receiving Cisco NetFlow statistics collected from various routers in the service provider's network. Cisco NetFlow is the most widely deployed DDoS identification and network traffic flow analysis technology for IP networks available today. It classifies packets by looking at the seven-tuple in the header—the information on ingress interface, IP protocol type, type-of-service indicator, source and destination IP addresses, and source and destination ports. This information describes a baseline profile of normal traffic across the entire network, enabling network managers to see abnormal flows as they develop anywhere. Its data collection does not affect network performance or reliability, and the data travel out of band to the Arbor Peakflow SP device, so coverage can be easily scaled in size.

Using the NetFlow data, the Arbor Peakflow SP device identifies anomalies using both signature analysis and

Cisco Network Foundation Protection

In today's competitive business climate, connecting to the Internet is imperative; however, this also exposes network elements and infrastructure to myriad risks and threats. To address the increasing complexity of attacks in this heightened security environment, Cisco has enhanced Cisco IOS Software features and capabilities for network elements as well as the infrastructure, helping to ensure their availability under any circumstances.

Cisco Network Foundation Protection (NFP) provides the tools, technologies, and services that enable organizations to secure their network foundation. This, in turn, enables controlling packet flows and protecting a service provider's network core against security threats such as DDoS.

A secure infrastructure also forms the foundation for service delivery. Continuous service delivery requires a methodical approach to protecting router planes. The router is typically segmented into three planes, each with a clearly identified objective. The *data plane* allows the ability to forward data packets; the *control plane* allows the ability to route data correctly; and the *management plane* allows the ability to manage network elements.

In securing the foundation, Cisco recommends that service providers take the "security toolkit" approach—selecting security tools and techniques based on assessing and identifying risks and threats to the network infrastructure. The security toolkit should also be flexible enough so that new tools and techniques can be

added when a reactionary approach is warranted to defend against a security threat.

With careful consideration to meet the objectives for each router plane, service providers can select the right tool for the right job when dealing with security incidents.

Data plane protection allows for detecting traffic anomalies and responding to attacks in real time. Some of the tools associated with securing the data plane are NetFlow, IP Source Tracker, access control list (ACL), Unicast Reverse Path Forwarding (uRPF), Remotely Triggered Blackhole (RTBH) Filtering, and QoS tools.

Control plane calls for defense-in-depth protection for routing control. Some of the tools for securing the control plane are Receive ACL (rACL) and Control Plane Policing (CoPP).

Management plane protection allows secure, continuous management of Cisco IOS network infrastructure. Among the tools for securing the management plane are CPU and memory thresholding and dual export syslog.

When it comes to securing the network foundation, Cisco NFP should be considered a proactive security measure. In addition, methodical segmentation of router planes combined with the security toolkit approach will go a long way in providing flexibility and strengthening the ability to provide tactical help with security issues. To learn more about Cisco NFP, visit cisco.com/go/nfp.

dynamic profiling, the two most effective methods deployed. The profiles, which are continually rebuilt as traffic patterns change over time, incorporate temporal and topological components to create sophisticated models of network behavior. If an anomaly exceeds user-defined thresholds for severity and duration, the device alerts network staff, who can choose to activate the Cisco Guard to stop the attack.

Mitigation: Diversion, Cleaning, and Injection

The Cisco DDoS Protection Solution mitigation function aims to accurately distinguish legitimate traffic from malicious traffic destined for mission-critical hosts (e.g., DNS servers, Web servers, and voice over IP softswitches), drop malicious traffic, and allow legitimate traffic to pass. Mitigation is accomplished by either a free-standing Cisco Guard XT appliance or the Cisco Anomaly Guard Service Module, which

resides in Cisco Catalyst 6500 Series switches or Cisco 7600 Series routers.

Upon receiving a request from the Cisco Traffic Anomaly Detector XT, Arbor Peakflow SP device, or via manual activation by network operations staff, the Cisco Guard XT sends the upstream router a Border Gateway Protocol (BGP) message to make it the next hop to the targeted zone. The upstream router diverts the dirty traffic (legitimate and malicious) destined for the targeted zone to the Cisco Guard for cleaning. The *zone* is a network element protected by the Cisco Guard against DDoS attacks; it can be a network server, client, or router; network link or subnet or an entire network; an individual Internet user or company doing business using the Internet; an ISP; or any combination or variant of these.

Continued on page 79

DDoS Protection, Continued from page 77

Once switched into zone protection, the Guard begins receiving the diverted traffic and applying its policies. The policies refer to and instruct the Guard's protecting system to carry out an action, which could range from notifying the user of suspicious traffic to directing the traffic to various Guard anti-spoofing or anti-zombie mechanisms to dropping the traffic. The zombie policies react when the measured traffic violates their threshold.

To perform its protective role, the Guard has an array of filters with different characteristics that are sensitively tuned. These filters enable the Guard and the user to filter out suspicious and malicious traffic and enable legitimate traffic to pass the zone. The Guard also has protection modules that follow up on traffic cleansing with their anti-spoofing mechanisms. The *analysis module* allows the traffic during protection to flow monitored but unhindered as long as no abnormalities are traced; the *basic module* has anti-spoofing and anti-zombie mechanisms that authenticate traffic; the *strong module* has more severe anti-spoofing mechanisms; the *drop module* drops malicious traffic; the *rate limiting module* limits the rate of a desired traffic flow or an overall zone traffic; and the *recognition module*, which coordinates between the Cisco Guard policies and the filter system, and samples the outgoing traffic for analysis. The Guard filter system and protection modules operate in a cyclic mode over the zone traffic flow.

After the traffic destined for the attacked zone has been scrubbed, the Cisco Guard transparently allows the zone traffic to be injected back to the zone. This state continues until the user ends the zone's protection. The method of injection depends on whether the network uses a Layer 2 or Layer 3 core topology. These methods, such as Policy Based Routing (PBR), Virtual Routing/Forwarding (VRF), and Generic Routing Encapsulation (GRE), are configured on the immediate downstream router and the Cisco Guard or on only the Guard. They ensure that the clean traffic doesn't loop back to the Cisco Guard. The Cisco Guard's mitigation process ends when the configurable timeout passes since the last dynamic filter created for dropping DDoS traffic is removed. Then, the traffic flows destined for the previously attacked zone will be no longer diverted to the Guard and resume their normal data paths.

To initialize a zone for DDoS mitigation, the Cisco Guard XT, like the Traffic Anomaly Detector XT and Arbor Peakflow SP, needs to be put in "learning" mode during peace time. In this mode, the device passively monitors traffic patterns and their rates and thresholds destined for resources within the zone to understand normal behavior and establish a baseline on which to compare zone traffic and trace abnormalities that might, in turn, become malicious when a DDoS attack occurs.

Deployment Models

Cisco's goal with its DDoS Protection Solution is to enable service providers to integrate the devices and their capabilities in the way that most cost effectively benefits the provider and its customers. There are three service deployment models:

- *Managed Network DDoS Protection*—Providers offer their customers protection against DDoS attacks on their last-mile connections and infrastructure networks. This model delivers an additional layer of business continuance assurance.
- *Managed Hosting DDoS Protection*—Hosting providers help protect critical managed Web and application servers from DDoS attacks. Detection is provided closest to the assets under attack.
- *Managed Peering Point DDoS Protection*—Enables providers to supply DDoS-free wholesale connections to their downstream ISP customers, maximizing bandwidth for legitimate traffic.

Another deployment option, *Infrastructure DDoS Protection* enables providers to defend their own network infrastructure from DDoS attacks. This model reduces directed attacks on vital places in the network, and protects critical servers in the provider's data center, including DNS, HTTP, and Simple Mail Transfer Protocol (SMTP) servers.

In addition to creating new revenue streams, the managed services deployment models benefit providers on several fronts. Among them: enhances the customer's view of the provider as a trusted partner who understands the security needs of its business; leverages existing network infrastructure capabilities (i.e., core assets) to offer the service with minimal capital expenditure (CapEx) investment.

Benefits to customers include proactive, effective, real-time mitigation of attacks before the last mile and data center resources are overwhelmed; avoidance of costly last-mile bandwidth upgrade due to traffic congestion or loss caused by DDoS attacks; and improved network uptime for greater business continuity and enhanced customer experience. ■

FURTHER READING

- Cisco DDoS Protection Solution
cisco.com/go/cleanpipes
- Cisco DDoS Protection Solution White Paper
cisco.com/packet/173_8b2
- Cisco Anomaly Detection and Mitigation Portfolio
cisco.com/packet/173_8b3
- Arbor Networks
arbor.net

Smart Connections

Red Hat has room to grow with Cisco Catalyst 4500 Series Switch.

By David Baum

In an industry rife with success stories, Red Hat, Inc. (redhat.com) is one of the most striking. What began as a better way to build software—based on openness, transparency, and collaboration—quickly shifted the balance of power in the entire computer industry. Today, the open source model pioneered by Red Hat is influencing how millions of organizations buy, sell, and use software, and Red Hat's Linux brand is now the most recognized open source operating system in the world.

"We've seen explosive growth for about five years now," says Stacy J. Brandenburg, manager of network operations at Red Hat, which is headquartered in Raleigh, North Carolina.

Since its founding in 1993, Brandenburg has watched Red Hat grow from 350 employees to more than 1,000 employees, and from five offices to 27 offices worldwide. Red Hat has also added two major engineering centers and several data centers for both internal and external usage. Because Red Hat runs a very lean IT organization, with a network operations staff of just three people, Brandenburg and his team insist on purchasing stable and efficient networking gear that is easy to upgrade.

"We need flexible network devices that can scale," Brandenburg emphasizes. "With the Cisco Catalyst 4500 Series, we have a consolidated switching infrastructure with a single point of management. It's all in one box, with one interface, and thus is very easy to deploy and maintain."

Flexible Solution

Currently, Red Hat has 14 Cisco Catalyst 4500 Series switches—primarily used to manage core routing functions at its remote offices. In some cases, the switches are connected to ancillary Cisco Catalyst 2900 Series and Cisco Catalyst 3500 Series switches to increase port density. "The Cisco Catalyst 4500 Series handles all of our internal routing needs in these locations," confirms Brandenburg. "We evaluated equipment from other vendors but we did not find anything comparable that we were comfortable with."

Formerly, Red Hat used Cisco Catalyst 2900 Series Switches as edge switches in its branch offices. While these network devices worked well, managing them became unwieldy—especially once Red Hat had 30 or 40 switches in each location. "We thought a



Courtesy of Red Hat

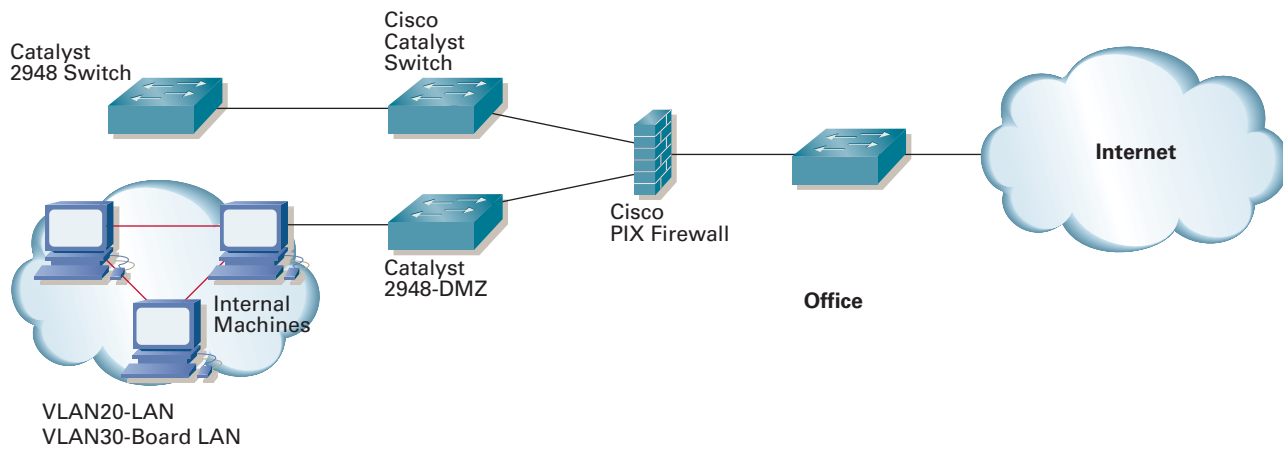
SPEED AND EFFICIENCY Red Hat's consolidated Cisco switching infrastructure has enabled the company to handle explosive growth over the past five years.

chassis-based switch architecture would make more sense for our needs," Brandenburg recalls. "The Cisco Catalyst 4500 Series is ideal in offices where we have between six and 60 people. It delivers the redundancy we need, and it has a complete set of Layer 2 and Layer 3 options."

The Cisco Catalyst 4500 Series supports a unified set of line cards that are operated by a powerful supervisor engine to provide access, aggregation, and small core functions. Each chassis holds the individual components that make up the system, and every element in the chassis is changeable. Some slots are reserved for special functions, such as supervisor engines. Most slots are flexible and can be configured for specific customer needs.

"The benefit of this modular architecture is the ability to add, remove, and change system elements over time," explains Ellen Brigham, a product manager in Cisco's Gigabit Switching Business Unit. "For example, line cards can be added and removed as needed, which ensures longevity and investment protection through reuse of system components."

With a traditional fixed-configuration switch, network administrators cannot change the type of access port. Fixed switches have one power supply and one fan, both of which are single points of

**OFFICE IN A BOX**

Cisco Catalyst 4500 Series switches are quick to deploy and upgrade, providing Red Hat's remote offices with a consolidated infrastructure and a single point of management. Ancillary Cisco Catalyst 2900 Series switches increase port density.

failure. There is no option for hardware performance upgrades, and there are fewer capabilities per port or per user.

"In a modular switching system, by contrast, every element in the chassis can be removed and upgraded," Brigham says. "You have redundant hardware capabilities, and you can increase the performance by upgrading the supervisor engine."

Red Hat's standard configuration includes redundant supervisor engines to help minimize network downtime. Thanks to stateful switchover technology, the secondary supervisor engine can immediately take over if the primary supervisor fails. During the switchover, Layer 2 links are maintained transparently, without the need to renegotiate sessions. Brandenburg is particularly impressed with the Cisco Catalyst 4000/4500 Supervisor Engine IV.

"The new supervisor engine extends all the functionality we need, some of which was not available in our former switching solution," he says.

Office in a Box

According to Brigham, the Supervisor Engine IV includes advanced Layer 3 switching capabilities and gigabit wire-speed performance, enhancing its suitability as a backbone device in branch-office networks.

"These capabilities make the Cisco Catalyst 4500 Series an ideal branch-office solution capable of meeting the needs of both small and large operations, as well as small enterprise applications," she notes.

For example, as Red Hat opens new branch offices, the Cisco Catalyst 4500 is used as an "office-in-a-box" switch. Network engineers set up each chassis with multiple power supplies and redundant supervisor engines, then select line cards appropriate to the needs of the office at hand. After that, configuring the switch is a simple matter.

"We just set the IP name, deal with a couple of host issues, and upload our standard configuration to the switch," says Brandenburg. "We can unpack a switch and prepare it for deployment in 30 or 40 minutes."

Adding a new line card is even simpler. "We just pop it in and configure it, without taking the switch off the network," Brandenburg adds. "We handle all implementation and upgrade activities in-house. The deployment process is straightforward. As long as the port density is not a major factor, the 4500 is a great choice for SMBs [small and midsized businesses]."

Brigham agrees, particularly as small to midsized firms adopt converged networks supporting IP telephony. "Converging voice, data, and video traffic on a single network lowers the overall cost of ownership and simplifies administration—assuming you have a switching infrastructure that can distinguish each traffic type and manage it according to its unique requirements," she points out.

When combined with Cisco IOS Software, the Cisco Catalyst 4500 Series can deliver this type of control. It also includes integrated Power over Ethernet (PoE) in the chassis to send both power and network connectivity to IP devices such as telephones, wireless access points, video cameras, and other appliances—over a single wire (see figure on page 82). "PoE makes it possible to place devices in unique locations without having to provide new outlets or electrical circuits," explains Brigham.



Moreover, PoE allows businesses to isolate critical devices on a single power system, limiting potential points of failure and enabling the entire system to be supported by uninterruptible power supply (UPS) backup.

While Red Hat plans to use these IP telephony capabilities in the future, today the company primarily relies on the Cisco Catalyst 4500 Series switches to support high-density gigabit access at Layer 2 and robust routing and switching at Layer 3. In most cases, the switches handle both access and distribution, with support for up to 30 users, with six to eight ports for each user. As additional users are added, Red Hat deploys fixed configuration switches as well, and the 4500 becomes the distribution/core switch.

Gigabit to the Desktop

Because Red Hat develops operating system software, users tend to transfer very large files around the network. “Gigabit to the desktop is a must for everyone,” says Brandenburg. “When our developers are working on an OS kernel, they routinely exchange three or four gigabytes of data. Our productivity is curtailed if software engineers have to wait on these files. Having a very fast and efficient network is important to us.”

The Cisco Catalyst 4500 Series easily extends gigabit bandwidth to the desktop. It’s available with 48- and 24-port triple-speed autosensing and autonegotiating 10/100/1000BASE-T line cards, supporting up to 384 Fast Ethernet or Gigabit Ethernet ports in a chassis—either fibre or copper.

At a time when global collaboration and distributed work teams are commonplace, more companies depend on these types of low latency, high throughput connections. Ad agencies, architectural firms, design shops, radiologists, and a variety of engineering companies routinely transmit large files across the network. High-density switches such as the Cisco Catalyst 4500 Series allow these companies to deploy Gigabit Ethernet over their existing infrastructures of Category 5 cable, eliminating the need to rewire the LAN.

“Users implementing Gigabit Ethernet to the desktop realize significant improvements in performance and productivity,” Brigham notes.

Integrated Management

Efficiency and self-reliance are hallmarks of the Red Hat culture. Its employees take pride in being “the biggest small company in the world.” Red Hat stands among giants, competing and partnering with companies many times its size. Its continued success is proof of the power of the open source model—and also of the resourcefulness of Red Hat’s staff.

This attitude certainly pervades the IT staff. Brandenburg and his colleagues are in a constant state of alert as they support the bandwidth needs of their worldwide organization.

“In the past, many of our sites did not have the IP functionality that we needed within the Layer 3 cards of the 4000 series,” he admits. “We had to install a second router, which meant two points of management for IP connectivity. We also had a problem with line cards requiring another set of switches.”

Today, the Cisco Catalyst 4500 Series presents a single point of management for multiple networking functions.

Convenient Upgrades

When it comes to upgrades, higher-layer functional enhancements are possible on all system ports without replacing existing line cards and wiring—unlike conventional switching products, where complete equipment upgrades are typical during migration. Line cards can be added and deleted from the system as needed, and a supervisor engine upgrade immediately enhances all ports on all line cards.

These architectural advantages extend the useful life of Cisco Catalyst 4500 Series line cards, which can be used interchangeably among Cisco Catalyst 4500 Series and Cisco Catalyst 4000 Series switches. “Our cost of ownership with these switches is very inexpensive in the long run,” Brandenburg maintains. “The network typically runs at 80 to 100 percent utilization, yet we’ve seen a very low failure rate.”

Red Hat foresees using 10 Gigabit Ethernet in its branch sites as well as at its central engineering sites, especially as they develop and test new drivers for their Red Hat Linux implementations.

“We will be going to 10 Gigabit Ethernet within a year,” Brandenburg confirms. “These modular switches have been a good investment for us.” ■

FURTHER READING

- Cisco SMB Class Network Foundation Solutions
cisco.com/packet/173_9a1
- Cisco Catalyst 4500 Series Switches
cisco.com/packet/173_9a2

Mobility and Pervasive Networks

How can networks provide pervasive connectivity?

By Russ White

It's customary to think of today's wireless networks as untethered—of offering a degree of freedom and flexibility unthought of 20, or even 10 years ago. The next age of networking, however, will be geared toward the second kind of tethers that hold today's networks together: services. In a pervasive network, devices are always attached to a network that is always on, or, perhaps a more accurate description would be that *devices form networks as they are capable of communicating*; services become a part of the network itself.

But how do you attain pervasive networks? Two components must be built before you can truly have pervasive networking:

- Pervasive applications: Whenever users connect to “the network,” they should be able to use services (applications) based on reachability and policy.
- Pervasive connectivity: Users need to be able to connect to “the network” from anyplace, anytime.

Today's applications are generally server-centric; services are run on “well known” servers, and you must be connected to a network with a well known server attached to use specific services. For example, to connect to another user through instant messaging, you must be able to connect to an instant messaging server as a “meeting place,” and for a voice call you must be able to connect to a call manager for voice services. The advent of point-to-point networks illustrates how valuable a more decentralized approach could be in the real world, where services are atomized, and as the atoms collect, they create a larger and larger service based on available connectivity. These peer-to-peer service models provide us with a glimpse of what “the network” could be if pervasive connectivity were combined with atomized services.

Pervasive connectivity is the second piece of the pervasive network puzzle: the piece that falls squarely within the networking world. How can networks provide pervasive connectivity? If we define pervasive connectivity as the ability to connect any device into any collection of devices at any time, pervasive connectivity sounds very similar to device mobility—and, in fact, mobility is one of the keys to pervasive networking.



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 Cisco's *IP Journal*, and can be reached at riw@cisco.com.

Several mobility technologies are being developed and deployed at the moment, including:

- Mobile IP and Network Mobility (NEMO), both of which assume a fixed infrastructure to which services are attached and mobile devices will always have connectivity. Mobile IP and NEMO allow a single device to transparently reach back to services available within “the network” by allowing a device to register with a home agent, which in turn redirects traffic to the actual location of the device at that moment. Mobile IP and NEMO are standardized through the efforts of the Internet Engineering Task Force's (IETF) Mobile IP Working Group; Cisco products support Mobile IP and Mobile Router features in many different configurations.
- Mesh networks, which provide Layer 2 switched networks on an ad-hoc basis, normally through radios of some type. Mesh networks are typically concerned with backhauling connectivity into infrastructure networks with greater simplicity than routed networks. Mesh standards are being developed in the IEEE, the ITU, and other standards bodies. Many companies also provide proprietary solutions for building Mesh networks, and some large-scale deployments of Mesh networks are already occurring.
- Mobile Ad-Hoc Networks (MANET), which are Layer 3 routed networks based on random connectivity patterns. Routing protocols for MANETs are being standardized through the efforts of the IETF, specifically in the MANET and Open Shortest Protocol First (OSPF) working groups.

Each of these three technologies has advantages and disadvantages under specific network conditions. Mobile IP is the most mature of these, with widespread deployment in many environments. Mesh and MANETs are the subject of a great deal of current research, prototypes, and testing. Large-scale Mesh and MANET networks are projected for deployment in the next several years.

In the end, a single device that can connect to any type of mobility network, Mesh, MANET, Mobile IP, or NEMO, will be the ideal client to provide pervasive networks when combined with a network that supports all of these technologies.


Cisco is at the forefront of pervasive IP version 4 (IPv4) and IPv6 connectivity, and is active in the standards organizations and technical teams designing Mobile IP, NEMO, and MANET and Mesh solutions. Mobile IP in Cisco products is already deployed in many environments, including cellular telephones, transportation, and public safety. In the coming years, Cisco will remain a leader in this area, developing and deploying solutions in each area, with a final goal of an always-on, always connected, pervasive network. ■

SPOTLIGHT ON:

Cisco Shared Port Adapters and Cisco SPA Interface Processors

Cisco now offers a portfolio of interface flexibility (I-Flex) options that are interchangeable across Cisco carrier-class routing platforms. Cisco I-Flex combines shared port adapters (SPAs) and SPA interface processors (SIPs)

in a design that can be used to prioritize voice, video, and data services for intelligent, flexible, secure networking.



The fully compatible SPAs and SIPs offer consistent feature support and accelerated service delivery for the Cisco 7304, Cisco 7600 Series, Cisco XR 12000 Series, and Cisco 12000 Series routers as well as the Cisco CRS-1 Carrier Routing System.

The modular architecture allows customers to mix and match interface types from copper T1/E1 to optical 10 Gbit/s and densities of up to ten optical ports per SPA and up to six SPAs per slot.

A single interface can support multiple service types including premium broadband services, IP VPNs, and IP Security (IPSec). Network convergence is supported by interfaces in the SPAs and SIPs for ATM, Frame Relay, and packet over SONET.

For service providers and large enterprises, the Cisco I-Flex design offers line-rate performance while maximizing connectivity options, providing a rich set of quality of service (QoS) features, protecting investments, and reducing the overall cost of network ownership.

cisco.com/go/spa

Core and Multiservice Edge Routing

Cisco XR 12000 Series Router

The new Cisco XR 12000 Series routers scale from 2.5 Gbit/s to nX 10 Gbit/s per slot to help scale services in next-generation IP/Multiprotocol Label Switching (IP/MPLS) networks. These routers are powered by the Cisco IOS XR Software, a fully modular, distributed self-healing operating system for carrier-class routers. The Cisco XR 12000 Series also supports secure virtualization, continuous system operation, as well as scalability and distributed processing for delivering multiple services on a single platform. Router models offer the choice of a 16-, 10-, 6-, or 4-slot chassis for deployment in a service provider's point of presence. For more on the Cisco XR 12000 Series Router and IOS XR Software, see page 43.

cisco.com/go/12000

Edge Routing, Access, and Aggregation

Cisco 7604 Router

The Cisco 7604 Router provides a 4-slot, modular chassis for high availability, density, and nX 10 Gigabit Ethernet performance. This new router can be used for enterprise WAN aggregation or by service providers as an edge device to deliver multiple services and security features between IP and MPLS networks. Two configuration options are available: a single supervisor engine and up to three line cards, or dual supervisor engines and up to two line cards for high availability and redundancy. The compact router also accommodates line cards from DS0 to OC-48 as well as 10/100/1000 Ethernet routing speeds.

cisco.com/go/7600

Cisco uBR10012 Universal Broadband Router: New Half-Height Cards

The Cisco uBR10012 Universal Broadband Router helps to meet needs for Gigabit Ethernet deployment with new support for half-height cards. The Cisco Gigabit Ethernet Half-Height Line Card offers an IEEE 802.3z-compliant Ethernet interface that can run up to 1 Gbit/s in full-duplex mode. The card uses a Small Form-Factor Pluggable (SFP) Gigabit Interface Converter (GBIC) module that supports the SX, LX/LH, and ZX interface types for Gigabit Ethernet. The Cisco Half-Height Line Card Carrier holds up to two half-height cards in a single Cisco uBR10012 router slot, which increases interface density and modularity in a cable headend. cisco.com/packet/173_npd2


Cisco 2600/2800/3700/3800 Series Routers: Application-Oriented Networking Module

The Cisco 2600/2800/3700/3800 Series Application-Oriented Networking (AON) Module brings Cisco AON advantages to enterprise branch offices. With a single processor, the module occupies a single slot in Cisco 2600XM and Cisco 2691 routers as well as all Cisco 2800 Series, Cisco 3700 Series, and Cisco 3800 Series routers. Multiple Cisco AON modules may be deployed in the same chassis and managed as a single virtual cluster. Cisco AON is covered in greater detail on page 59. cisco.com/go/aon

Switching Cisco Catalyst 6500 Series Switch: Application-Oriented Networking Module

The Cisco Catalyst 6500 Series AON Module can be installed in any Cisco Catalyst 6500 Series Switch to take full advantage of the switch's high availability, security, and traffic management capabilities. The module's dual-processor architecture supports hardware acceleration for Extensible Markup Language (XML) and cryptography operations. Occupying a single slot in the switch chassis, the Cisco Catalyst 6500 AON Module is compatible with both fabric and data-bus architectures. Multiple modules may be deployed in the same chassis and managed as a single virtual node. Cisco AON is covered in greater detail on page 59. cisco.com/go/aon

Security and VPNs Cisco WebVPN Services Module

The Cisco WebVPN Services Module offers high-speed Secure Sockets Layer (SSL) virtual private network (VPN) processing for Cisco Catalyst 6500 Series switches and Cisco 7600 Series routers. Supporting up to 32,000 SSL VPN users and 128,000 connections per chassis, the Cisco WebVPN Services Module can cost effectively meet the capacity requirements of service providers and large enterprises. For security control, the module also offers unique virtualization capabilities that simplify policy creation and enforcement for diverse user communities. The module also supports the Cisco Secure Desktop (September 2005 release), which assesses device security before connecting to the VPN. cisco.com/packet/173_npd3 



Cisco Traffic Anomaly Detector Services Module and Cisco Anomaly Guard Services Module


The Cisco Traffic Anomaly Detector Services Module, available for Cisco Catalyst 6500 Series switches and Cisco 7600 Series routers, quickly detects network attacks and automatically activates the Cisco Anomaly Guard Services Modules to initiate mitigation services. Designed to process attack traffic at gigabit rates, the Cisco Anomaly Guard Services Module includes an anomaly recognition engine that identifies and mitigates distributed denial of service (DDoS) and other attacks. The combined solution provides a scalable, flexible, cost-effective method to help preserve business integrity and operations. These new modules are covered in greater detail on page 75.

Cisco Traffic Anomaly Detector Services Module: cisco.com/packet/173_npd4
Cisco Anomaly Guard Services Module: cisco.com/packet/173_npd5

Cisco Clean Access Out-of-Band Solution

The Cisco Clean Access Out-of-Band software solution works with selected Cisco Catalyst switching models to automatically detect, isolate, and clean infected or vulnerable devices that attempt network access. Designed to increase protection for enterprise networks, the out-of-band features make it easier to enforce security policies and deploy access-control capabilities across a WAN. Cisco Clean Access is covered in greater detail on page 88. cisco.com/packet/173_npd6

Voice and Video Cisco AS5400XM and Cisco AS5350XM Universal Gateways

The new Cisco AS5400XM and Cisco AS5350XM universal gateways offer unparalleled capacity and the flexibility to deploy advanced voice, fax, and remote-access services with double the performance of the current generation AS5400HPX. These gateways provide simultaneous support for voice, fax, and remote-access services using multiple protocols including Session Initiation Protocol (SIP), H.323, and Media Gateway Control Protocol (MGCP). They are an important component in many solutions including voice over broadband; Cisco Customer Voice Portal (CVP 3.0) for IP-enabled call centers; unified communications; and Cisco Voice Infrastructure and Applications (VIA) architecture. Both models offer dual, autosensing 10/100/1000 Gigabit Ethernet interfaces and 512 MB memory. The Cisco AS5400XM supports PSTN connectivity up to a CT3, and the Cisco AS5350XM offers density up to eight PRI. cisco.com/go/iad 



Wireless

Cisco 2700 Series Wireless Location Appliance

For enhanced network visibility and control, the Cisco 2700 Series Wireless Location Appliance simultaneously tracks thousands of IEEE 802.11 wireless devices within an enterprise wireless LAN. Powerful device tracking to within a few meters is supported by Cisco's patent-pending RF fingerprinting technology that locates people or resources; provides asset tracking and inventory management; helps enable workflow automation; and delivers source identification for wireless security threats. This appliance works with the Cisco Wireless Control System (WCS) and uses Cisco Aironet lightweight access points. For more on the Cisco 2700 Series Wireless Location Appliance, see page 68. cisco.com/packet/173_npd7

CiscoWorks Wireless LAN Solution Engine Express

The CiscoWorks Wireless LAN Solution Engine (WLSE) Express software combines a wireless LAN management solution and a server for user authentication, authorization, and accounting (AAA). This solution is suitable for managing wireless LANs based on the Cisco Distributed WLAN Solution with Cisco Aironet access points in small and midsize businesses or enterprise branch offices. CiscoWorks WLSE Express can manage up to 100 Cisco Aironet access points while the AAA services provide reliable and localized user access control.

cisco.com/packet/173_npd8 



Network Management

Cisco AON Development Studio and Cisco AON Management Console

Cisco AON Development Studio (ADS) is a Windows-based tool that helps developers configure how application messages are handled at runtime. Cisco ADS includes a set of preconfigured functions (called bladelets) for creating message

plans, one-button synchronization of message plans with the Cisco AON Management Console, and separate additional kits for creating custom bladelets and adapters. Cisco AON Management Console (AMC) is a Linux-based, Web application that provides role-based access control features for centralized management of the Cisco AON system. Cisco AMC helps administrators maintain consistent, up-to-date configurations across all Cisco AON devices through numerous functions for monitoring and management. For more on Cisco AON, see page 59.

cisco.com/go/aon

Cisco Network Planning Solution

Cisco Network Planning Solution (NPS) software provides a "what-if" impact analysis tool that helps users to optimize IP network capacity and performance, improve resilience, plan for new technology deployments, improve application and service continuity, and validate planned configuration changes. Cisco NPS reduces the risks associated with network growth, migration, and consolidation by supporting numerous planning and operational decision studies. Study data is drawn automatically from Cisco routers and switches, as well as selected Cisco applications for network management. The software serves network staff in midsize and large enterprises.

cisco.com/packet/173_npd9

Cisco Configuration Assurance Solution

Cisco Configuration Assurance Solution (CAS) software automatically diagnoses device misconfigurations, policy violations, inefficiencies, and security gaps using a combined visualization of network topology, traffic, and configuration. Cisco CAS improves network availability and security, and documents compliance with best practices and important regulatory and IT governance requirements. Cisco CAS offers the unique ability to analyze the entire IP network configuration from individual devices to network-wide operations, and to predict the network's ability to maintain integrity and security under failure conditions.

cisco.com/packet/173_npd10

Cisco Application Analysis Solution

Cisco Application Analysis Solution (AAS) software provides a visual, quantitative

ABOUT NEW PRODUCT DISPATCHES

Keeping up with Cisco's myriad new products can be a challenge. To help readers stay informed, *Packet* magazine's "New Product Dispatches" provide snapshots of the latest products released by Cisco between May and July 2005. For real-time announcements of the most recently released products, see "News Archive, News Releases by Date" at newsroom.cisco.com/dlls/index.shtml.

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

breakdown of the complex interactions among applications, servers, and networks to troubleshoot and reliably deploy networked applications. Cisco AAS identifies whether performance problems are due to network or application issues. Cisco AAS helps network and application staff to pinpoint bottlenecks, diagnose performance problems, explore the impact of proposed changes, and predict application performance under varying configurations and network conditions.

cisco.com/packet/173_npd11

Cisco Network Assistant Version 2.0

The Cisco Network Assistant software offers centralized management and configuration capabilities for routers, switches, and access points in networks that connect up to 250 users. New features in version 2.0 include: a Smartports Advisor tool that discovers Cisco devices and applies Cisco-recommended configurations, a drag-and-drop feature for installing Cisco IOS Software upgrades, the ability to apply common services across network devices by creating a community, and enhanced topology views.

cisco.com/go/cna

Using Cisco Clean Access for Network Admission Control

The Cisco Networking Professionals Connection is an online gathering place for Cisco experts and networking colleagues. Following are excerpts from a recent Ask the Expert forum, "Using Cisco Clean Access for Network Admission Control," moderated by Cisco's Nick Chong. To view the full discussion, visit cisco.com/packet/173_10a1. To join other live online discussions, visit cisco.com/discuss/networking.

Q: *What method is used by the out-of-band solution for controlling network access based on role? What is the direction for non-Cisco or Cisco switches older than the Catalyst 2950 or 3550 series, specifically the XL series that are common in ResNet?*

A: The Cisco Clean Access out-of-band product provides two forms of access assignment for authenticated users: port-based or role-based virtual LAN (VLAN) mapping. The end user to CAS (server connection) is UDP (discovery) and SSL (communication), while the CAM (manager) to switch is based on Simple Network Management Protocol (SNMP) v1/2c/3. The current out-of-band release (starting April 15, 2005) supports the 2950, 3550, 3560, 3750, 4500, and 6500 series. Cisco Clean Access inband (agnostic to switch/router) will continue to support all non-Cisco and older Cisco switches. You can have either or both in your network.

Q: *Is there another way to distribute the Clean Access Agent instead of the Web login? Can the agent be updated automatically?*

A: If you have an internal update server or software deployment tool, download our agent from the Cisco download and upload to your update server. Use your common update method for providing the update to users. Or you can create a "link" button on the actual Web login page (username, password). Then, create a link such as "Click here to download the agent." Users can then download prior to login. With release 3.5.1 (May 20, 2005), Cisco Clean Access introduced easy auto upgrade of Clean Access Agent client software. Whenever a newer version of the agent is available, the Clean Access Agent informs the user and automatically performs the upgrade.

Q: *Can you use Cisco Clean Access without any Cisco network devices? Where do we install Cisco Clean Access in our LAN? Which components are required?*

A: Cisco Clean Access supports networks with or without Cisco switches/routers. For non-Cisco networks, you can deploy Cisco Clean Access Server in inline mode, as either a virtual gateway (Layer 2) or real-IP gateway (Layer 3). You can also deploy the server at your edge network or central (switch). In most instances, you would see the Cisco Clean Access Server connecting between an edge switch and a core switch, or hanging off a core switch where the edge switches are trunked to using 802.1q. There are three components to Clean Access: manager, server, and the

optional client. Refer to documentation at cisco.com/go/cca.

Q: *Can you exceed the 1,500 user recommended limit? Does Clean Access Manager or Clean Access Server manage the switches? How many switches can a server/manager manage out of band?*

A: The 1,500 user limit is based on our lab projection using 1 GB of memory, a single Pentium processor, and concurrent user logon with posture assessments using the agent. Cisco recommends staying within the guidelines. For out of band, the manager controls the switches. The server provides the authentication, posture assessment, and remediation services. Cisco recommends 20 servers to a manager.

Q: *Does multicast traffic working through current Clean Access Servers work with the Clean Access Server in the virtual gateway, real-IP gateway, and NAT [Network Address Translation] modes? How will the out-of-band deployment work with it?*

A: You have various options to support multicast in either mode. For inband, configure virtual gateway using VLAN mapping. In out-of-band mode, also use virtual gateway with VLAN mapping if you have multicast traffic during the authentication/quarantine role (e.g., some PXE imaging applications used for quarantine remediation). If multicast traffic is post-auth/post posture assessment, you can use real-IP for /30 subnets during auth/quarantine. After that, the Clean Access Server is bypassed, and multicast traffic traverses via the access VLAN.

Q: *Does the Clean Access Agent have to be on the same Layer 2 network as the Clean Access Server? Are VPN users supported?*

A: Effective in release 3.5.3 (July 2005), the Clean Access Server can support Clean Access agents that are on the same Layer 2 network, as well as agents that are multihop IP away via the inband configuration. VPN users are supported in the same release with an additional single sign-on (SSO) feature when using the Cisco VPN Concentrator.

Do you have a question about Cisco Clean Access and Network Admission Control? Ask the NetPro Expert. Send your question to packet-netpro@cisco.com, with the subject line "Cisco Clean Access." ■



NICK CHONG is a technical marketing manager for the Cisco Clean Access product family. He has more than 15 years of technical networking experience, with hands-on support knowledge in IP routing, SNA networking, network intrusion protection systems, and network access control. He can be reached at nchong@cisco.com.

IP Phones, Continued from page 57

Among the most significant cost benefits of VoIP is remote management. To service its PBXs and add new features, the bank previously relied on contracts with in-market providers. In the markets where Bank of America has migrated to Cisco CallManager, the service is now rolled into the data service contract. "To deploy new features, we can field our own centralized team or outsource to a provider; either way we can manage it remotely," says Hinkley.

Alignment Between IT and Lines of Business

Hinkley won the support of the bank's executive team by presenting the technology solution in the context of its cost savings and business value, such as increased customer satisfaction, increased wallet share, and potential for opening new customer markets.

"We made it a point to 'de-technify' what we discussed with executives," Hinkley says. "They don't really care about OC-48 or IP networks. They care about the business value of the solutions."

Hinkley's team and EDS also educated the executive team about the reliability of VoIP.

"When an application doesn't work, many people wrongly assume that it's a network problem, which makes them hesitant to use the network for voice," says Gilligan, of EDS. "We showed the executive team how we could build a network with the availability and resilience to meet or exceed what Bank of America had experienced with a TDM-based phone."

Hinkley tailored his presentations for each line of business. For the investment banking group, for example, responsiveness is an important competitive advantage. Hinkley conveyed the advantages of IP communications by explaining how associates could receive calls directed to their phone number no matter where they were—home, office, or hotel—using a Cisco VPN 831 concentrator and Cisco IP phone. The solution is in use today.

"The major advantage of VoIP is not the phone, but rather that by packetizing voice, we can deliver it anywhere there's an IP presence," Hinkley concludes. "We're educating our lines of business to consider the flexibility that VoIP technology brings and how it can improve business processes." ■

FURTHER READING

- Cisco IP Communications Solutions cisco.com/go/ipc
- Packet article on network security for the financial industry (Second Quarter 2005) cisco.com/packet/173_7a1
- iQ Magazine IP communications issue cisco.com/packet/173_7a2



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
American Power Conversion (APC)	http://promo.apc.com (key code c986x)	4
BellSouth Business	www.bellsouth.com/business/nobrainier	OBC
Boson Software	www.boson.com	A
BroadHop	www.broadhop.com/cisco	34
Cisco Marketplace	www.cisco.com/go/marketplace/packet	6
Cisco Press	www.ciscopress.com	B
Cisco Systems E-Learning	http://cisco2.elementk.com	10
eiQ Networks	www.eiqnetworks.com/cisco	74
Empirix	www.empirix.com/cisco	78
GL Communications	www.gl.com	64
Ipcelerate	www.ipcelerate.com	16 / 66
NetScout	www.netscout.com/ad/cii	54
Network General	www.networkgeneral.com/Cisco1	58
New Edge Networks	www.newedgenetworks.com/products	42
OPNET Technologies	www.opnet.com	70
Panduit	www.panduit.com/dp16	IBC
SBC	www.sbc.com/ipt	80
Solsoft	www.solsoft.com/packet	8
Statseeker	www.statseeker.com	18
SurfControl	www.surfcontrol.com/go/threatshield	F
Trend Micro	www.trendmicro.com/cisco	60 / 61
Websense	www.cdw.com/websense	24

CACHE FILE

Snippets of Wisdom from Out on the Net

CYBER QUOTE

**“Drive thy business
or it will drive thee.”**

—Benjamin Franklin

High-Speed Web Access Migration Continues Momentum in US

The number of US users migrating from dialup to broadband access continues to grow, and the momentum is expected to continue through the end of the decade, according to JupiterResearch (jupiterresearch.com). Broadband reached 32 million US households, just below half of residential online accounts in 2004. That figure is expected to grow to 88 million, or 78 percent, by the end of 2010. The cable versus DSL war will also continue, reports Jupiter.

Audio Computer Games on the Rise

Michael Feir is an avid gamer who spent so much time playing games in college that he created his own online publication, *Audyssey Gaming Magazine*. But Feir doesn't play the best-selling games—because he's blind. A growing library of audio computer games has been built especially for blind gamers, using sound instead of visuals to let players know what is going on. Because audio games are limited in visuals, programmers must find creative ways to use sound to signify events. Experts estimate that the demand is such that the niche has grown from text-based games coded by hobbyists to between 30 and 50 professional audio game developers who sell 3,000 games a year (cisco.com/packet/173_13a1).

Net Lingo

Evernet—The name given to the always-on, high-speed, broadband, multiformat Web; coined by author and columnist Thomas Friedman (netlingo.com).

Internet Backbone Router Sales Keep Healthy Pace

Sales of high-end (10 Gbit/s) Internet backbone routers worldwide grew at a healthy pace in the first calendar quarter of 2005, according to a report by the Dell'Oro Group (delloro.com). Following a steady growth trend that began in 2003, revenues from backbone router sales reached US\$420 million in the first quarter, representing a quarterly growth rate of 9 percent. “Even though the first quarter of the year is seasonally weak, high-end router demand was exceptional,” notes Shin Umeda, Dell'Oro's principal analyst for routers research. “Internet and telecommunications service providers around the world are investing heavily in their backbone networks to keep up with the onslaught of network traffic coming from broadband users.”

Online Video Poses Rich Opportunity

According to a recent study conducted by the Online Publishers Association (online-publishers.org), online video poses an increasingly rich opportunity for interactive advertisers and publishers. The study, conducted in partnership with Frank N. Magid Associates, surveyed 27,841 Internet users age 13 and older on 25 different publisher Websites. It found 51 percent of respondents watch online video at least once a month, 27 percent watch Internet video at least once a week, and 5 percent watch it on a daily basis. Among the study findings was a strong positive reaction to video ads. Of 70 percent of the respondents who said they had seen a video advertisement online, 44 percent said they had taken some kind of action as a result of seeing that ad.

THE 5TH WAVE



“You should check that box so they can't profile your listening and viewing habits. I didn't, and I'm still getting spam about hearing loss, anger management and psychological counseling.”

©The 5th Wave, www.the5thwave.com