

PACKET

CISCO SYSTEMS USERS MAGAZINE

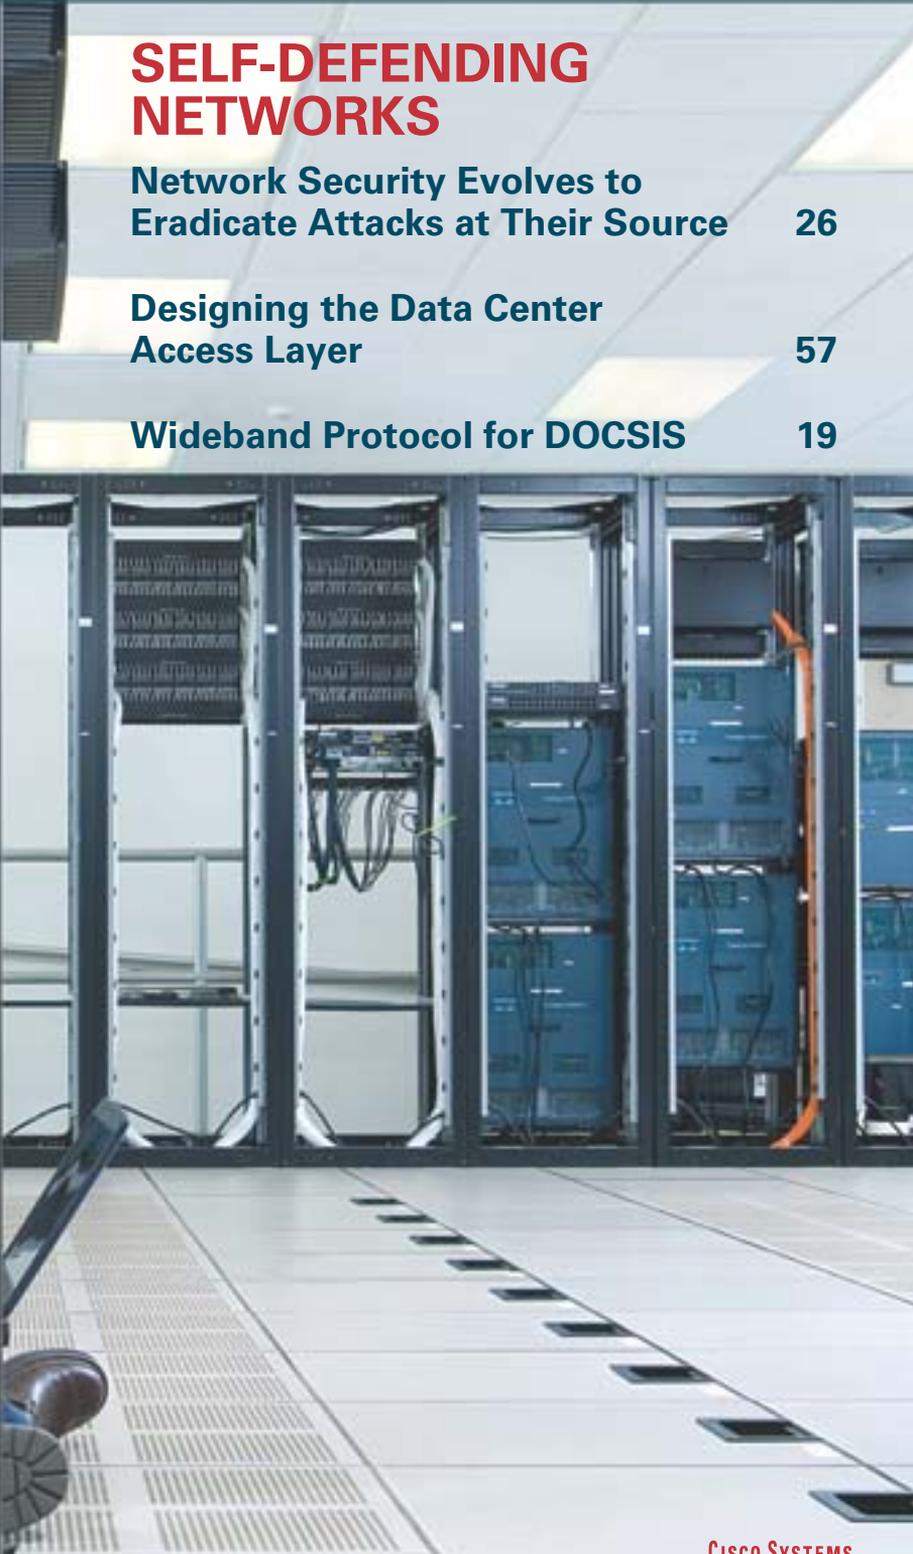
SECOND QUARTER 2005

SELF-DEFENDING NETWORKS

Network Security Evolves to Eradicate Attacks at Their Source 26

Designing the Data Center Access Layer 57

Wideband Protocol for DOCSIS 19



CISCO SYSTEMS

CISCO.COM/PACKET



PACKET

CISCO SYSTEMS USERS MAGAZINE

SECOND QUARTER 2005
VOLUME 17, NO. 2



ON THE COVER

In Self Defense

26

Network security grows adaptive, reaching inside Web applications and excising attacks at their source. New security products from Cisco aim to protect every packet and every packet flow on a network.

Keeping Voice Confidential

34

The key risks of voice over IP joining the network and how to mitigate them.

Eradicating Wireless Intruders

39

Multilayered RF monitoring leaves no room at the wireless LAN table for uninvited guests.

SAN Security: Beyond Zoning

42

Interconnected storage area networks and IP-based access heighten the urgency of SAN security.

Stopping Bad Behavior at Endpoints

47

Cisco Security Agent prevents attacks on servers and desktop PCs by enforcing behavioral policies.



6



19



73

IN EVERY ISSUE	
Mail	3
Acquisitions	5
Calendar	5
Networkers	F
Tech Tips	18
Advertiser Index	89
Cache File	90
The 5th Wave	90

TECHNOLOGY

BROADBAND: Wideband Protocol for DOCSIS	19
Cable operators get ten times the bandwidth at one-tenth the cost of today's cable data service—over existing networks.	

ROUTING: Who's Afraid of DUAL-3-SIA?	23
Cisco IOS Software enhancements improve EIGRP active route processing.	

ENTERPRISE SOLUTIONS

The RFID-Ready Network	53
IP-based network connectivity for RF Identification deployments.	

Data Center Networking	57
Designing the server farm access layer.	

SERVICE PROVIDER SOLUTIONS

Safe Metro Aggregation	61
Innovative Catalyst switch security and QoS features bring reliable, resilient, secure, high performance to the metro aggregation layer.	

The Service Exchange Framework	65
Mastering services means comprehending and controlling every packet and policy in your network. Here's how to do it.	

IP/MPLS Interprovider	69
Extending network infrastructures and services beyond administrative boundaries.	

SMALL AND MIDSIZED BUSINESSES

Integrated Services Routers in the Small Office	73
Cisco extends integrated services routers with new models and integrated wireless across the portfolio.	

Buying Strategies	76
Purchasing refurbished hardware from a reputable source can pay off long term for SMBs.	

Digital Security	77
Financial institutions manage risk and regulatory compliance proactively, with Cisco Self-Defending Networks.	

DEPARTMENTS

From the Editor	1	Technically Speaking	83
Security Is as Security Does		Trends in cluster and grid computing	
User Connection	5	New Product Dispatches	84
New Cisco Powered Network Designation • Cisco Connected Car • Certifications Update		What's new from Cisco over the past quarter.	
Tech Tips & Training	9	NetPro Expert	88
Troubleshooting Cisco IPCC • Deploying Cisco Security Agent • Reader Tips		Configuring and troubleshooting dial-related issues.	

PACKET MAGAZINE

David Ball
Publisher and Editor in Chief

Jennifer Redovian
Managing Editor

Susan Borton
Senior Editor

Joanie Wexler
Contributing Editor

Robert J. Smith
Sunset Custom Publishing
Project Manager

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

Jeff Brand
Art Director

Emily Burch
Designer

Ellen Sokoloff
Diagram Illustrator

Bill Littell
Print Production Manager

Valerie Marliac
Promotions Manager

Achille Bigliardi
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, iO, Linksys, *Packet*, and PIX are registered trademarks or trademarks of Cisco Systems, Inc., and/or its affiliates in the USA and certain other countries. All other trademarks mentioned in this publication are the property of their respective owners.

Packet copyright © 2005 by Cisco Systems, Inc. All rights reserved. Printed in the USA.

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

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

This magazine is printed on recycled paper.



10%
TOTAL RECOVERED FIBER

FROM THE EDITOR

Security Is as Security Does

In the past few months, I've downloaded more than one spyware-combating program on my home PC and become reacquainted with Cisco Security Agent on my work laptop. While these are just two relatively simple security measures I've taken as a full-time remote worker and fan of Web surfing, they are important reminders of the pervasive, dynamic role security plays in today's networks.

Fifteen months ago when we first wrote about Cisco's Self-Defending Network strategy, we emphasized the business need for *integrated security*—that is, embedding security capabilities into network elements such as routers and switches—and *collaborative security*, whereby those embedded security capabilities are linked across the infrastructure and extended to user endpoints. In this issue, we elaborate on the Cisco products and technologies launched over the past year that strengthen these areas of security, as well as address the third prong of Cisco's Self-Defending Network strategy: *adaptive threat defense*—which aims to protect every packet and every packet flow on a network.

Increasingly, security attacks are being introduced from within Web-enabled applications, which open the door to application abuse as traffic traverses multiple networks. For this reason alone, networks require distributed security capabilities, networkwide awareness of the context of endpoint credentials as host behavior and status change, and authentication mechanisms ensuring that those credentials can be trusted. In March, Cisco introduced several products that advance a company's ability to protect every packet and every packet flow on the network. Among their attributes, the new products protect against application abuse, identify and thwart intrusions at Layer 7, and eliminate the sources of attacks—using capabilities such as deep-packet inspection, networkwide event correlation, context-based policies, and policy auditing. Read all about these products and technologies beginning with the article "In Self Defense," page 26.

As far as identifying and mitigating security risks for advanced technologies, check out "Keeping Voice Confidential" (page 34) and "Eradicating Wireless Intruders" (page 39). In the article "SAN Security: Beyond Zoning" (page 42), you'll find best practices for securing storage area networks and learn why Gartner awarded the Cisco MDS 9000 Series multilayer SAN switches an A++ for Fibre Channel SAN fabric security. "Stopping Bad Behavior at Endpoints" (page 47) lays out the latest capabilities in Cisco Security Agent Version 4.5 software that use behavioral policies to thwart spyware and adware, among other intrusive network behaviors. Version 4.5 also adds compatibility with operating systems outside the US, and expands platform support to include Linux servers and desktops as well as Windows clusters.

Security will always be an evolving process for companies. In addition to hackers, there will inevitably be more virulent viruses, clever worms, and surreptitious network behaviors to contend with. Companies with a watchful eye and strategic consideration for the three primary areas of a Self-Defending Network will be as ready as they can be to thwart the next new security threat.

David A. Ball

David Ball
Editor in Chief
daball@cisco.com



Rob Brodman

MAIL

Allergic to Spyware

We have Cisco Security Agent installed on our system but some of our PCs are still being infected with the Virtual Bouncer and VX2 spyware. Isn't the CSA product supposed to help prevent our PCs from having their registries edited by this kind of junk?

—Marty Browne, Allergy & Asthma Associates, Houston, Texas, USA

Cisco's Josh Huston, author of the NetPro Expert column, "Boosting Network Security Using Cisco Security Agent," (Fourth Quarter 2004) responds:

Spyware of the types you mention is usually installed along with programs that end users want to install. Because users want to install these items, they will likely choose "yes" to any popups provided that require interaction. There are many ways you could provide more direction to your users on these items. As one example, you could modify your policies to put tighter restrictions on the HKEY_LOCAL_MACHINE\software\microsoft\windows\currentversion\run registry key to stop any writing to that key unless it is a corporate approved installer. You could also modify the pop-up text to indicate that this looks like a possible spyware program.

Tips and Techniques

Recently, reader Alain Moretti submitted a script to clear a CLI session using SNMP ("Reader Tips," First Quarter 2005). My organization is using a similar script but our community string includes a character (\$) that the router does not interpret correctly. How do you modify the script to use special characters in the community string?

—Jeff Lavender, Verizon Federal Network Systems, Stuttgart, Germany



The following is a response from Cisco technical support:

The problem is that the invoking shell wants to interpret any \$xxx string as a variable. To avoid this, put single quotes around the offending string, for example: script router_id 'some\$string' 0.

Home Networking

I have just wired my house for Ethernet Category 5 wiring and have connected everything using an eight-port hub. I have been told that a router would allow me to access the Internet with my network and still maintain a bit of security. Most all of the routers for sale today must use DSL or cable to connect to the Internet, but I would like to use my dialup connection. What would your recommendations be? If you could give me a few URLs to research, I would be most appreciative.

—Charles T. Olinda, Atlantic County Institute of Technology, Mays Landing, New Jersey, USA

You can set up most small, home office routers to act as a firewall (security device) to protect your computer resources. Cisco has a wide variety of routers to choose from. Check out this educational site at Linksys, the division of Cisco that markets home and small office networking solutions: cisco.com/packet/172_2a1.—Editors

For Beginners

I have been working in networking for the past year and a half and am a faithful Packet reader, but it is often difficult to digest the complicated scenarios you describe. Please consider publishing more material for beginners, specifically on Multiprotocol Label Switching (MPLS), IP Security virtual private networking (IPSec VPN), and IP telephony.

—Sumedh Dharwadkar, Tata International Ltd., Pune, India

We appreciate your desire to read more articles in Packet written for beginners, and it's our goal to strike a balance in content that satisfies our readers' range of technical experience. Regarding your topics of interest (MPLS, IPSec VPN, and IP telephony), following are a few links that you may find useful:

- "MPLS FAQ for Beginners" cisco.com/packet/172_2a2
- "Cisco IOS Software and Multiprotocol Label Switching" cisco.com/packet/172_2a3
- "How Virtual Private Networks Work" cisco.com/packet/172_2a4
- "Decoding IPSec: Understanding the Protocols of Virtual Private Networks" cisco.com/packet/172_2a5
- "A Primer for Implementing a Cisco Virtual Private Network" cisco.com/packet/172_2a6
- "Cisco IP Telephony Overview" cisco.com/packet/172_2a7

—Editors

Send your comments to Packet

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

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

New IP VPN Multiservice QoS Certification for Service Providers

A new Cisco Powered Network Quality of Service (QoS) certification helps organizations that are converging voice and data networks evaluate providers of managed virtual private network (VPN) services. The new certification verifies that a managed IP VPN service meets Cisco-defined best-practice criteria for delivery of real-time voice and video services.



MARK OF QUALITY The Cisco Powered Network mark indicates multiservice IP VPN services that meet Cisco quality standards for real-time voice, video, and other business-critical applications.

The new QoS Certification requires service providers to undergo an annual onsite assessment to validate that the service provider follows best practices for delivering recommended levels of network performance (including latency, jitter, and packet loss) and customer support. This assessment includes the following intra-continental performance requirements for delivery of real-time voice or video packets from customer edge to customer edge:

- Maximum 150 ms one-way delay for voice/video packets
- Maximum one-way packet jitter of 30 ms for voice/video traffic
- Maximum voice/video packet loss of 1.0 percent

For more information, visit cisco.com/packet/172_3a1.

About the Cisco Powered Network Program

The Cisco Powered Network program identifies service providers that use Cisco equipment in their networks end-to-end and meet Cisco standards for service quality and support. Cisco Powered Network program members are committed to maintaining high

levels of network quality and providing services that offer unmatched interoperability with enterprise networks built on Cisco equipment. These service providers bring businesses specialized expertise, greater efficiency, end-to-end security, cutting edge technologies and access to global network resources, allowing businesses to extend the power of their Cisco network and to focus on their core business.

Currently, more than 390 service providers worldwide are members of the Cisco Powered Network program. For more information, visit cisco.com/cpn. ■

Cisco Acquires Topspin Communications

Cisco has announced the acquisition of privately held Topspin Communications, Inc., of Mountain View, California. Topspin is a leading provider of intelligent server fabric switches, a new class of server networking equipment that promotes resource flexibility and dramatically reduces equipment and management costs.

With this acquisition, Cisco will be able to provide its customers with an end-to-end data center switching capability with specialized networking technology and services that allow them to build their data centers in a flexible, grid-like fashion. Topspin's Infini-Band-based data center switching solutions complement Cisco's existing network and storage switching solutions, including the Catalyst switching platform and the MDS Family switches for storage area networks. Topspin products and technology are an integral part of solutions offered by leading system vendors including Dell, HP, IBM, NEC, and Sun.

The server fabric switch market is an emerging market opportunity within the data center driven by customers' need for an intelligent, high-performance fabric for server virtualization, clustered enterprise applications, and grid/utility computing. For more information on cluster and grid computing, see "Technically Speaking," page 83.

Topspin's 135 employees in California and Bangalore, India, will join the Cisco Data Center, Switching, and Wireless Technology Group. ■

CISCO WORLDWIDE EVENTS

June 19–24, 2005	Networkers, Las Vegas, Nevada, USA
July 18–22, 2005	Networking Solutions Technical Conference (NSTC), Montreal, Quebec, Canada
Aug. 29–Sept. 1, 2005	VoiceCon, San Diego, California, USA
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. 19–21, 2005	Networkers China, Beijing, China

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

Cisco Connected Car Demonstrates IP Technology on the Move

A technology showroom on wheels, the Cisco Connected Car has been touring Northern Europe for the past several months demonstrating to public safety organizations how wireless communications can help police, fire, and ambulance personnel improve their responses to emergencies and provide more effective citizen services. The car is a Volvo V70 rigged with a Cisco 3200 Series Wireless and Mobile Access Router in the trunk, which can roam across wireless LANs, mobile phones, and TETRA networks to provide communication between emergency headquarters and the moving vehicle, regardless of physical location. Also in the car is a notebook computer with connection to external networks provided by a General Packet Radio Service (GPRS) modem mounted in the car. A global positioning system (GPS) device allows the car to be traced wherever it goes, and a network camera with IP video surveillance software from Milestone Systems lets emergency staff film incidents and instantly transmit the footage to headquarters using the in-vehicle router.

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



CISCO CONNECTED CAR The model car shows how IP technologies make it possible to react more quickly and effectively to accidents, disasters, riots, or crimes.

Cisco Certifications Roundup

New IP Contact Center Express Specialist Certification

The new Cisco IP Contact Center Express Specialist certification is the latest focused certification to be offered by the Cisco Career Certifications program. The new certification was created in response to heightened customer and channel partner demand for knowledgeable network professionals who can successfully plan, install, configure, troubleshoot and manage Cisco's IP Contact Center (IPCC) Express Edition.

Cisco's IP Contact Center Express Edition provides a feature-rich, tightly integrated automatic call distributor (ACD), Interactive Voice Response (IVR) services, and computer telephony integration (CTI) on a single platform, allowing organizations to significantly reduce operating costs and improve customer satisfaction.

To earn the IP Contact Center Express Specialist certification, candidates must hold an active Cisco associate-level CCNA certification and pass one additional exam. To help candidates prepare for the exam, the IPCC Express and IP IVR Deployment (CRSD) Version 3.5 course covers planning and designing an IPCC Express deployment; installing the IPCC Express software; building self-service work flows, call routing and queuing logic; configuring agent skills and supervisor teams; creating agent screen pops; integrating to CRM databases, setting up silent monitoring and recording; generating reports; and troubleshooting.

Enhanced Recertification Policy for CCSP

Cisco has enhanced the recertification policy for the CCSP professional-level certification. Now, in addition to passing any CCIE written exam or achieving the CCIE expert-level certification individuals can be recertified for the CCSP professional-level certification by taking the Cisco SAFE Implementation exam.

"In today's fast paced Internet economy, recertification is a tangible indication to both IT professionals and the organizations that employ them that Cisco certification holders are current on the latest technology trends," says Don Field, director of certifications at Cisco.

The CCSP certification indicates training in advanced knowledge of securing Cisco networks. With a CCSP, a network professional can secure and manage network infrastructures to protect productivity and reduce costs. The content emphasizes topics such as perimeter security, virtual private networks, and intrusion protection, as well as how to combine these technologies in a single, integrated network security solution.

CCSP certifications are valid for three years from the date they are awarded, and must be renewed prior to their expiration date. To determine the current status of a CCSP certification, visit the Cisco Career Certifications Tracking System at cisco.com/go/certifications/login.

For complete details about the training and exam requirements for the CCSP and other certifications, visit cisco.com/go/certifications.

Training can be purchased using Cisco Learning Credits, an easy method of payment for Cisco authorized training, redeemable through participating Cisco Learning Partners worldwide. For information about Cisco Learning Credits, visit cisco.com/go/learningcredits.

CCIE Sizzles with Hottest Certification for 2005

CertCities.com, an online publication for certified IT professionals, recently awarded Cisco's CCIE expert-level certification the "Hottest Certification for 2005" and "Most Respected High-Level Certification" awards. The annual CertCities.com Readers' Choice Awards combine information obtained from readers on their intent to pursue certifications with survey scores reflecting general enthusiasm for certification programs.

The "Hottest Certification" is awarded to programs that industry analysts and professionals expect to grow the fastest.

For the third year in a row, the CCIE program has swept this category, thanks to the almost mythic difficulty of the CCIE lab exam, a grueling, eight-hour, hands-on practicum. Currently, 8,852 network industry professionals have achieved CCIE certification.

"Customers respect the CCIE certification and they appreciate the certification process," says Rami Kandah, content engineer with Cisco Customer Advocacy. "Four out of ten customer support requests ask for a CCIE."

Career Builders

Cisco offers three levels of career certification—Associate, Professional and Expert—in areas such as Routing and Switching, Network Security, and Voice over IP. In addition, numerous Cisco Qualified Specialist certifications are available in specific technologies, solutions or job roles. The CCSP was the second place winner for the "Best Security Certification" award and is the only curriculum accredited by the US National Security Agency.

The CCNA certification was a finalist in the "Best Entry-Level Certification" category and the CCNP was rated second for "Best Mid-Level Networking Certification."

Cisco Certifications represent three levels of expertise: Associate, Professional, and Expert.

For more information on Cisco career certifications, visit cisco.com/go/certifications. For additional information on CertCities.com's Third Annual Readers' Choice Awards, visit cisco.com/packet/172_3b1. ■

Minimizing Abandoned Calls in Cisco IP Contact Center

By Sandeep Gupta

Cisco IP Contact Center (IPCC) Enterprise Edition (formerly Cisco IP Contact Center), delivers intelligent contact routing, call treatment, network-to-desktop computer telephony integration (CTI), and multichannel contact management over an IP infrastructure. By delivering functionality in a unified solution, Cisco IPCC enables companies to rapidly deploy a distributed contact center infrastructure.

Reducing abandoned contact center calls is a key factor in maintaining customer satisfaction, so you should become familiar with IPCC Call Disposition 1 (CD1). A CD1 call is one that is abandoned—or dropped—before terminating at a target device such as an automatic call distributor (ACD) system or agent desktop. Abandoned calls occur when a caller hangs up while on hold. This article helps you

diagnose the causes of abandoned calls and provides troubleshooting tips.

Before Diagnosis

Make sure you are familiar with the five major components of Cisco IPCC and the system architecture. For details, see the product information for Cisco IPCC at cisco.com/packet/172_4a1.

The system architecture of Cisco IPCC differs considerably from Cisco Intelligent Contact Management (ICM) configurations using legacy ACD, and these differences affect IPCC reports, which you use to diagnose abandoned calls. For details on the differences, see the Cisco IPCC Administration Guide at cisco.com/packet/172_4a2.

DIAGNOSE ABANDONED CALLS

POTENTIAL CAUSE

The Agent Reservation timer is set too low. The timer is used to set the maximum allowable time between agent selection and when the agent actually receives the call. Any call that exceeds the time is reported as CD1.

Agents enable call forwarding to voice mail or to another agent phone. No events are reported to Peripheral Gateway (PG), resulting in CD1 calls.

In each of the following three scenarios, the PG cannot match a call-arrival event with a pre-call indicator-event, and the call is reported as CD1:

- Agent goes into Talking state (presses the headset or speaker button) when the ICM router is about to send a call to that agent.
- Agent uses the ACD DN (the Directory Number on which the agent receives routed calls) while in the AVAILABLE state.
- Agent-to-agent calls occur while agents are in the AVAILABLE state.

Caller disconnects/hangs up while the call is being routed to an agent. This rarely happens in complete IP telephony networks, but happens more frequently in the time division multiplexing (TDM) world, where call setup might take longer. The PG cannot match a call-arrival event with a pre-call indicator-event, and the call is reported as CD1.

Call is sent to an incorrect label for an agent (device target). IPCC sends the label (or DN) to its PG to route the call to an agent phone or Cisco IP Interactive Voice Response (IP IVR). PG or IP IVR passes this label to Cisco CallManager. If Cisco CallManager is unable to route the call to a device identified by this label, the call does not establish and is reported as CD1.

POSSIBLE FIX

The timer default is 7 seconds. Adjust the timer (up to 12 seconds).

Note: The agent is placed in a NOT READY state after two consecutive calls are not received.

- Collect and analyze Termination Call Detail (TCD) and Route Call Detail (RCD) to identify CD1 calls.
- Identify the label/Directory Number (DN) and check for any call forwarding set on an agent DN. You might see multiple CD1s for this agent DN.
- Advise agents to not use call forwarding.
- Set up a phone button template for agent phones; remove the CFwdALL key.

- Provide each agent with a second line to be used for non-IPCC calls.
- Do not assign the ACD DN to the first line on the agent phone.
- Block agent-to-agent, internal and PSTN calls from ACD DN by using Calling Search Space and partition.
- Use a translation pattern that routes calls to a CTI route point that invokes a script for agent-to-agent calls.
- Advise agents to enter the NOT READY state when they make non-IPCC calls. IPCC will not route calls to these agents.

1. Analyze IPCC reports to identify whether the abandoned call rate is within an acceptable range.
2. If the rate is unacceptable, identify a pattern (time of day, day of week, a particular agent or group of agents, call volume changes, and so on).
3. If you suspect the problem, start to capture logs and analyze them.

1. Examine the IPCC configuration (device target labels, routing clients, and translation routes). IPCC sends a label to a routing client; the client is responsible for routing and completing the call.
2. If the labels are set correctly, analyze Cisco CallManager and IP-IVR logs to determine the cause.

POTENTIAL CAUSE	POSSIBLE FIX
You are using Call Pick groups and Call Park on ACD DN. These features are not supported and result in CD1 calls.	Do not use Call Pick groups and Call Park on ACD DN.
You are using ACD DN as a shared line. Call Waiting, Call Park, Call Pickup, and Call Manager Pilot Point and Hunt groups are not supported and result in CD1 calls.	Do not use ACD DN as a shared line.
Improper Calling Search Space (CSS)/Partition on CTI route point, Gateway, and agent ACD DN. Cisco CallManager uses CSS and Partition to provide class of service (COS) to phones, gateways, and applications so that they can call limited, selected, or all devices. Based on the CSS assigned on the device, it is possible that a call may not complete when PG sends the label to Cisco CallManager. The PG cannot match a call-arrival event with a pre-call indicator-event, and the call is reported as CD1.	<ol style="list-style-type: none"> 1. Collect and analyze TCD and RCD to identify CD1 calls. 2. Identify the label/DN and check its partition in Cisco CallManager. This DN partition may be different than that defined on other working DNs. 3. Change the CSS on the Gateway or CTI route point to include ACD DN partition.
Improper design and sizing of bandwidth for location (Cisco CallManager configuration) to route calls to remote agents. Cisco CallManager provides Call Admission Control (CAC) using bandwidth to avoid congestion on WAN. If agent calls cannot be completed due to bandwidth, PG cannot match a call-arrival event with a pre-call indicator-event, and the call is reported as CD1.	<ol style="list-style-type: none"> 1. Collect and analyze TCD and RCD to identify CD1 calls. 2. Identify the label/DN and check if any or all Gateways are using a location. 3. Enable Cisco CallManager tracing for the location (using Service Parameters) and examine Cisco CallManager logs. 4. If possible, increase bandwidth in location based on WAN bandwidth sizing, and/or check for bandwidth leaks.
Improper Region configuration and lack of transcoders to route calls to agents. Cisco CallManager uses Region to allow different codecs to be assigned to remote devices, a phone, or IP IVR. This saves bandwidth on the WAN by sending G729 calls to agents in remote locations. Transcoders are used by devices that are unable to negotiate codecs dynamically. Lack of transcoding resources result in no events being reported to PG, and calls are reported as CD1.	<ol style="list-style-type: none"> 1. Collect and analyze TCD and RCD to identify CD1 calls. 2. Identify whether label (DN) is in a different region than that of Gateway or CTI route point. Region is assigned on Device pool. 3. Ensure that the device pool has a Media Group Resource List that contains a Media Resource Group with transcoder.
An Automated Alternate Routing (AAR) configured in Cisco CallManager to route calls in the event of network congestion results in a setup delay to an agent phone. Cisco CallManager uses AAR to route calls via PSTN when it is not possible to route a call based on CAC (bandwidth defined in location). AAR is not tested and supported to work with IPCC since there is a strong possibility of delay in call setup via PSTN, resulting in CD1 calls.	<ol style="list-style-type: none"> 1. Collect and analyze TCD and RCD to identify CD1 calls. 2. Identify label and remove AAR from Cisco CallManager configuration.
Network issues (WAN outage, congestion, latencies, and so on) lead to a delay in call setup. PG cannot match a call-arrival event with a pre-call indicator-event, and the call is reported as CD1.	<ul style="list-style-type: none"> ■ Verify that the network interface card (NIC) speed and switch ports for all devices, servers, and desktops are hard coded to 100 MB full duplex. ■ Provision WAN for outage, glitch, sizing, and any additional bandwidth requirements.

Diagnosis and Troubleshooting

The table above lists the most common causes of abandoned calls, along with possible fixes. Before attempting these procedures, all agents on your team should be properly trained.

Best Practices

You can further optimize Cisco IPCC functionality and provide better stability by following these best practices, which are based on successful customer implementations:

- Reroute busy/failed calls to voice mail/attendant by setting up Forward on CTI route points.



SANDEEP GUPTA is a customer support engineer at the Cisco Technical Assistance Center (TAC) and a technical leader for the Cisco IPCC Enterprise team. He can be reached at sandgupt@cisco.com.

- Ensure that IPCC can pull a call back into the queue and reroute unanswered calls by setting up a Forward On No Answer (FONA) timer in Cisco CallManager that is higher than the Redirect On No Answer (RONA) in ICM.
- For best system performance, hard code NIC speed and switch ports to 100 MB full duplex for Cisco CallManager, IP IVR, PG, agent desktops, etc.
- Avoid or minimize device registration and call processing on the Cisco CallManager publisher node. Avoid call processing on the Cisco CallManager publisher, and perform configuration changes only on the publisher while subscribers are processing calls. Any major device addition or deletion performed by subscribers can cause issues.
- Set up device pools (using Cisco CallManager Groups) so that the device weight per Cisco CallManager node is within guidelines defined in the Cisco CallManager Solution Reference Network Design Guide (cisco.com/packet/172_4a3).

Analyze IPCC Reports and Logs

To generate effective reports and logs, set up IPCC tracing to capture as much information as possible about calls, call flow, network topology, etc. If you enable tracing for troubleshooting purposes, this might have an adverse impact on system performance, particularly during peak traffic hours. For detailed information about setting up traces, refer to the following documents:

- **Configuring Cisco CallManager Trace Parameters:**
cisco.com/packet/172_4a4
- **CTI Manager:**
cisco.com/packet/172_4a5
- **SDL (Cisco CallManager and CTI):**
cisco.com/packet/172_4a6
- **IPCC PG:**
cisco.com/packet/172_4a7
- **IP IVR (enable SS_TEL, LIB_ICM, and LIB_MEDIA debugging bits):**
cisco.com/packet/172_4a8
- **TCD and RCD from SQL (make these queries as similar as possible with respect to time):**

```
Select * from Route_Call_Detail where DateTime >=
"00:0000/00/0000" and DateTime <= "00:00 00/00/0000"
```

```
Select * from Termination_Call_Detail where DateTime >=
"00:00 00/00/0000" and DateTime <= "00:00 00/00/0000"
```

Below are examples of using logs to identify the causes of abandoned calls. In general, trace any failed call in the logs and check the call against known causes of CD1 (ICM configuration, Cisco CallManager, IP IVR, etc.).

- **Insufficient latency set in Agent Reservation timer.** Search for "DtAbort" or "ProcessReservedTimeout" in PIM logs. The following errors in EAPIM logs confirm the problem:

```
TelephonyDriver::ProcessReservedTimeout: No call
arrived to match PreCall message.
```

```
MISSED 2 consecutive routed calls. FORCING TO NOT READY
STATE.
```

- **Wrong label for device target.** View EA PIM event logs; you should be able to associate the ProcessReservedTimeout with a single device target. It is possible that IPCC has a correct label, but Cisco CallManager is still unable to route. Refer to logs from IP IVR, PG, and Cisco CallManager, and track calls as defined in the call flow. The following errors in the IP IVR – MIVR logs confirm the problem:

```
MIVR-SS_TEL-3-REDIRECT_FAILED:Redirect failed:
All Call ids=CallID:32
```

```
MIVR-SS_TEL-3-EXCEPTION:com.cisco.jtapi.
InvalidPartyExceptionImpl: Redirect failed because of
an invalid destination.
```

- **Lack of transcoders or incorrect region/codec defined in the device pool.** Track the mismatch by reviewing snippets in Cisco CallManager logs. The following confirms the problem:

```
CCM|SPROC - Incompatible Regions or Capabilities
```

```
CCM|SPROC - Origination Side: Region=Default|
```

```
CCM|SPROC - Origination Side: Capabilities
MaxFramesPerPacket|
```

```
CCM|SPROC - Media_Payload_G711Alaw64k, 20|
```

```
CCM|SPROC - Destination Side: Region=Default|
```

```
CCM|SPROC - Destination Side: Capabilities
MaxFramesPerPacket|
```

```
CCM|SPROC - Media_Payload_G711Ulaw64k, 160|
```

- **Insufficient bandwidth.** Check snippets in Cisco CallManager logs. Confirm that location tracing is enabled in service parameters. The following confirms the problem:

```
cdccPID=(1.14.372335) Orig=1 not enough bw. bw=24
curr=18 max=210|
```

- **CSS/Partition issue.** Search for the following pattern in Cisco CallManager and block it just before Cisco CallManager performs digit analysis:

```
StartTone tone=37(ReorderTone)
```

You can minimize abandoned calls in Cisco IPCC by applying known fixes in your IPCC solution, using the troubleshooting table in this article to identify causes of abandoned calls, and applying the recommended fixes. By following best practices, including setting up effective reports and logs, you should notice a significant decrease in dropped calls. ■

FURTHER READING

- **Cisco IPCC Enterprise Administrator Guide**
cisco.com/packet/172_4a9
- **Cisco ICM Database Schema**
cisco.com/packet/172_4a10
- **Cisco IPCC Enterprise Designing, Sizing, and Planning**
cisco.com/packet/172_4a11

Deploying CSA

A Guide to Successfully Implementing Cisco Security Agent

By Brian Cincera

Host Intrusion Prevention (HIP) software is fast emerging as a standard for corporate desktops. Like the antivirus software that preceded them, HIP applications enable organizations to mitigate the growing risks of malware, malicious code, spyware, hacker attacks, and their ilk at a lower cost with less human intervention.

Intrusion prevention applications are the latest in security defense software, and several of them move beyond detection of security threats by signature to detecting threats by their behavior. The natural benefit is that organizations can significantly improve their threat defense while dramatically reducing their management requirements for patches and signatures.

Formerly the Okena StormWatch product, Cisco Security Agent is a leader in the host intrusion prevention market and has a place on any desktop. (For a related article, see “Stopping Bad Behavior at End-points,” page 47.)

Power and Protection

Most organizations today include remote offices and remote users, which have special needs that can stress or even break internal support systems. In addition to providing enhanced security capabilities to remote users, Cisco Security Agent is one of those rare products that can actually make supporting remote offices and “road warriors” easier in the following ways.

Deployment. Agent deployments can be initiated using a link to a Website sent in e-mail for installation of the agents. The user “pull” deployment model allows end users on low-speed connections to pick the best time to initiate deployment of the agent. Traditional deployment methods such as group policies or software deployment tools can also be used.

Offline capability. Once deployed, Cisco Security Agent is fully functional when disconnected from an organization’s network, offering complete protection for users who connect at airports, hotels, or client locations. When remote systems are reconnected to the home network, logs are transmitted and updates,

if needed, are sent—all with negligible impact on even the lowest bandwidth WANs.

Policy tuning. Policies can be tuned to meet specific business or operational conditions. For example, users can be offered options to allow behavior that Cisco Security Agent flags as anomalous to keep false positives from preventing legitimate but uncommon activities.

Many organizations value the ability to limit or eliminate altogether any interaction between Cisco Security Agent and end users. For example, agents can be deployed to desktops that limit a user’s ability to install or uninstall applications, connect external peripherals such as USB drives, run file-sharing applications, or even connect to other computers. In a distant location, where onsite IT support is not available, this level of control over an end user’s PC can help eliminate a common set of help-desk-intensive, misuse conditions.

In addition to providing enhanced security capabilities to remote users, Cisco Security Agent is one of those rare products that can actually make supporting remote offices and “road warriors” easier.

Management Environment

The Cisco Security Agent management environment allows for grouping user-level agents by many profile factors, including status as a remote-office user or location. By creating and enforcing policies in a group fashion, management adjustments to agents can be made specifically to suit remote users. A single change can be pushed to all group members, relieving support staff from having to manage end users as individual units.

Deployment Challenges

There are two primary challenges to consider when deploying Cisco Security Agent into remote office environments. The first is enforcement of agent deployment. When deploying the agent software,

organizations have a choice of traditional software distribution capabilities, an end-user “pull” from an e-mail link, or manual installation. For remote users and offices that are connected by low-speed links, initial agent deployment can be time consuming. Although policy updates and log transmissions are very small, initial agent deployment files can be 5 megabytes (MB) or more. Organizations need to balance the real need to deploy the agent for security purposes with the user inconvenience associated with a forced installation at an inopportune moment.



BRIAN CINCERA is vice president for security solutions at Greenwich Technology Partners. He can be reached at bcincera@greenwichtech.com.

ABOUT GREENWICH TECHNOLOGY PARTNERS

Greenwich Technology Partners

(greenwichtech.com) is a vendor-independent IT professional services firm that helps organizations maximize the return on their IT investments. GTP delivers strategic business-focused IT solutions in the practice areas most vital to large and medium enterprises, including information security, multiservice networks, infrastructure optimization, and application resilience.

The second challenge is the support tradeoff that organizations make when determining desktop policies. Deploying the recommended Common Security, Required Windows System, and Desktop modules provides protection against 90 percent of the common threats. Policy adjustments made to enhance security protections or limit user actions increase the chances that end users will encounter false positives or attempt a function that Cisco Security Agent prevents.

While such actions prevent a system breach or the infection of a computer, they will also likely result in a call to the help desk anyway. Thus, organizations must consider the balance of protection and support effort when considering the optimal remote-office agent policies.

Sound decision making on these two issues that is appropriate to the environment will help smooth and speed the deployment to remote offices and remote users. The remaining architecture, implementation, and operations decisions for supporting remote offices are the same as they are for any desktop user community.

Lessons Learned

Having performed many large and small Cisco Security Agent implementations, Greenwich Technology Partners has encountered many different scenarios and learned some important lessons.

Provide reasonable management redundancy. It is a good idea to install two or more management consoles for Cisco Security Agent, but it is not necessary to have fault-tolerant redundancy. Agents continue to operate fully, even when disconnected from the management station. Policy updates and log collection are the only communications between the agents and the management station.

Deploy default agent policies. The default desktop and server protection policies are predefined to address roughly 90 percent of likely threats without blocking legitimate use. Tuning might still be necessary, but using the default is a good starting point. Many organizations tinker with the controls as soon as they install the first agent, but aside from very specific policy enforcement issues, it is best to try the default settings first and adjust from there.

Deploy in test mode first. Test mode allows real agents with real protection policies to operate in “detect and log” mode. This allows you to see what *would* be blocked without actually performing blocking. By monitoring log activity, you can easily determine the agent policy tuning that will bridge the last 10 percent of protection without using end users as your test subjects.

Consider policy enforcement, not just threat prevention. Most organizations realize the greatest benefits when they begin to use the agents to enforce policy. You can configure Cisco Security Agent to enforce policies beyond threat prevention. Functions such as limiting instant messaging, blocking file-sharing applications, or preventing users from copying proprietary information to removable drives are all possible with Cisco Security Agent. The combination of threat prevention and policy enforcement capabilities is powerful.

Evaluate the tradeoffs of end-user interaction. You can configure Cisco Security Agent to allow or deny end-user interaction with the agent.

In its most interactive form, Cisco Security Agent momentarily blocks malicious behavior and allows users to decide if they want to proceed. However this tends to be problematic, because most users do not understand the warning and therefore do not make the correct choice about whether to proceed.

Intrusion prevention technologies are an important step toward improved resilience in the face of growing threats.

In its most protective form, Cisco Security Agent immediately blocks malicious behavior without any warning to end users. This is also problematic because in the event they were performing a legitimate action that Cisco Security Agent suspects is malicious, users are offered no clues as to why their actions were unsuccessful.

Most organizations must plan carefully and understand the impact on end users and the potential increase in help-desk calls that either decision can create. At Greenwich Technology Partners, we have found that most organizations try a starting position somewhere in the middle and gravitate toward less user interaction.

Build a business case. Technology alone can only take you so far. The organizations that have been the most successful in deploying Cisco Security Agent have been able to show how it returned more value than it cost in resources to deploy. In the security realm this has always been difficult to demonstrate.

At Greenwich Technology Partners, we have found the best business case arguments are made on a cost avoidance model. Easy targets are to calculate costs associated with common vulnerabilities such as Blaster, SQL Slammer, or I Love You. Each of these was the result of vulnerabilities that were ultimately resolved by an operating system patch.

If your organization is like most, the emergency response to system recovery or to emergency patch deployment was staggering. (One of our clients reports an incremental cost of US\$50 per desktop for Blaster remediation.)

While products such as Cisco Security Agent do not eliminate the need to patch, they do eliminate the need for emergency patches, at the same time that they prevent these types of threats from damaging computers. Many clients can build a business case for intrusion prevention with payback in less than two years.

Improving Business Resilience

Intrusion prevention technologies are an important step toward improved resilience in the face of growing threats. These technologies also support steps toward policy enforcement that are often mandated by government regulation.

Cisco Security Agent provides intrusion prevention capabilities that can meet needs in both areas. In addition, the application makes a great choice in remote office and remote user settings where local IT support is minimal or nonexistent.

As with any important desktop application deployment, sound decision making during the design stage and taking advantage of lessons learned by others can help make the implementation process easier and more successful. ■

FURTHER READING

- Cisco Security Agent
cisco.com/packet/172_4b1
- Cisco Self-Defending Networks
cisco.com/go/sdn
- Cisco Security and VPN
cisco.com/go/security

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 Reconfiguring Layer 2 Switch Addresses

Here is a suggestion for changing Layer 2 switch administrative addresses on the fly. We had the additional problem that console access was out of the question because of the geographical distances between the switches. The crucial points that made it possible were: We had Layer 2 connectivity from a Layer 3 device to the switch, and we had remote access to the Layer 3 device via Telnet, and we kept the virtual LAN (VLAN) containing the new address undefined during the configuration actions.

We were faced with the problem of rearranging Layer 2 switches in a different topology. Distances between the switches are up to tens of kilometers, using Ethernet tunneling over a provider network. Rearranging switches involves requesting additional Ethernet tunnels before the rearrangement and requesting terminating superfluous tunnels after the rearrangement.

We wanted to align the administrative addresses of the switches with the new Layer 2 ring topology that they would be part of. In preparation, we wanted to change the administrative addresses of the switches prior to actually changing the topology. Links between the Layer 2 switches and from a Layer 2 switch to a Layer 3 device are all (802.1Q) trunk links.

We used the following procedure:

1. Prepare a new VLAN (VLAN-new) with an IP address on the Layer 3 device. This acts as the gateway for the switches.
2. Include this VLAN in the trunk from the Layer 3 device to the concerned Layer 2 device(s).
3. Define the new VLAN in all the switches in the same Layer 2 structure(s) in the current topology as the switches needing new IP addresses. This is done to ensure that all the trunks allow the new VLAN. If there are VLAN restrictions on trunks, you may have to change them to accept the new VLAN.

Repeat the following procedure for each switch:

1. Connect from a Layer 3 device within the current subnet of the switch concerned to the switch (we used Telnet).
2. Schedule a reload in 30 minutes (to avoid permanent loss of connectivity to the switch in case of mishaps during reconfiguration).
3. Make sure the new VLAN (VLAN-new) is not defined in the switch:

```
no vlan <VLAN-new> (if necessary)
```

4. Change the administrative VLAN and IP address in the switch:

```
interface vlan <VLAN-new>
no shut
ip address <new address> <new mask>
exit
ip default-gateway <address of Layer 3 device in the
new VLAN>
```

5. Define the new VLAN on the switch to make the connection effective:

```
vlan <VLAN-new>
(optional) name <description of VLAN>
exit
```

The connection is now lost. After the spanning tree reconverges, you can reach the switch again by ping-ing the new address. Connect to the switch again using the new address.

6. Optionally, delete the former administrative VLAN in the switch (no interface <old-administrative-VLAN>).
7. Check that the conversation with the switch continues functioning.
8. Save the configuration.
9. Cancel the scheduled reload.

—Paul De Valck, Imtech Telecom NV,
Brussels, Belgium

TIP Saving Very Large Configuration Files

When trying to load a very large configuration (approximately 2 MB) into a Cisco 3725 Router (Cisco IOS Software Release 12.3.9b with the

SSH/3DES feature) it is impossible to store the configuration in the machine's NVRAM in plain or compressed form. The Cisco 3725 does not recognize the IOS command `boot config flash`. Maybe the flash file system on the 3700 Series Router cannot read the flash devices at boot time.

The standard recommendation is to load the configuration from an external TFTP server using the `boot network` or `boot host` command.

To avoid relying on an external host, use this procedure:

1. Perform basic configuration tasks on the router.
2. Configure the two network load commands into the normal NVRAM startup configuration:

```
service config
boot network tftp://10.10.10.10/host-config
boot host tftp://10.10.10.10/startup-config
```

The 10.10.10.10 IP address is the Fast Ethernet address.

3. Configure the following command in the NVRAM startup configuration:

```
tftp-server flash:/startup-config alias startup-config
tftp-server flash:/host-config alias host-config
```

The `host-config` file is a placeholder (it contains remarks, but no IOS instructions) but is necessary because the router attempts to load it sending a local broadcast TFTP packet.

4. Save the configuration into NVRAM.
5. Load the big configuration into the router. If your image supports `scp` you can load the file from your workstation using the `scp` command:

```
scp big-file user@10.10.10.10:flash:/startup-config
```

Alternatively, you can use a perl script that loads the configuration using SNMP. In this case I used the `scp` method because the configuration is a large group of `rtr` commands that another engineer has generated using an external perl program.

At the next reload your router will load the large configuration from the flash file using the local

activated TFTP server. From here you can save the configuration:

```
copy running-config flash:startup-config
```

—*Andrea Montefusco, Kyneste S.p.A., Rome, Italy*

Troubleshooting

TIP Handling Mistyped Commands

Mistyping commands is a common and annoying problem that causes the router to respond as if you typed a hostname. For example:

```
MyRouter#shwo
Translating "shwo"...domain server (10.1.1.2)
% Unknown command or computer name, or unable to find
computer address
MyRouter#
```

In this example, the word `show` is mistyped. To correct this problem, change the preferred transport method:

```
! Console port
line con 0
transport preferred none
! VTY Ports
line vty 0 5
transport preferred none
```

The output shows the lack of a failed connection based on the mistyped keyword:

```
MyRouter#shwo
^
% Invalid input detected at '^' marker.
```

—*Shahzad Rana, ORIX Leasing Pakistan Limited, Karachi, Pakistan*

TIP Determining the Committed Information Rate in a Frame Relay Network

Recently, our service provider changed a Frame Relay port to the remote office from a 256k port to 512k. Unfortunately, we assumed that the provider would automatically change the 128k Committed Information Rate (CIR) to 256k. After many trouble tickets for an unusually high rate of "Discard Eligibles" error messages before we ever came near our CIR, we learned that the provider never adjusted the CIR and it was still set at 128k. We also discovered that it was maladjusted at several other sites. However, the following Cisco IOS Software command will effectively give you a reading on your CIR.

```
New-Chester> show frame-relay map

Serial10/0.102 (up): point-to-point dlci, dlci
102(0x66,0x1860), broadcast, BW =

128000
```

—A.G. Teslicko, Amscan, Inc., Chester, New York, USA

TIP Monitoring Link Quality

Cisco Bit Error Rate Testing (BERT) is a very helpful utility for monitoring link quality. Just loop the suspected media at its far end, configure a BERT profile, and run a test of any 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 and get pass or fail results. The tool enables you to remotely verify your media from a remote command line interface and isolate problems quickly.

```
!
<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 Ltd, 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

Visit the new Cisco Technical Support & Documentation website. Cisco's online technical support and product documentation has been integrated in a new Website on Cisco.com, enabling users to find product, support, and technical information in the same place.

cisco.com/packet/172_4e1

Identify jitter and typical voice quality symptoms. This new TAC Case Collection item describes how to get a general determination of jitter in the network by using the IOS command show call active voice.

cisco.com/packet/172_4e2

Best practices for administration of Cisco Secure ACS for UNIX. These practical guidelines are culled from actual design and deployment experiences of development engineers.

cisco.com/packet/172_4e3

New Q&A for Cisco Traffic Anomaly Detector and Cisco Guard. This new question-and-answer format document provides specific detailed answers to common questions about configuring these two security products.

cisco.com/packet/172_4e4

Detect disconnected calls in Cisco IP IVR. This document describes a script for detecting disconnected calls in the Cisco IP Interactive Voice Response (IP IVR) software.

cisco.com/packet/172_4e5

Receive the latest product information with Cisco Product Alert Tool. Set up a profile in this tool and you will receive automatic e-mail updates about reliability, safety, network security, and end-of-sale issues for the Cisco products you are interested in (available to registered users only).

cisco.com/packet/172_4e6

Wideband Protocol for DOCSIS

Cable operators get ten times the bandwidth at one-tenth the cost of today's cable data service—over existing networks.

By Janet Kreiling

How can cable operators leapfrog fiber to the home (FTTH) and deliver 50 or 100 Mbit/s both up and downstream to each residence (or business) on their hybrid fiber-coaxial (HFC) networks? By simply rearranging equipment that's already in place, thanks to *Wideband Protocol for DOCSIS*, a new technology from Cisco that frees up bandwidth already in that fiber and coax.

Lindsay Schroth, senior analyst for broadband access technologies at the Yankee Group, calls Wideband Protocol for DOCSIS “absolutely a disruptive technology, especially in Europe and Asia where telcos are competing now with very-high-bandwidth DSL.” The technology will become just as important in the US, she adds, but because DSL speeds are currently much lower, most US cable providers will likely wait until the release of the DOCSIS 3.0 standard, which will incorporate it. (DOCSIS—Data over Cable Service Interface Specification—is a set of CableLabs standards that govern delivery of data over cable networks.)

In Europe and Asia, where housing is often very dense, installing FTTH is cost effective, and telcos are delivering 10- or 100-Mbit/s Ethernet in urban areas; several have 1 Gbit/s in sight. In some countries and cities, customer demand is reinforced by governmental mandates requiring telecom carriers to provide very high bandwidth. So, there is already intense interest in the new technology from Asian and European providers. Even in the US, where cable companies serve about two-thirds of broadband homes, there's incentive now. “Users of bandwidth-hungry applications will go with whatever carrier gives them the quality of service they want,” says John Mattson, director of marketing for cable products at Cisco. “If interactive gamers, for instance, can't get the bandwidth for good graphic resolution or the low latency they want on their current service, they'll switch to another provider.”

One of the next big applications is going to be downloading movies, adds Mattson. “High-definition streaming video can consume 20 Mbit/s, or with compression, perhaps 10 to 12 Mbit/s. Downloading to storage will take a few seconds on a 50-Mbit/s link compared to much longer times on a traditional high-speed Internet connection, or even several hours on a lower speed link. Wideband DOCSIS will let cable companies get in on the ground floor.”

Blows Away Speed and Capacity Limitations

Wideband Protocol for DOCSIS offers higher throughput downstream pipes at significantly lower cost, by allowing downstream channels to be added independent of upstream ones, notes John Chapman, the Cisco Distinguished Engineer who created this new wideband technology. “Yet it works with today's DOCSIS



1.x and 2.0 cable modem termination systems, and it takes advantage of the decline in prices for external QAM [quadrature amplitude modulation] devices,” he says. There's a terabit of capacity in the HFC serving a typical 100,000-person city, adds Chapman, “and only 1.9 percent of it is being used.”

Decoupling downstream and upstream channels gets away from the ratio of one down to four or six up, so cable operators can economically offer whatever bandwidth a subscriber wants by grouping together down or upstream channels as needed to form a larger “wideband” channel. The techniques in Wideband DOCSIS for combining channels up and downstream differ somewhat, but both are consonant with current DOCSIS protocols and very economical.

Downstream wideband channels can use external (edge) QAMs, which, because they have less functionality than a cable modem termination system (CMTS), cost less per port. CMTS handles both DOCSIS (all-digital) and non-DOCSIS (analog) traffic such as video on demand and regular broadcasting. Edge

QAMs couple the downstream digital channel onto the analog HFC network.

The core of Wideband DOCSIS is the formatting of DOCSIS frames into 188-byte MPEG-TS packets; the packets are broken into pieces that are transmitted simultaneously by up to 24 or 48 QAM channels. Chapman calls this technique “striping” the packet across the parallel channels. Transmitting these large packets in multiple chunks simultaneously ensures that wideband doesn’t introduce latency. A sequence number embedded in each packet enables the transmit framer to stripe packets on channels as needed, and the receive framer to reassemble them. The QAM channels do not need to be adjacent. If certain QAM channels have already been assigned to non-DOCSIS uses, Wideband DOCSIS uses what’s available.

MPEG-TS packets were chosen as the carrier medium rather than bytes or ordinary packets, Chapman says, because they permit bonding of channels at the transmission convergence layer, above the physical layer and below the MAC layer. Because it does not affect either layer, Wideband DOCSIS is transparent to traditional DOCSIS protocols. “This is very powerful as it has the potential to maximize re-use of the existing DOCSIS environment,” he points out.

Downstream signaling takes place via the standard DOCSIS downstream signaling channel, so both wideband and present-day cable modems can co-exist in the network.

Upstream and downstream transport are different, because the equipment originating signals is different at each end; the CMTS and edge QAM transmit downstream, and the cable modem upstream. Upstream, data does not travel in MPEG-TS packets. Rather, IP packets are placed into a Packet Streaming Queue service flow, which is then chunked and transmitted to the CMTS over a wideband channel that is dynamically allocated to different upstream QAMs as resources are available.

The Packet Streaming Queue is a construct of the Packet Streaming Protocol, a new concept also introduced by Chapman. Packets may be sorted into queues according to quality of service (QoS) level or other policies and travel on service flows dedicated to the different service requirements; the CMTS manages QoS at the cable modem, as in the current DOCSIS release, and prevents head-of-line blocking (where a higher-priority packet might get stuck behind a lower priority one). This is the major difference between Wideband DOCSIS and the earlier versions of the

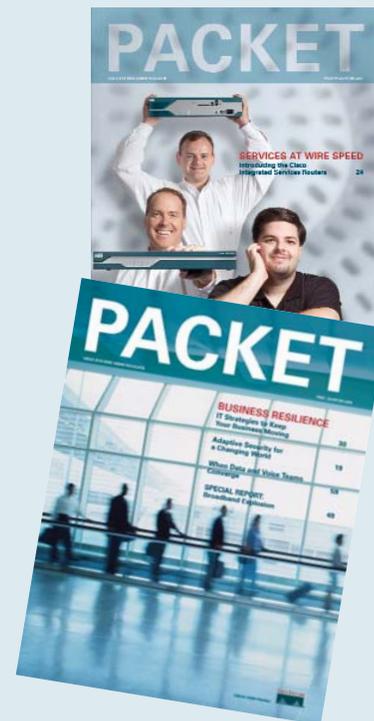
New!

Packet Magazine Digital Edition

The digital edition of *Packet* will be delivered directly to your PC every quarter.

- Read online or download to read anytime, anywhere
- Click on live “Further Reading” links and e-mail addresses
- Print individual articles or the entire issue (.pdf format)
- E-mail articles to colleagues
- Keyword search the entire magazine

Check it out at www.cisco.com/packet/digital



standard. Otherwise, transport adheres to DOCSIS practices for signaling between the cable modem and CMTS regarding bandwidth—launching requests from the cable modem to the CMTS and receiving allocation grants—and is compatible with DOCSIS 2.0 concatenation and fragmentation.

Putting Wideband DOCSIS into Your Network

Installing Wideband DOCSIS in your network is largely a matter of installing new WAN interface cards in your CMTS. These cards will implement the MAC and framing tasks, initially managing up to 24 or 48 upstream and downstream QAM carriers. External QAMs couple the signal onto the HFC. In addition to the MPEG-TS packets, they can also support DOCSIS-based IP services. Both upstream and downstream channels are inherently highly available. Downstream, if one of the QAMs in the wideband channel is down, the Wideband DOCSIS protocol simply doesn't stripe across it. If more availability is needed, the protocol can invoke an RF switch and a redundant QAM. Upstream, if one of the line cards bearing a service flow fails, the wideband channel can be dynamically reconfigured around it.

DOCSIS supports a variety of load-balancing features through Dynamic Channel Change (also invented by Chapman), which was devised primarily for voice traffic. But these techniques work best when used with a group of two to four channels. Wideband DOCSIS, which creates one large channel, better serves large numbers of QAMs and bandwidth-hungry traffic such as video and gaming.

Upstream and downstream wideband channels can be dynamically configured, making the new protocol especially responsive to the customers' need for short-term high bandwidth. "Cable operators can offer a 'turbo button' subscribers can use when gaming or doing peer-to-peer file transfers," says Schroth of the Yankee Group.

Wideband DOCSIS is en route to becoming part of DOCSIS 3.0. This is partially in response to the many cable operators who would like to evolve to a wideband service within the DOCSIS framework to reuse their current DOCSIS infrastructure, mix wideband and traditional services on common downstreams during the transition to wideband, and save operating costs by avoiding rewiring and moving customers to new systems.

Moreover, DOCSIS offers very definite benefits, says Andy Page, product manager in Cisco's Broadband Edge and Midrange Routing Business Unit. "Wideband DOCSIS leverages excellent features in provisioning, billing, security, and other areas. For example, the DOCSIS protocol is very hard to hack and makes stealing service very difficult. Cable providers can choose the billing paradigm—flat rate, time-based, or volume-based. Its provisioning is much more

Cable Operator Trials Wideband Protocol in Japan

Himawari Network, Inc. is testing Wideband Protocol for DOCSIS at the Toyota Dream Home at Aichi prefecture in Japan. Based on the Cisco uBR10012 CMTS platform, the trial showcases the ability to converge video and data traffic onto a single IP-based, high-speed service offering. Himawari will use the technology in parallel with existing modem deployments to provide a migration path to additional high-speed service offerings such as video on demand and online gaming. For more, see cisco.com/packet/172_5c1.

streamlined than DSL, and it makes offering different flavors easy, which helps providers differentiate their offerings and tailor them to subscribers."

Says Page, "Cisco plans to offer the technology to the industry via DOCSIS 3.0 rather than locking in the intellectual property, as part of our philosophy of open systems." Wideband DOCSIS, he adds, "is the logical migration path for cable operators to offer all services over a common IP infrastructure." US trials will take place in the second half of this year, and products should be available in the first half of 2006.

Optimizing Revenue per User

Mattson cites a DFC study from July 2005 that projects worldwide gaming revenues will increase from US\$1.96 billion in 2003 to \$5.2 billion in 2006 to \$9.8 billion in 2009. DFC also predicts that customers for on-demand movies over the Internet will increase from under 3 percent of US households in 2005 to upward of 16 percent by 2008. A recent MDR/Instat survey reveals that 64.6 percent of US homes are now sharing files via broadband; 43.9 percent view pictures; 42.1 percent listen to music; 30.6 percent watch videos; and 29.9 percent do some IP telephony. With Wideband DOCSIS, says Mattson, "You can let video-Napster happen without losing sleep, and in fact you can profit handsomely from it."

"The driver for all of this is optimizing the average revenue per user," Chapman notes. "Having a highly adaptive, cost-effective architecture that accommodates changing traffic patterns, services, and customer needs is immensely valuable. The cable industry has an obvious advantage in DOCSIS, which has a historic focus on service bundling and compelling content and which can now standardize a very wide pipe."

He proposes a five-year goal of at least 1-Gbit/s downstream data capacity and 100-Mbit/s upstream capacity. Are you up to the challenge? ■

Who's Afraid of DUAL-3-SIA?

Cisco IOS Software enhancements improve EIGRP active route processing.

By Russ White

Cisco's Enhanced Interior Gateway Routing Protocol (EIGRP) is widely deployed in large-scale networks, particularly in enterprise financial and retail organizations. EIGRP is also well known for something else—Stuck in Active routes (SIAs)—and yelling “Stuck in Active” in a network operations center is likely to gain a reaction similar to the one you would get if you yelled “Fire” in a crowded theater. When an EIGRP route is in SIA state the DUAL-3-SIA error message occurs.

To reduce the problems with SIAs, Cisco has changed how EIGRP handles the active route process in Cisco IOS Software Release 12.1(5). With the help of these changes in the EIGRP code, you can greatly reduce the scope and number of SIAs, although they still will not be completely eradicated.

This article examines EIGRP's active process and discusses the recent modifications that almost eliminate the situations in which you will see SIAs in your network.

Network Example

The figure on this page illustrates an example of a small network running EIGRP as a routing protocol.

In the simple network in the figure, Router B prefers the path through Router C to reach the IP address 10.1.1.0/24. Router A receives two routes to 10.1.1.0/24, one through Router B and one through Router C. Assume that Router A chooses the path through Router B as its best path, and the path through Router C is not marked as a loop-free path because of the metric differentials in the two paths.

If the link between Routers B and C fails for some reason the following sequence of events occurs:

- Router B examines its local EIGRP topology table for other loop-free paths toward 10.1.1.0/24.
- Failing to find any alternate loop-free paths, Router B queries each of its remaining EIGRP neighbors to determine if it can find a new loop-free path to 10.1.1.0/24. At this point, Router B sets a three-minute active timer and sends a query to Router A, asking if another path to 10.1.1.0/24 exists.

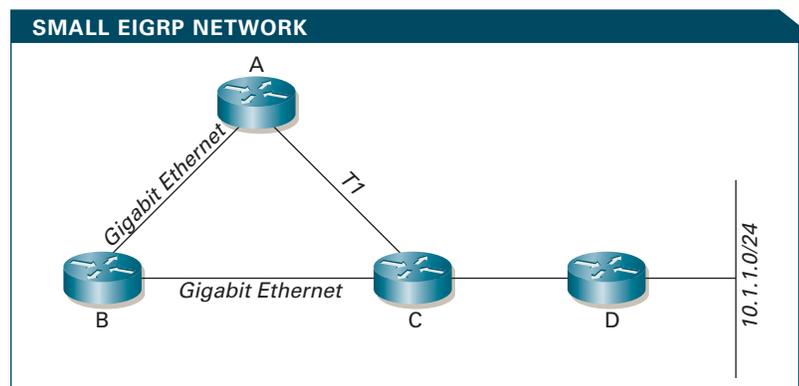
- Router A, which used Router B as its best path, examines its local topology table and finds it has no other loop-free path. Router A sets a three-minute active timer and sends a query to Router C.

At this point, the route to 10.1.1.0/24 is active at Routers A and B. Usually, Router C examines its local routing table, determines it has a path toward 10.1.1.0/24 that is still loop free, and replies to Router A with this information. Router A then installs the route back into its local routing table, takes the route out of active state, and stops the three-minute active timer. Router A, in turn, replies to Router B and Router B follows the same steps.

But what happens if for some reason Router C never responds to Router A? The A to C link could be very dirty, causing most of the packets on the link to be lost, although enough packets make it through for the EIGRP neighbor relationship to stay up, for instance. The entire time that Router C is trying to transmit its reply about an alternate route to Router A, Router B's active timer is still running. Eventually, Router B's active timer will expire and Router B will reset its neighbor adjacency with Router A.

But wait—the problem is not between Routers A and B. The active process, while bringing the network

BEST PATH In the active state, a router sends queries out to its neighbors requesting a path to the lost route.



back into a known state from which to work (the active timer essentially limits the amount of time the network can take to converge), has also caused a network failure where none previously existed: between Routers A and B.

New Active Process

The enhancements to EIGRP in Cisco IOS Software Release 12.1(5) allow the active process to respond to this correctly, by pushing the neighbor reset down to where the actual problem exists. Instead of setting a three-minute timer when a route is marked active, the router sets a one-minute timer. When the active timer expires, instead of resetting the neighbor adjacency, the router sends another query to make certain that the neighbor it is waiting for is still working to resolve the active route.

To illustrate, let's walk through the process using the same network example as before.

- The link between Routers B and C fails. Router B examines its local topology table and discovers it has no alternate loop-free paths to 10.1.1.0/24. Router B marks the route active, starts a one-minute active timer, and sends a query to Router A.
- Router A receives this query, examines its local topology table, and discovers it has no alternate loop-free paths to 10.1.1.0/24. Router A marks the route active, sets a one-minute timer, and sends a query to Router C.
- Router C acknowledges this query but fails to respond to it.
- When Router B's one-minute timer expires, it sends another query to Router A. Router A acknowledges this query, so Router B resets its one-minute active timer.
- Router A continues to wait for Router C to answer the query. Its one-minute timer expires, so it sends another query to Router C.



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 can be reached at ruwhite@cisco.com.

- Router C's acknowledgement to this second query does not get through, so Router A resets its neighbor relationship with Router C.
- Router A, on resetting its neighbor relationship with Router C, examines its local topology table and finds that Router B is waiting for a reply for an earlier query. Router A notes that it is not waiting for any of its neighbors for a reply to this query and it has no alternate loop-free path, so it responds to Router B.

Routers A and C now reset their neighbor relationship, rather than Routers A and B, putting the symptom (a neighbor reset) where the problem is.

Since the change in EIGRP's active processing, the number of Technical Assistance Center (TAC) cases involving EIGRP SIAs has plummeted, along with the number of network managers calling for development support to deal with networks with large numbers of SIAs. EIGRP's new active process means no longer does anyone need to fear SIAs. ■

FURTHER READING

- Cisco EIGRP
cisco.com/packet/172_5b1
- Cisco IOS Software
cisco.com/packet/172_5b2
- Cisco IP Routing
cisco.com/packet/172_5b3
- Cisco Internet Protocol Journal
cisco.com/ipj



A FEW SHORT YEARS AGO, network security was built on standalone products at the physical perimeter of a network, where the LAN met the WAN and corporate networks hooked up to the Internet. Operating system patching and continuous antivirus software updates rounded out the typical corporate security strategy.

However, the concept of a definable network boundary is evaporating. User devices often connect to multiple networks, rendering the perimeter a moving target. Communications among customer and partner extranets is common, for example. The productivity gains afforded by wireless, mobile, and remote-access networks are also fueling the multi-network connectivity phenomenon.

In Self Defense

Network security grows adaptive—reaching inside Web applications and excising attacks at their source.

The security challenge is that user laptops link to other networks and the Internet from home offices, public hotspots, and hotel rooms, for example, and pick up an infection. Then, a user might return to the office and reconnect directly to the corporate network via an Ethernet port or by associating with a wireless LAN access point, inadvertently passing along the bad code. Meanwhile, there is a rapidly shrinking window of time between when that network anomaly arrives and propagates across the corporate network to cause serious consequences. By the time networking personnel detect a virus, worm, Trojan horse, or other unwelcome intruder and attempt remediation, it's often too late to avoid network downtime and losses in productivity or sales.

SECURITY

“This is why security has evolved into a strategic systems issue,” says Kevin Flynn, security products and systems group manager in Cisco’s Products and Technology Marketing Group. “Security has now become indistinguishable from other IT and network operations.”

Networks have grown too complex for a single mechanism to reliably keep them secure. Modern networks require distributed security capabilities, networkwide awareness of the context of endpoint credentials as host behavior and status change, and authentication mechanisms ensuring that those credentials can be trusted.

Protecting Every Packet, Every Flow

Cisco’s Self-Defending Network strategy, which comprises three phases, has been rapidly gaining new components to fulfill these requirements. The third phase, called *Adaptive Threat Defense*, for example, got underway with several important product announcements in March of this year, many of which are described in this article. Adaptive Threat Defense aims to protect every packet and every packet flow on a network.

The first phase of the Self-Defending Network strategy involves integrating security capabilities directly into network elements, such as routers, switches, wireless access points, and standalone network appliances. The second phase, which includes the industry-wide Cisco Network Admission Control (NAC) effort, involves security-enabled network elements communicating with one another in a collaborative manner, such as an intrusion prevention system (IPS) telling an access control list (ACL) to deny access to a connection. It also extends the security capabilities to the user endpoint devices that connect to other networks and might infect the corporate network.

Why has it now grown necessary to protect every packet and flow? One reason is that, increasingly, security attacks are being introduced from within Web-enabled applications, which use HTTP’s port 80 to communicate.

“Web applications, while empowering users, open the door to application abuse as traffic traverses multiple networks and potentially picks up virulent code,” says Jayshree Ullal, senior vice president of Cisco’s Security Technology Group.

CISCO SELF-DEFENDING NETWORK STRATEGY		
PHASE 1: Integrated Security Launch 2000	PHASE 2: Collaborative Security Launch 2003	PHASE 3: Adaptive Threat Defense Launch March 2005
Security capabilities are embedded into network elements such as switches and routers	Embedded security capabilities are linked across the network and extended to user endpoints	The network gains the ability to protect every packet and every flow and eradicate attacks at their source

FIGURE 1 Security is an infinite, ever evolving process that encompasses all three phases of the Cisco Self-Defending Network strategy.

A slew of new Cisco products protect against application abuse, identify and thwart intrusions at Layer 7, and even eliminate the sources of attacks. To do so, they leverage capabilities such as deep-packet inspection, networkwide event correlation, context-based policies, and policy auditing. To combat application abuse, for example, application inspection firewalls have been added to the Cisco PIX 7.0 Firewall appliance and to the Cisco IOS Software firewall in Cisco IOS Software Release 12.3(14)T, as well as to a new, next-generation appliance that combines several Cisco-leading security technologies: the Cisco ASA 5500 Series Adaptive Security Appliance.

Application inspection firewalls now check port compliance for HTTP (port 80) and e-mail (port 25). In other words, the engines inspect traffic on these Layer 4 ports to make sure that it is, indeed, the type of traffic intended for that port. “This helps network operators control port misuse by rogue applications that hide traffic inside Web and e-mail applications to avoid detection,” says Ullal.

Overarching Security Monitoring and Response System

A pivotal advancement in Cisco’s ability to protect every packet and every flow is its recent introduction of a networkwide security management system called the Cisco Security Monitoring Analysis and Response System (CS-MARS).

“CS-MARS basically enables, for the first time, the comprehensive, centralized management of a Cisco Defense-in-Depth network,” says Greg Simmons, customer solution manager in the Network Management Technology Group at Cisco.

The system, a fruit of Cisco’s recent acquisition of Protego Networks, Inc., collects security event data from every network element configuration, host log, and TCP and UDP session (packet flow) in real time. It then correlates them all with one another and with corporate security policies to determine whether each event is legitimate.

“You could shut down one offending laptop centrally using CS-MARS,” explains Simmons.

All security capabilities integrated into Cisco routers, switches, firewalls, appliances, and Cisco IOS Software—including many new features described herein—continue to act as “soldiers,” each defending against individual attacks on a particular endpoint, explains Timothy Smith, technical marketing engineer in Cisco’s Network Management Technology Group. These devices, in addition to some non-Cisco devices, continually feed event data to CS-MARS.

The CS-MARS system, by contrast, behaves as the “general” by overseeing the entire security battlefield. It cross-relates all the security activity, creating and unleashing a top-down combat strategy. All devices in a given TCP or UDP session between any two hosts, both Cisco and non-Cisco devices, report data up to CS-MARS, which can identify every device in the path of that session.

Having this information allows network managers to identify an attack, alert the user, and shut down the source of the attack, says Smith. CS-MARS will send the network administrator the appropriate command to execute an action to excise the problem from the network at its source. By contrast, the job of the soldiers—the various individual security products—is to act on the immediate

effect of the attack at its point of potential impact at the various layers of the network.

Additionally, by knowing the full path an attacker has traversed, CS-MARS can protect the network in other ways. For example, it might anticipate that the CPU or memory in a given router, access switch, or application server might be on the verge of maxing out because of a flood of packets.

“CS-MARS would then take action, delivering a command, say, to leverage another network element in the path of the failing device to save it and prevent denial of service,” says Smith.

CS-MARS also serves as a Cisco NetFlow collection engine. NetFlow is a Cisco IOS Software capability for counting the packets in individual traffic flows, helping network operators quickly spot traffic patterns and account for usage of network resources. “CS-MARS might use NetFlow counts to determine that a port is experiencing a sharp rise in traffic, which is one variable,” explains Smith. Then CS-MARS will determine, based on other data—such as intrusion protection signatures and firewall logs—if an attack is occurring on a particular device, for example, in Building 6, Floor 2, Conference Room B (displayed on the CS-MARS console in these easily understandable parameters).

“Not only do I learn which devices are causing the anomalous traffic, CS-MARS gives me the command to shut those devices down,” Smith says.

At the high end, the CS-MARS family includes the CS-MARS 200, which processes 10,000 events per second or 300,000 NetFlows per second, and scales down to the CS-MARS 20, which processes 500 events per second. From a configuration standpoint, a CS-MARS 200 could serve a single large site; alternatively, multiple smaller devices could be distributed at remote sites and report up to a CS-MARS Global Controller, explains Smith.

Multifunction Security Appliance

A new “soldier” to the Cisco security portfolio is an integrated device that addresses the Adaptive Threat Defense landscape by combining multiple Cisco security technologies into a single, extensible appliance:

- Firewall technology
- VPN capabilities—both IP Security (IPSec) and Secure Sockets Layer (SSL) technologies
- Inline IPS technology

The Cisco ASA 5500 Series includes the Cisco ASA 5540 at the high end (650 Mbit/s of firewall throughput, even as new services are added), the Cisco ASA 5520 for mid-range performance, and the Cisco ASA 5510 at the low end. Each platform provides application firewall services and flexible IPSec and SSL VPN connectivity. The optional Advanced Inspection and Prevention Security Services Module (AIP-SSM) supports IPS and network-based antivirus, worm, and spyware protection, according to Michael Jones, product line manager at Cisco.

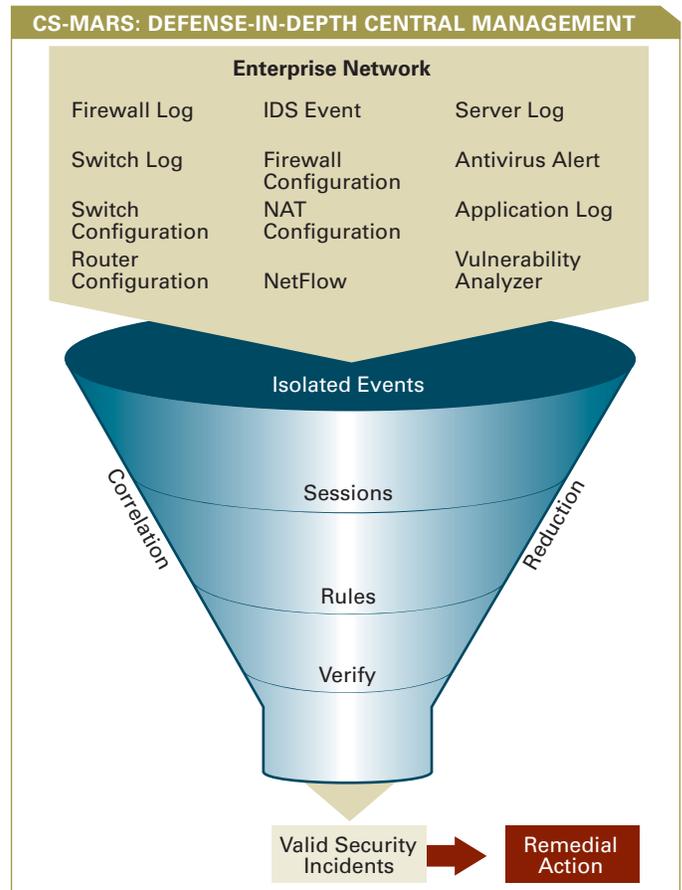


FIGURE 2 By aggregating, correlating, and analyzing security data from devices network-wide, the CS-MARS security and management system helps users readily identify and eliminate valid network attacks at their source.

A unique, extensible services architecture, called Adaptive Identification and Mitigation (AIM), is at the heart of the Cisco ASA 5500 Series design. This architecture allows network operators to apply specific security and network services on a per traffic flow basis, providing extensive policy control. Furthermore, the AIM services architecture enables the integration of future threat identification and mitigation services—further extending investment protection and allowing businesses to defend their networks against new threats as they arise.

The Cisco ASA 5500 Series’ ability to satisfy a broad range of security roles yields reduced deployment and operations costs. “For example, you can standardize on this single appliance for many of your security needs, including firewalling, VPN connectivity, and intrusion prevention,” says Jones. In addition, a unified, Web-based user interface for all ASA functions decreases management complexity and lowers overall operational costs.

“For folks designing new networks, the adaptive security appliances are a great solution for putting all services in one location,” says Jones, “or for refreshing an existing site with additional services.”

VRF-Aware Firewalling

As mentioned, the new base firewall code is also included in Cisco IOS Software Release 12.3(14)T. This step has rendered the Cisco IOS Firewall virtual routing and forwarding (VRF)-aware.

Continued on page 32

SECURITY

Self-Defending Network, Continued from page 29

In other words, a router that is running multiple routing instances (functioning, in effect, as multiple routers within a single chassis), can now also run multiple Cisco IOS firewalls within that chassis to match, explains Tom Guerrette, product manager in Cisco's IOS and Router Security Marketing Group.

The new software release applies Cisco IOS Firewall functionality to each VRF interface, allowing customers to configure per-VRF firewalls. The firewall inspects IP packets that are sent and received within a VRF. A few noteworthy capabilities about the VRF-aware IOS firewall:

- It supports overlapping IP address space, thereby allowing traffic from nonintersecting VRFs to have the same IP address.
- It supports per-VRF (rather than global) firewall command parameters and denial-of-service (DoS) parameters. In the case of a service provider managed service, for example, the VRF-aware firewall can run as multiple instances allocated to various VPN customers.
- It performs per-VRF URL filtering.
- The VRF-specific syslog messages it generates can be seen only by a particular VPN, allowing network administrators to manage the firewall.
- It supports the ability to limit the number of firewall sessions per VRF.

The same capabilities apply to the Cisco PIX 7.0 Firewall and Cisco Adaptive Security Appliances, as well.

Checking Conformance to Policy, Best Practices

The Cisco Security Auditor is a new component of Cisco's security management suite that enables customers to cost effectively audit their network security infrastructure postures. With the tool, they

can automatically check for conformance to their corporate security policies and, simultaneously, check against multiple industry best practices such as those set by Cisco, the US National Security Agency (NSA), and the Center for Internet Security (CIS).

"Having the ability to measure, compare, and report on the security status of a dynamic network helps efficiently manage networking security risks and meet government security mandates," says Flynn.

The software allows automated auditing of thousands of devices, significantly reducing the time required to audit a network. The Cisco Security Auditor also provides easily understood security improvement recommendations, including those required to correct deviations from security policy, which would have previously required time consuming manual analysis and the use of scarce experienced staff.

Cisco NAC Initiative

The industrywide Cisco NAC initiative embeds endpoint scanning technology directly into Cisco network elements, such as switches, routers, and, now, Cisco VPN 3000 Series concentrators, from partners such as antivirus developers McAfee, Symantec, and Trend Micro. Combined with Cisco Trust Agent software—which uses IEEE 802.1X and RADIUS authentication technology—it resides on client endpoints to keep all connections to the corporate network free of infection.

NAC APIs have been opened up to independent software vendors (ISVs), who have joined the NAC initiative for automatically having their Cisco network infrastructures check for patch management or to handle authentication or software management via their own software management systems. Among the ISVs shipping NAC-enabled products are IBM (the IBM Tivoli Security and Identity Management Product Suite); Computer Associates (eTrust AntiVirus and eTrust PestPatrol), and InfoExpress (CyberGatekeeper Server 3.1 & CyberGatekeeper Policy Manager 3.1).

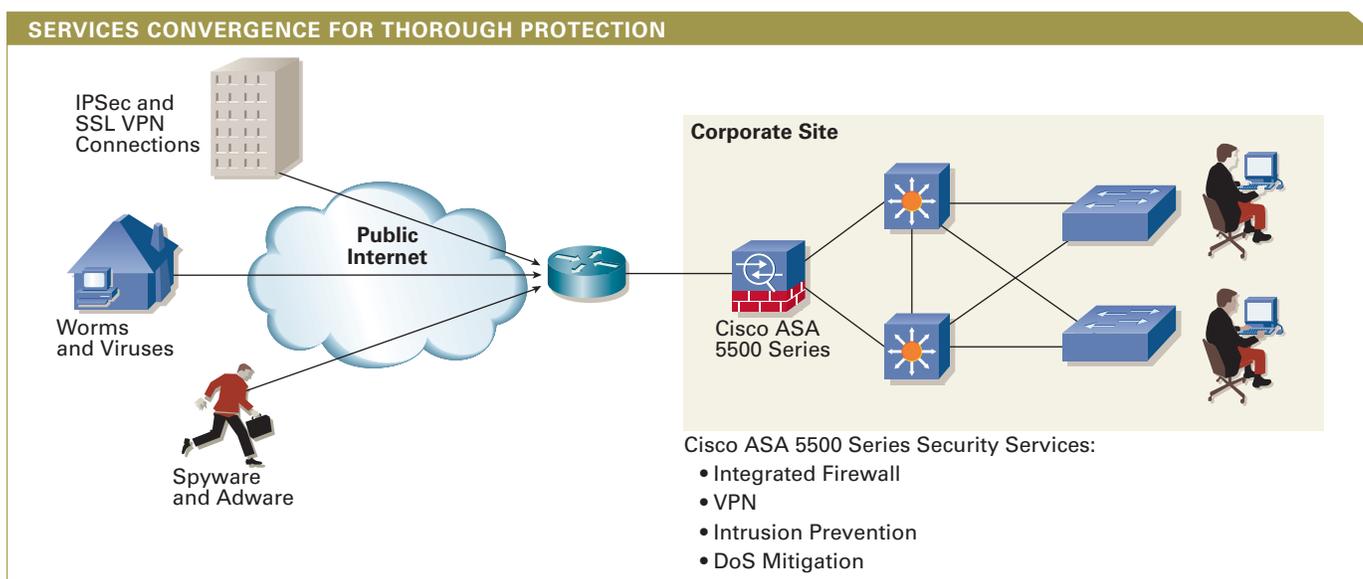


FIGURE 3 The Cisco Adaptive Services Appliances' integration among security functions allows different per-user policies based on IPSec or SSL credentials.

APPLIANCES VERSUS INTEGRATED SOFTWARE

The form factor in which you choose to deploy security likely will depend on your organizational structure, the state of your existing equipment, and your technology preference and philosophy. Here are some considerations:

Are your network operations and security operations handled separately within your organization? If so, it might be simpler to have the separate groups administer policies using separate devices.

Are you in an upgrade cycle with your router software and resources? If so, now may be the time to add security features. If not, it might be simpler and less expensive to go the appliance route.

Similarly, if you have maxed out the number of slots in your router, yet aren't considering an upgrade for a while, you might consider an appliance.

If you are considering products for a branch office, you might prefer integrating your capabilities in an all-in-one software-based router/switch approach to reduce capital expenditures (CapEx) and to conserve real estate. You might prefer to put your resources instead into a discrete appliance at your headquarters and other very large sites.

Cisco Clean Access: NAC in Appliance Form

Cisco now also offers Cisco Clean Access, an all-in-one NAC appliance option, resulting from its recent acquisition of Perfigo, Inc. Falling beneath the NAC umbrella, Clean Access is a self-contained solution for conducting posture assessment and automatic checks for the latest antivirus updates and critical OS patches. In recent deployments of Cisco Clean Access, the number of computers requiring intervention by IT staff due to viruses and worms reduced dramatically. At Arizona State University, for example, the number of computers requiring IT staff intervention plunged from 6000 to 50 after Clean Access was deployed.

Clean Access doesn't require 802.1X authentication or client software. Client software is available, however, and deeper scanning capabilities and simplified remediation are possible by downloading the agent, says Irene Sandler, Clean Access marketing manager at Cisco.

In the in-band product, all hosts attempting to gain access to the network traverse the Clean Access server. In the out-of-band configuration, which began shipping in April, Clean Access works together with Cisco switches to provide Layer 2-based quarantining for non-compliant machines, which are then repaired by the Clean Access server. After properly remediated, the now-compliant host is placed back on the network, "out of band," to the Clean Access server.

The out-of-band Clean Access version also allows network administrators to create and enforce policies through a central interface. Policies can be defined on a per-role basis.

"This makes it extremely easy for an administrator to assign a certain level of permission or compliance requirements to employees, for example, while applying a separate level of compliance for guests," Sandler says.

Intrusion Prevention Improvements

Cisco IOS IPS, introduced in Cisco IOS Software Release 12.3(8)T, delivers a new level of inline accuracy to identify and halt more threats businesses face without impacting legitimate traffic. Cisco IOS IPS goes beyond traditional IPS products by using risk-based analysis and real-time correlation to improve prevention accuracy, says Guerrette. The system divides supported signatures into signature micro engines (SMEs). In Cisco IOS Software Release 12.3(14)T, three new SMEs were added that represent "where most of the new attacks seem to be going," according to Guerrette. "As a result, new attacks will be discovered and stopped more rapidly," he says.

Meantime, Cisco Security Agent Version 4.5 host-based IPS software adds compatibility with operating systems outside the US and expands platform support to include RedHat, Inc.-based Linux servers and desktops and Windows clusters. Management scalability for large enterprises has been increased to 100,000 agents from a single Cisco Security Agent Management Center. Advanced integration with NAC allows policies to be dynamically changed based on the NAC security posture, logged-on user, or location of the end device.

21st Century Security

With the addition of the Cisco Adaptive Threat Defense phase to the Self-Defending Network strategy, multiple layers of built-in network security now reach from an Ethernet port to the interior of a Web application. With this phase comes a much improved security paradigm for the 21st century.

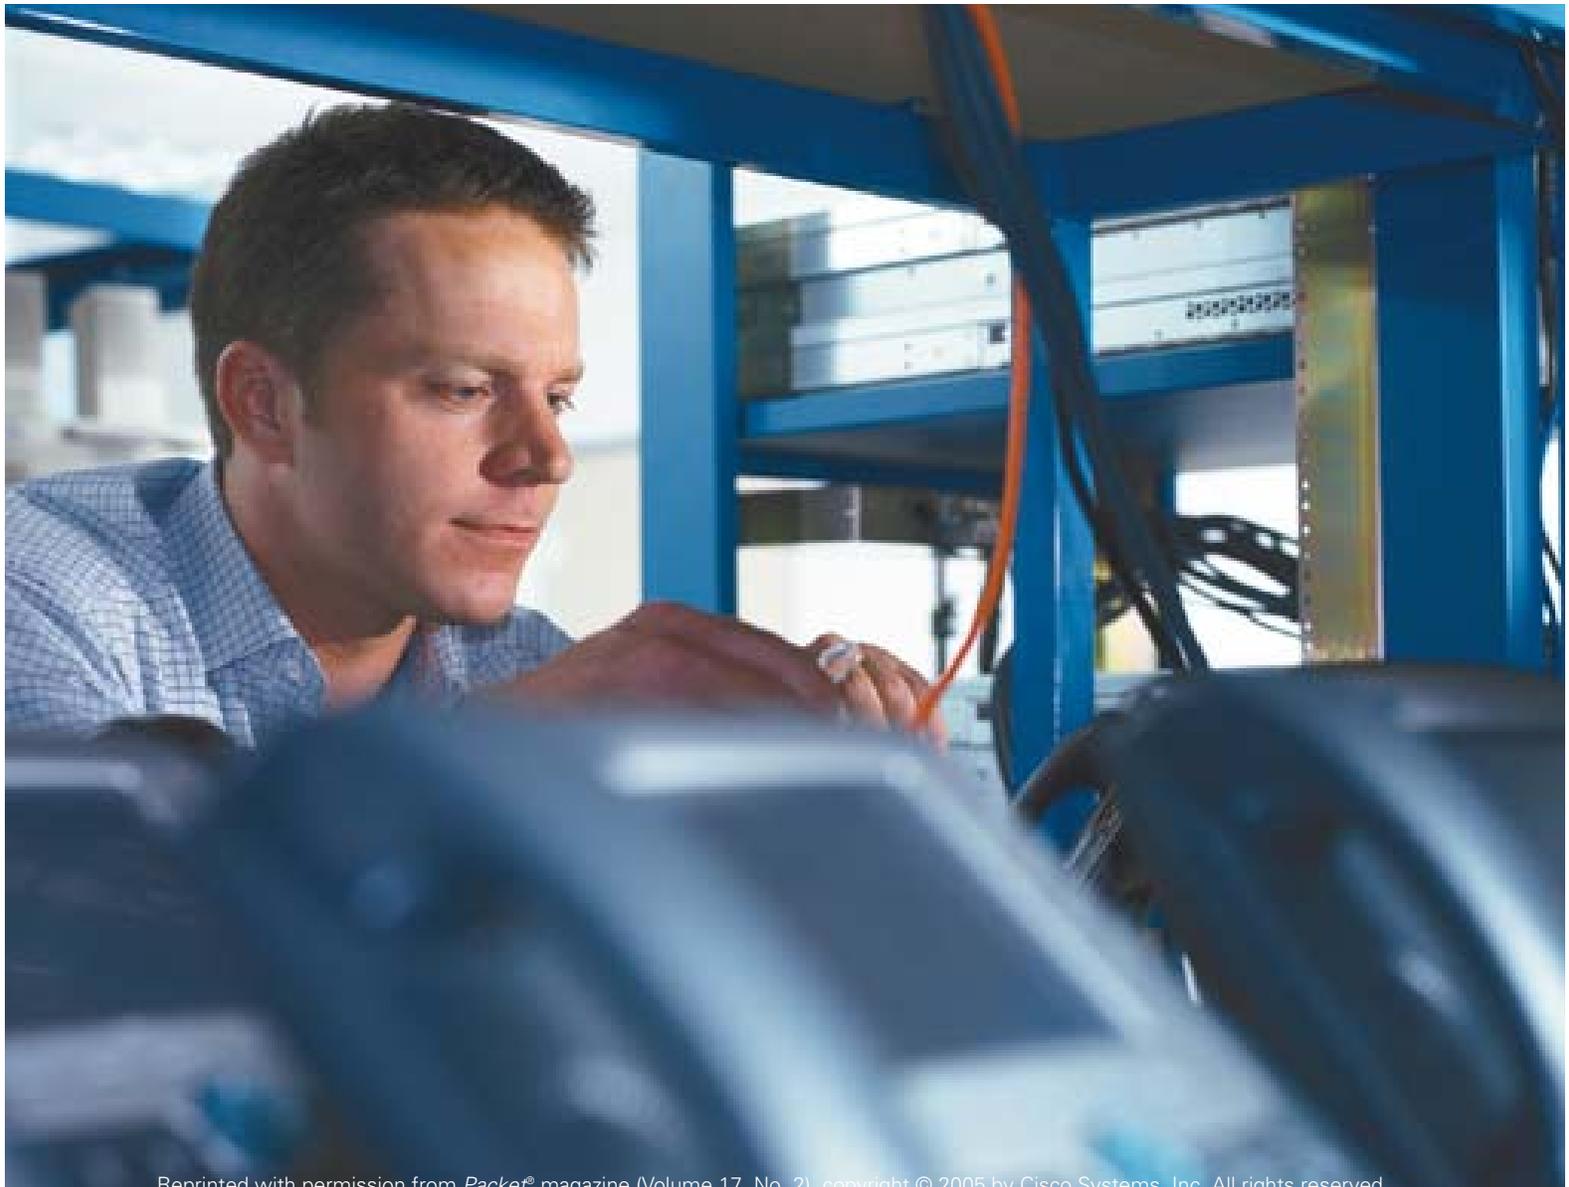
"Protection is no longer dependent on just antivirus software and signatures," says Ullal. "It is built on behavioral and trusted clients that work closely and collaboratively with the network."

With the disappearance of a definable network perimeter and security threats coming at networks from every angle, point products alone no longer are an adequate defense. An integrated and proactive multilayered system makes the Self-Defending Network—now a requirement to ward off the consequences of rapid-propagating attacks—possible. And security will be an ongoing process that will likely be forever evolving as networks, applications, and threats themselves change. ■

FURTHER READING

- Cisco Security Product Reference Sheet
cisco.com/packet/172_6a1
- Cisco Self-Defending Network
cisco.com/go/security

Voice Keeping Confidential



The Key Risks of VoIP Joining the Network and How to Mitigate Them

IP VOICE PACKETS are rapidly spilling over from converged WAN services and onto enterprise LANs. In the US alone, IP phone extensions are expected to more than triple from about 14 million in 2004 to about 45 million in 2008, according to the Insight Research Corporation in Boonton, New Jersey.

As installations grow, so do user worries about security ramifications. Indeed, businesses will certainly want to continue taking precautions to prevent “stolen” long-distance phone service and to ensure that access to voice resources, such as voice mail, remains private.

The integrated, systems approach used to secure IP networks today already protects a large number of business-critical applications, and there is a great deal of experience protecting separate circuit-switched PBX systems in the traditional voice community. So the remaining question is: How do we best leverage the strengths of both groups to deliver advanced IP communications voice services that are as secure, or more secure, than the legacy systems they are replacing?

Identifying the Threats

Let's start with identifying the threats to the voice network components and service, then discuss the ways to mitigate them. Some risks threaten the availability and quality of the voice service itself and the privacy of voice conversations. Others come in the form of using the voice network as a catalyst to penetrate the data network, which is where the most valuable assets for theft generally reside.

The most common threats in a premises-based voice-over-IP (VoIP) environment are in the following primary categories:

- Theft of service, or toll fraud
- Unauthorized access to voice resources, such as voice mail
- Compromise of the data network using voice devices and infrastructure as an entry point
- Downtime/denial of service (DoS) and associated loss of productivity
- Invasion of call privacy (eavesdropping)

VOICE SECURITY

As in the case of the data network, most susceptibilities can be thwarted by applying standard network design and configuration best practices to voice network elements in a multilayered, or defense-in-depth, manner.

The first two threats are handled in an IP environment in the same way they would be in a circuit-switched environment. Preventing hackers from dialing into your PBX and back out onto a long-distance network is usually a matter of disallowing extension transfers to outbound ports, for example.

Keeping access to voice resources restricted to intended recipients is important so that outsiders don't gain insight as to who is communicating with whom and learn about potential business partnerships, acquisitions and mergers, confidential R&D information, etc. Strong access codes, passwords, and encryption assist here.

The remaining threats—data theft (a voice-assisted network break-in), denial or degradation of service, and eavesdropping are more closely related to voice and data packets sharing a common infrastructure. Among the defenses against these remaining vulnerabilities:

- Voice virtual LANs (VLANs), or logical network segmentation
- Disablement of certain features on certain phones, depending on their location and use
- Use of application-layer firewalls and access control lists (ACLs)
- Use of resource-limiting to tame the impact of DoS attacks
- Media and link encryption

Let's take a look at how to apply each.

Segmentation Lowers Risk

Assigning voice traffic to specific VLANs to logically segment voice and data traffic is an industrywide accepted best practice. As much as possible, devices identified as

NUMBER OF US IP PBX PHONE EXTENSIONS INSTALLED (IN MILLIONS)

2004	2005	2006*	2007*	2008*
13.6	19.4	26.6	35.3	45.2

Source: Insight Research Corp., Boonton, N.J.

*Projected

EXTENDING ITSELF As IP telephony grows, so will security worries.

voice devices should be restricted to dedicated voice VLANs. And, as such, they could communicate only with other voice resources. More importantly, explains Roger Farnsworth, senior systems marketing manager in Cisco's IP Communications Security Group, the voice traffic is kept away from the general data network where it might more easily be intercepted or tampered with.

"VLANs help prevent rogue devices from plugging into the network. VLANs have specific membership criteria, and devices joining a VLAN must meet the criteria in order to authenticate to them," Farnsworth explains. "Nontelephony devices would ideally be kicked off a VLAN configured to carry voice traffic only."

From a DoS perspective, segmenting voice into their own logical VLANs will limit the probability of an attack, says Greg Moore, a technical marketing engineer in Cisco's IP Communications Security Group. Hackers tend to write viruses for the most popular software, which, to date, have been general-purpose data server operating systems.

"The scope of responsibility for call-processing components such as a Cisco CallManager server is more limited than that of a general-purpose data server," Moore notes. "So keeping voice logically separate reduces voice's susceptibility to these intrusions."

Voice VLANs, by the way, carry a strong quality-of-service (QoS) side benefit; they can be prioritized over data VLANs in a switch/router's priority queue to consistently reduce VoIP latency.

When building VLANs, Farnsworth recommends not using the default VLAN address, so that the VLAN number isn't easy for a hacker to guess. "And use non-contiguous-numbered VLANs for the same reason," he adds.

Authentication and Security of IP Phones

An important component of a secure voice network is that IP phones—handsets and softphones—are authenticated as legitimate participants in the IP telephony network by registering with the call server (Cisco CallManager or CallManager Express).

"Management of identity is a crucial component of voice security," says Farnsworth.

Moore also recommends some physical best practices for IP phones. One is disabling phone ports to which downstream PCs connect in environments where people other than employees could gain network access, such as in a lobby, cafeteria, loading dock, and guard shack. Similarly, he advises disabling Web access to IP phones in such public areas.

"You can potentially learn a whole lot about a network by compromising an unprotected phone," he explains. "The XML applications on the phone use the same HTTP port 80 that Web applications do."

Though voice and data might be logically separated, port 80 might be open to IP phones, allowing an intruder from the Web onto the phone and, from there, onto the voice network.

"ACLs should be written so that only XML servers can get to phones, and phone-to-Web access over port 80 should be disabled," he advises. "You can configure all this using Cisco CallManager."

DoS and Resource-Limiting

Using rate limiting and QoS tools can also help secure voice networks against downtime caused by DoS attacks and other types of packet flooding, says Farnsworth. Limiting the amount of processor and link resources that can be consumed by a given protocol, for example, curbs the impact that a DoS attack can have, he explains.

Moore adds that LAN Microflow Policing, a Cisco QoS capability, handles this function on a per-flow basis, allowing users to limit one IP address, or session, to a certain amount of bandwidth, for example. Doing this would prevent an attack emanating from a particular IP source address from consuming more than the maximum amount bandwidth specified (always leaving some capacity left over for production traffic).

In addition, a suite of features called Catalyst Integrated Security (CIS) foils these attacks before they happen. Among these tools is Dynamic ARP Inspection, which watches ARP requests for contradictions in the binding table at Layer 2, and IP Source Guard, which watches for Layer 3 contradictions. Both features are available in the Catalyst and Cisco IOS Software operating systems. Upon finding any such contradictions, these features can be programmed to drop associated packets or disable the associated port, Moore says.

Access Control and OS Protection

First-generation firewalls looked only at IP address information and matched it with ACLs for a permit/deny decision. However, application-layer firewalls can now inspect port-layer and Layer 7 application information in IP headers to make more informed decisions about whether traffic is indeed legitimate.

Because VoIP signaling protocols can often hop across many ports, for example, application-layer firewalls are necessary to follow the session and open and close ports as needed, rather than leaving a range of ports open and potentially vulnerable to intruders.

Firewalls and associated ACLs, then, must be configured uniquely for VoIP to allow signaling to take place between IP handsets or softphones and Cisco CallManager or CallManager Express call servers (and Real-time Transfer Protocol [RTP] media servers, in the case of streaming applications). Also, certain applications are uniquely aligned with VoIP, such as Cisco's Attendant Console automated call-routing application and Web Dialer click-to-dial application. The Layer 4 session ports used by these applications must be known and configured to grant privilege for them to work through the firewall.

As is part of any network security best practice, operating systems must be kept up to date with patches. Also, host-based intrusion prevention in the form of Cisco Security Agent software and network-based intrusion detection and prevention available in several form factors (see cover story, page 26) help protect the operating systems on servers and the integrity of the network. Appropriately configured versions of Cisco Security Agent are currently supplied for for Cisco CallManager, Cisco Contact Center, and Cisco Unity host operating systems, as well.

Traffic Interception and Media Encryption

Intercepting traffic is not a trivial exercise, says Moore. "If all the other security layers described have been covered, encryption is not a necessity for everyone." However, media and link encryption are available as yet another layer of protection that some organizations require.

Cisco enables encryption end to end, between IP phones, using the Secure Real-Time Protocol (SRTP). Designed specifically for voice packets, SRTP supports the Advanced Encryption Standard (AES) and is an Internet Engineering Task Force (IETF) standard (RFC 3711). Media encryption using SRTP is more bandwidth-efficient than IPSec, an important consideration for latency-sensitive VoIP transmissions.

Cisco supports media encryption on a wide range of Cisco IP phones, including the Cisco IP Phone 7940G, 7960G, and 7970G. These phones also come with support for industry-standard X.509 digital certificates capable of authenticating the

end device for encryption as opposed to relying on manual entry of encryption key data, easing scalability in large installations.

Media encryption is also available on a wide range of Cisco media gateways including the Cisco 1800, 2800, and 3800 Series Integrated Services Routers. In addition to supporting SRTP media encryption, these Cisco gateway products support encryption of call setup information using IPSec.

The privacy protection of encryption can also be applied to voice messages via secure private messaging feature for the Cisco Unity unified messaging system. Messages can be marked private and secure such that only the intended recipient can decrypt and listen to messages.

Rise to the Challenge

IP telephony deployments, now quickly on the rise, expose the enterprise to new but manageable risks. Having an understanding of those risks is the first step toward a successful defense. And because the IP telephony service interacts with the IP data infrastructure, mapping defensive measures to the overall networkwide security framework and strategy is critical. Many of the steps used to protect the data network are imperative for also protecting the voice network; likewise, many of the same risks to the circuit-switched voice network need mitigation in the IP network.

From there, running voice VLANs, configuring application-aware firewalls and ACLs to be voice-aware, disabling certain features on phones in public places, and using encryption for privacy are some of the measures that keep the voice network functioning and prevent it from being a conduit to hacking corporate data resources. ■

FURTHER READING

- Media Authentication and Encryption Using Secure RTP
cisco.com/packet/172_6c1
- "Securing your Network for IP Telephony" White Paper
cisco.com/packet/172_6c2



Eradicating Wireless Intruders

Multilayered RF monitoring and wireless intrusion prevention systems leave no room at the wireless LAN table for uninvited guests.

WIRELESS NETWORKS offer tremendous mobility, flexibility, and productivity improvements to customers worldwide. Unfortunately, from a security standpoint, they also mean that you are about to become a network broadcaster and introduce new security threat vectors into your network. In fact, a whole new breed of wireless-specific intrusion detection and denial-of-service (DoS) attacks are challenging wireless LAN (WLAN) vendors to create new levels of detection and prevention services in their WLAN intrusion detection system (IDS) offerings.

Compounding some of the challenges of broadcasting data over the air is the fact that most wireless-enabled laptops running Windows 2000 or XP, by default, are actively seeking out any and all Wi-Fi capable connections with any access point—regardless of whether that access point is authorized or unauthorized. So the challenge becomes not only ensuring safe “air space,” but also safe and secure connections from your computers.

To help guard against these security breaches through rogue devices and intrusion attempts, Cisco offers several wireless best practices to advise network administrators on how to install intrusion detection and protection safeguards to counteract the security ramifications of this situation.

“We can’t simply create systems that pop up lots of alarms and alerts and call it the ‘bells and whistles’ approach to IDS that we saw on wired networks over the last 10 years. Our WLAN customers are really looking for wireless intrusion detection and now, wireless intrusion prevention and protection capabilities,” says Bruce McMurdo, product marketing manager in Cisco’s Wireless Networking Business Unit.

Now, detection and remedial action take place in a protective model that is both distributed and hierarchical and in which each network layer, from client to data center, plays a role in defending against network threats (see figure). To protect against intrusions, Cisco’s wireless architecture offers the following capabilities:

- Access point authentication
- Disablement of unauthorized access points
- Client containment
- Client policy enforcement
- Location-based intrusion detection

In the Cisco architecture, Cisco clients and Cisco Compatible Extensions clients, in addition to Cisco Aironet Series access points and Cisco lightweight access points, serve as network intrusion prevention system (IPS) sensors. While all WLAN IDS vendors rely on access points for rogue device detection, Cisco is unique in its ability to extend rogue detection to clients—allowing the Cisco solution to provide more comprehensive detective capabilities as these clients move throughout the WLAN environment.

Identifying and Disabling Rogues

Cisco access points, acting as IPS sensors, report the discovery of rogues to Cisco wireless control and management devices. Then, based on the network administrator’s policy, these devices automatically suppress the unauthorized, or rogue, devices that are connected to the network.

Cisco’s solution supports both an integrated and overlay intrusion prevention system. An overlay intrusion prevention system monitors the air space using separate distributed radio sensors. In this case,

“We recommend an integrated intrusion prevention system in most instances—where access points act as sensors while simultaneously supporting client transmissions.”

—Bruce McMurdo, Product Marketing Manager, Cisco Wireless Networking

IT staffs usually must choose between cost and effectiveness: Mapping radio sensors to active WLAN access points at a 1:1 ratio ensures the best network coverage but can be cost prohibitive; using fewer sensors can ease the cost but could result in monitoring coverage holes.

“Having a common infrastructure for both your monitoring and your wireless data network gives enterprises better visibility into their networks,” says McMurdo. “Therefore, we recommend an integrated intrusion prevention system in most instances—where access points act as sensors while simultaneously supporting client transmissions.”

Sensors will discover many rogues. “The key is to correctly identify neighboring networks that are *friendly* and pay vigorous attention only to those actually connected to the corporate network that shouldn’t be connected or access points placed in locations in your environment that shouldn’t be there,” says McMurdo.

With rogue access point suppression, the sensors detect wireless-device information, aggregate it, and pass it up to elements in the network that can correlate it and act

upon it. When a wireless access point is detected on the network, the WLAN intrusion prevention system sends RF management frames that disassociate any clients that connect to it and attempt to trace and shut down the switch port to which the rogue is connected.

With Cisco, customers can deploy an intrusion prevention system using a distributed solution or a lightweight solution. The distributed solution uses Cisco Aironet 1230, 1200, 1130, or 1100 series access points deployed with a Cisco

Wireless LAN Solution Engine (WLSE) Release 2.9 and higher and Cisco Catalyst 6500 Series Ethernet switch with a Wireless LAN Services Module (WLSM). With this solution, all legitimate Cisco and Cisco Compatible wireless client devices, as well as Cisco Aironet access points, gather information about all wireless devices in their immediate vicinity.

The lightweight solution, based on the recently acquired Airespace product portfolio, uses Cisco 1000 Series lightweight access points with a Cisco Wireless LAN Controller and Cisco Wireless Control System (WCS).

Both solutions can be deployed as an IPS where the access point serves the dual purpose of forwarding Layer 2 packets and also acting as a monitoring sensor on the network.

What Are the Threats?

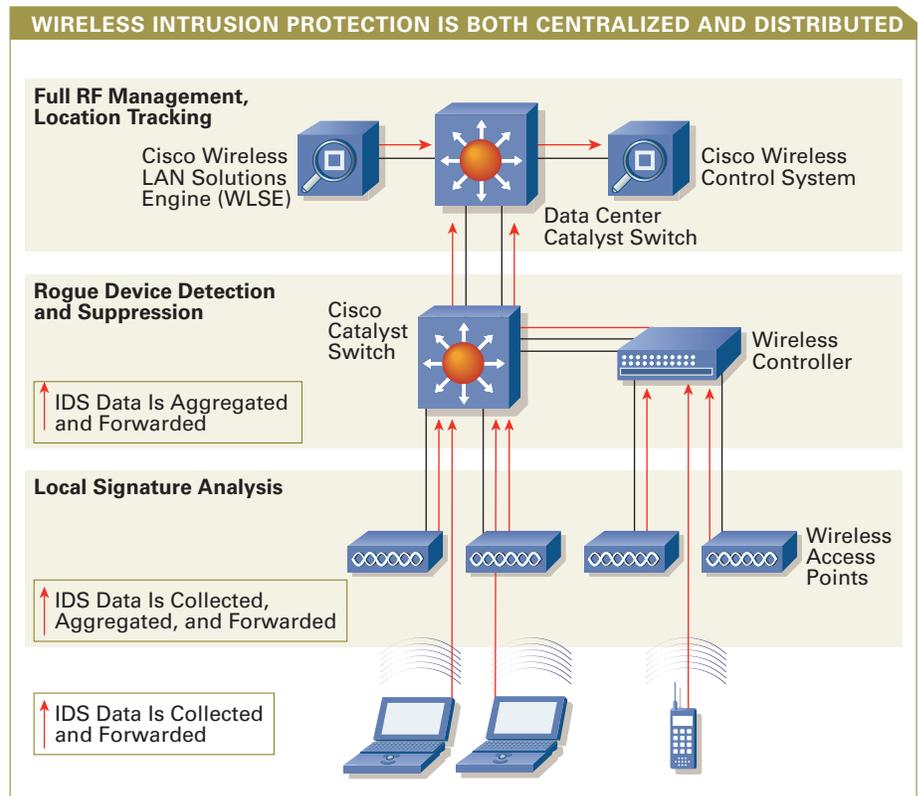
The intrusion-oriented threats of wireless LANs are caused by the difficulty in containing airborne transmissions within physical boundaries of a given organization’s walls and the nature of wireless client devices to automatically connect to the strongest signal they can find.

“When it comes to theft of data and network break-ins, deploying the encryption and authentication measures in 802.11i—

the latest extension to the IEEE 802.11 suite of wireless LAN security standards—is the best protection,” says Jake Woodhams, a technical marketing engineer at Cisco. “However, for thwarting data hijacking and service disruptions enabled by the widespread existence of rogue devices, scanning the airwaves to identify and shut down unauthorized devices is an important practice,” he says.

Specifically, here’s what intrusion detection and prevention seek to mitigate:

- **Impact of rogue infrastructure access points.** These radios, which plug directly into Ethernet switch ports or wireless LAN controller ports, might be maliciously installed by intruders for corporate hacking. More often, though, they are naively deployed by employees for easy wireless access. Because of the self-deploying nature of client devices, clients in a nearby office, parking lot, or coffee shop might connect with these rogues by intent or by happenstance. Because most enterprises currently have open Layer 2 ports on their Ethernet switches, plugging the rogues into the wired LAN infrastructure gives all clients associating with them access to at least some corporate resources.
- **WLAN traffic “attack” signatures.** Volumes of undesirable traffic can make their way onto the wireless transport path, flooding legitimate wireless devices with management frames, authentication requests, and many other types of packets, depending on the type of attack.
- **Ad-hoc networks.** When clients associate directly with one another, they form ad-hoc networks. These peer-to-peer connections can pose a risk if an unauthorized client(s) should automatically associate with a legitimate client storing sensitive data, because it could gain access to that device’s hard drive. In addition, the association could result in one client piggybacking onto the other’s connection to internal wired network resources.
- **Accidental associations.** These take place when a neighboring access point’s coverage area bleeds into the legitimate organization’s air space, triggering the



LAYERS OF PROTECTION Each layer in the network helps protect corporate resources from data hijacking and denial of service attacks.

organization’s legitimate wireless client devices to connect to the neighboring access point. Once the connection is made, a client device on the neighboring network can fairly easily gain access into the legitimate client and the organization’s resources, unless the legitimate client device is protected by software such as a personal firewall or Cisco Security Agent. Cisco Security Agent has recently been enhanced to detect if a client has connected to an invalid subnet range, such as that in an unauthorized network.

To contain the behavior of unauthorized clients connected to rogue access points, Cisco access points in sensor mode can send a deauthentication packet to the client so it will disassociate and find a legitimate access point to which to connect, says Woodhams.

- **Unwanted Spanning-Tree Bridge Loops.** In addition to the open-access threat, Windows XP laptops contain a zero-configuration default setting. “This can cause a spanning-tree bridged loop with older Ethernet switches that can paralyze the entire network,” explains Woodhams.

If spanning-tree protocol information is leaked outside the walls of the organization, for example, freeware tools can sniff it and launch a spanning-tree attack that, in effect, results in denial of service for the network backbone, he says.

“Wireless monitoring systems, by detecting any network protocols leaking into the air outside of corporate borders, arms network managers with the knowledge to take corrective action.”

Best Practices and Tools

The most failsafe monitoring solution involves deploying networks sensors that can detect the presence of all wireless LAN devices and their activities. Among the key practices and tools:

Continued on page 89

SAN Security: Beyond Zoning

By Rhonda Raider

NOT LONG AGO, the security of storage area network (SAN) solutions was considered secondary to performance, connectivity, and port count. Today, SAN security has moved to center stage. One reason is that more companies are extending their SANs outside the data center to provide business resilience (e.g., disaster recovery), partly in response to US security regulations such as the Gramm-Leach-Bliley Act, HIPPA, and the Sarbanes-Oxley Act, and the European Privacy Directive. These regulations compel organizations to protect private company and customer data as it travels between SANs.

The ever-increasing presence of Trojan Horses, worms, and denial of service (DoS) attacks has increased device security awareness. "A single compromised SAN-attached host has the potential to disrupt other hosts attached to the SAN, access unauthorized data within the SAN, or bypass existing firewalls and intrusion detection systems," says Lincoln Dale, technical marketing engineer in Cisco's Storage Group. The myth that Fibre Channel is inherently more secure than IP or Ethernet has been debunked, adds Dale. "Most of us don't hear about Fibre Channel security problems, but that's not because they don't exist," he says. "It's simply that attacks on SANs aren't mainstream due to the relative size of SANs compared to that of IP-based networks."

Another reason for heightened attention to SAN security is the growing trend toward using IP to transport storage traffic. "Using a common IP infrastructure and FCIP [Fibre Channel over IP] for SAN extension between data centers for disaster recovery and business continuance

provides a lower-cost alternative to dedicated connectivity," says Dale. "Most replication solutions transfer data unencrypted, so there is certainly a requirement to protect sensitive data if it's going over an unsecured network." Similarly, there is an increasing trend toward providing lower-cost access to storage using SCSI over IP (iSCSI). "iSCSI is popular because hosts and servers can connect to the network using their built-in Ethernet card, eliminating the costs of a host bus adapter [HBA] and Fibre Channel port on the switch," says Dale. "However, if the storage data is sensitive, both iSCSI and FCIP introduce the need to protect SAN traffic traveling over IP networks."

Finally, it is just as important for SAN security to prevent accidental data loss and corruption as it is to protect against intruders. "A zoning configuration error can cause as much devastation as if an intentional breach occurred," notes Dale.

These factors have spurred IT groups to attend to SAN security with the same thoroughness long paid to LAN and WAN security. Until now, one of the few widespread SAN security practices has been *zoning*, or enforcing access controls within Fibre Channel. "Zoning is better than nothing, but there are relatively easy ways to defeat it," explains Dale. "Soft zoning provides 'security through obscurity.' It's analogous to keeping a secret military center off a map, but not providing any guards to protect it if found. Hard zoning is better—every frame is checked as it passes through the switch. But even hard zoning cannot provide protection against spoofed addresses."

Many established LAN and WAN security technologies can be applied effectively to SANs, as well. These and SAN security technologies from Cisco provide companies with the economies of IP-based access and the confidence that their data is securely protected end to end. "The most critical attribute of a data center security plan is that it is end to end," says Dale. "A SAN cannot be secured independently from the LAN or WAN used to access it."

End-to-End Approach Earns A++

Within the Cisco MDS 9000 Family, Cisco has applied its expertise in LAN and WAN security measures to the unique challenges of SANs. In 2004, Gartner awarded the Cisco MDS 9000 Series multilayer SAN switches an "A++" for Fibre Channel SAN fabric security. What differentiates the Cisco MDS 9000 is its attention to all six areas of SAN security requirements (see Figure 1, page 44):

Fabric access—secure fabric access to fabric services

Target access—secure access to targets and logical unit numbers (LUNs)

SAN fabric protocol—secure communication and authorization for switch-to-switch Fibre Channel protocol communication

IP storage access—secure FCIP, used to interconnect SANs in two data centers for disaster recovery or application resilience, as well as secure iSCSI services, used for low-cost access to lower-end servers

Data integrity and secrecy—encryption of data in transit

SAN management access—secure access to management services

Interconnected SANs and IP-based access heighten the urgency of SAN security.

“You can’t have an effective solution if even one of these six elements is missing,” says Dale. “If management security is compromised, for example, an intruder could simply turn off other security mechanisms or make configuration changes to bypass them.”

Also unique in Cisco’s approach to SAN security is that most of the security features are built into the Cisco MDS 9000 Series Switch and do not require purchasing optional licenses (see table, “SAN Security Techniques,” page 44).

Fabric and Target Access Security

Unauthorized fabric and target access can compromise application data, LUN integrity, and application performance. The Cisco MDS 9000 Series multilayer switches provide the following security features to protect against these risks:

Fibre Channel zoning—Zoning restricts communication between devices within the same Fibre Channel fabric, preventing a host from gaining access to a disk used by another host and corrupting its data. Cisco MDS 9000 Series switches support both software-based zoning (soft zoning) and hardware-based zoning (hard zoning) for up to 2000 zones and 20,000 zone members. The switches enforce hard zoning by

applying hardware access control lists (ACLs) to every Fibre Channel frame as it is switched.

LUN Zoning and Read-Only Zoning—LUN Zoning, a capability unique to Cisco MDS 9000 Series switches, blends deep frame inspection and hard zoning. IT administrators can restrict access to explicit LUNs within a storage array. Read-Only Zoning is useful for systems such as multimedia servers that do not require write access to storage.

VSANs—VSANs increase the security and stability of the Fibre Channel fabric by logically isolating devices that are physically connected to the same set of switches. “Faults within one fabric are contained within a single VSAN and are not propagated to other VSANs,” explains Dale. No communication is possible between devices in different VSANs except where explicitly allowed through the use of the Cisco MDS 9000 Inter-VSAN Routing feature.

Port security—If an IT administrator enables port security for a particular port, devices can connect to that port only if they are listed as bound to the given port in the port security database.

Port mode security—This mode restricts the function of a port, for example, to prevent edge ports from inadvertently being used for inter-switch links (ISLs).

FC-SP DH-CHAP—FC-SP DH-CHAP helps ensure data integrity and authentication for host-to-switch and switch-to-switch communication. All major HBA vendors and some SAN switch vendors support FC-SP DH-CHAP. Authentication can either be performed locally in the switch or remotely through a centralized RADIUS or TACACS+ server. FC-SP DH-CHAP is the only technology available today that provides complete protection against spoofed addresses.

SAN Fabric Protocol Security

Many of the same features of the Cisco MDS 9000 Series switches used for fabric and target access security also help ensure SAN protocol security. Additional SAN protocol security capabilities include:

Disruptive Reconfigure Fabric Rejection—This feature protects against human errors by rejecting fabric reconfiguration requests that could cause an outage. These requests might come from misconfigured or new unconfigured switches when they are attached to an existing fabric.



SAN SECURITY

IBM Fiber Connection (FICON) Fabric Binding—Cisco MDS 9000 switches can restrict participation in a FICON fabric based on the switch and domain ID.

Fibre Channel ID Caching, Persistent Fibre Channel ID Allocation, and Static Fibre Channel ID Assignment—These features provide persistency to Fibre Channel IDs that are assigned to worldwide port names (pWWNs), regardless of switch restarts and physical port.

IP Storage Security

iSCSI provides SAN access at lower price point than that of Fibre Channel. When equipped with optional IP service modules or multiprotocol service modules, Cisco MDS 9000 switches can be configured to accept incoming iSCSI connections from hosts (iSCSI initiators), as well as use FCIP for IP SAN extension. Security features in the Cisco MDS 9000 switches include:

iSCSI authentication—Before establishing an iSCSI session, the switch authenticates the iSCSI initiator using CHAP.

iSCSI initiator persistent dynamic WWN and static WWN allocation—The switch can dynamically or statically map iSCSI initiators to virtual Fibre Channel initiators, enabling midrange and enterprise-class storage arrays to uniquely identify hosts connected via iSCSI in the same way they can identify hosts connected via Fibre Channel HBA.

iSCSI access controls—IT administrators can apply access controls to iSCSI initiators based on the target, VSAN, storage device, or interface. “In the latter case, individual iSCSI targets can be advertised on all or some Gigabit Ethernet interfaces, on subinterfaces, or VLANs,” says Dale.

FCIP—Cisco MDS 9000 switches also support FCIP, generally used for SAN-to-SAN traffic. “While FCIP itself does not have any explicit security, it can use all existing security mechanisms available to native Fibre Channel,” says Dale. “These include port security and FC-SP DH-CHAP switch-to-switch authentication.”

Data Integrity and Secrecy

Neither iSCSI nor FCIP protect data traversing IP networks. “If a rogue device in the path were able to eavesdrop, it could view storage data as it traveled across the link,” Dale warns. To protect data in

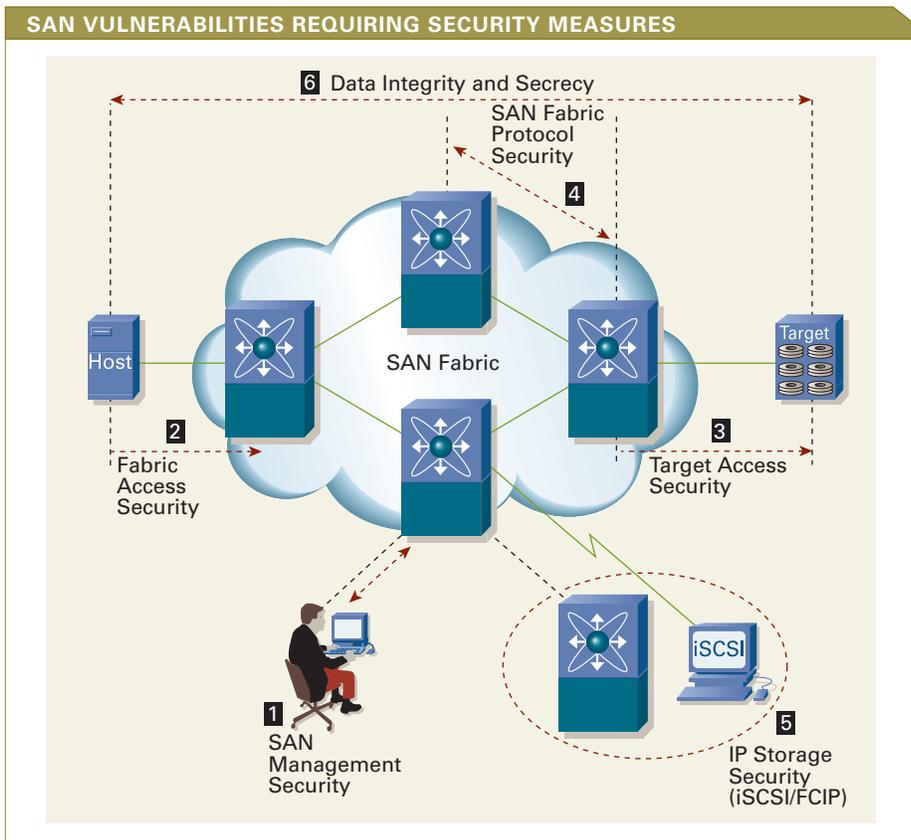


FIGURE 1 The Cisco MDS 9000 Series Switch addresses all six key areas of SAN security earning it the top grade from Gartner for Fibre Channel SAN fabric security.

SAN SECURITY TECHNIQUES

SAN SECURITY TECHNIQUE	VULNERABILITY IT ADDRESSES	CISCO MDS 9000 SERIES SWITCH FEATURES
Segregation of traffic destined for different server farms	Fabric and Target Access Security	Virtual SANs (VSANs) Hard zoning Fibre Channel port security
Authentication and integrity for switch-to-switch communication	Fabric and Target Access Security SAN Fabric Protocol Security Fibre Channel Port Security	Fibre Channel Security Protocol (FC-SP) Diffie-Hellman Challenge Handshake Authentication Protocol (DH-CHAP)
Encryption, to prevent data theft	Data Integrity and Secrecy	Integrated IP Security (IPSec) support
Traffic monitoring to identify malicious activity	SAN Management Security SAN Management Security	SPAN RSPAN Fibre Channel flow statistics Call Home RMON threshold alarms
Secure management, to limit the risk of an attacker gaining control of SAN devices		Authentication, Authorization, and Accounting (AAA) Secure Shell Protocol version 2 (SSHv2) Simple Network Management Protocol version 3 (SNMPv3) Syslog Network Time Protocol version 3 (NTPv3) Role-Based Access Control (RBAC)

BEST PRACTICES FOR SAN SECURITY

AREA OF VULNERABILITY	RECOMMENDED BEST PRACTICE
Fabric access	<ul style="list-style-type: none"> Use VSANs to isolate departments Use port security features everywhere Use FC-SP DH-CHAP authentication for switch-to-switch fabric access Hard-fix switch port administrative modes to assigned port function
Target access	<ul style="list-style-type: none"> Use zoning services for isolation where required Consider only allowing zoning configuration from one or two switches, to minimize access Use VSANs to divide and manage individual fabric configuration and resilience
SAN fabric protocol	<ul style="list-style-type: none"> Use WWN-based zoning for convenience and use port security features to harden switch access and limit zoning access to 1 or 2 SAN administrators Secure access to control protocol configuration using Cisco RBAC Enable port-security for locking of ISL ports Use FC-SP DH-CHAP for switch-to-switch authentication to block rogue ISLs Consider using static configuration (Domain_ID and Principal switch) for greater security than plug-and-play fabric protocol configuration
SAN management security	<ul style="list-style-type: none"> Use RBAC to grant adequate—not excessive—privilege to SAN administrators Use RADIUS or TACACS+ for centralized user account administration and auditing Use secure forms of management protocols (SSH, SFTP, SCP, SNMPv3, SSL) and disable others Enable NTP across all switches for consistent time stamping of events Log and archive all events, including syslog, configuration, and Call Home

flight, the Cisco Multiprotocol Switching 14+2 (MPS 14+2) line card and Cisco MDS 9216i Multilayer Fabric Switch offer integrated hardware-based IPSec support, providing wire-rate IPSec encryption and decryption with Advanced Encryption Standard (AES) and Triple Data Encryption Standard (3DES).

SAN Management Security

Unauthorized SAN management access is risky, because without effective safeguards, a malicious user could alter the network configuration. The three main areas of vulnerability in SAN management access are disruption of switch processing, compromised fabric stability, and compromised data integrity and secrecy. The Cisco MDS 9000 mitigates these risks:

AAA—Either TACACS+ or RADIUS can be used to provide both authentication and accounting for management access on a centralized basis. If no AAA server is used, a username/password database local to the Cisco MDS 9000 Switch may be used.

RBAC—RBAC allows each user to be assigned to a specific role, with management capabilities and restrictions potentially on a per-VSAN basis. “This approach allows companies to consolidate storage

while restricting administrators’ access to the fabric ‘island’ that they managed before consolidation,” says Dale.

SSHv2—An alternative to insecure protocols such as Telnet, rlogin, and FTP, SSHv2 provides secure remote access through authentication and encryption. It can be used with TACACS+ and RADIUS.

SSL Version 2 and TLS 1.0—Cisco MDS 9000 switches support the Storage Management Initiative Specification (SMI-S), the set of common interfaces based on Common Information Model (CIM) that allows multiple-vendor interoperability in a SAN environment. Management access via SMI-S is protected through SSL.

SNMPv3—An application-layer protocol, SNMP facilitates the exchange of management information between network devices. All Cisco MDS 9000 switches support SNMPv1, v2c, and v3. SNMPv3 (RFC 2271-2275) provides authentication and integrity using an MD5 MAC or SHA HMAC algorithm, and encryption with DES. The Cisco MDS 9000 also supports stronger AES 128-based encryption with SNMPv3 (RFC 3826).

Syslog—Syslog messages are unsolicited notifications that a network device can

save in a log file and/or direct to a server such as CiscoWorks Resource Manager Essentials (RME). Syslog messages include a timestamp from the syslog server, a device name, a sequence number, the timestamp from the network device, and the message itself.

Accounting log—Cisco MDS 9000 switches maintain an accounting audit trail of configuration commands. Commands can also be logged to centralized syslog and AAA servers through RADIUS or TACACS+ accounting messages. Critical audit logs are stored in NVRAM and are persistent across restarts and power loss.

Call Home—This feature can be used to send e-mail or pager notification to IT personnel when critical system events occur. It can also initiate Cisco AutoNotify services for direct case generation with the Cisco Technical Assistance Center (TAC).

Fabric Consistency Checker—Embedded within the Cisco MDS 9000 management suite, Fabric Consistency Checker highlights configuration deviations from the master policy switch and provides a mechanism to resolve the differences.

ACLs—Administrators can limit management and IP access to a subset of IP addresses by applying ACLs to various management and Gigabit Ethernet interfaces.

How Much Security Is Enough?

The degree of SAN security a company needs depends on its risks. “Ask yourself what would be the cost if a competitor or hacker got hold of the information,” advises Dale. “Usually it’s a combination of hard costs, such as a bank having to issue new credit cards, and soft costs, such as loss of customer trust. It’s important to understand the threats so you know the potential worth of the investment.”

When it comes to protecting against external, internal, and unintentional threats, Cisco recommends the practices shown in the table above. ■

FURTHER READING

- Cisco Storage Networking Solutions cisco.com/go/storage networking

Stopping Bad Behavior at Endpoints

By Gene Knauer



LIKE A NATURAL DISASTER, when the Sapphire Worm, better known as “Slammer,” was unleashed in December 2003, it shut down Websites, disabled automated teller machines (ATMs), flooded networks, and resulted in a massive loss of productivity, money, and peace of mind for numerous businesses and IT staffs. Meanwhile, however, some enterprise networks, including the University of California, Berkeley, remained uninfected, even though their servers and desktop PCs had yet to be patched to prevent Slammer from being transmitted and exploiting buffer overflow vulnerability in computers running Microsoft’s SQL Server or Microsoft SQL Server Desktop Engine 2000.

Cisco Security Agent prevents attacks on servers and desktop PCs by enforcing behavioral policies.

These networks were protected by an end-point security product from the small intrusion prevention software company Okena, which that same year was acquired by Cisco and the product rebranded as Cisco Security Agent.

Behavior-Based Endpoint Security

Cisco Security Agent software is considered an intrusion *prevention* tool. Working from network endpoints such as desktops and servers, it is designed to correlate appropriate and suspicious behavior and prevent new attacks, even before a security patch or “signature” can update the network’s antivirus or other security software. In sum, Cisco Security Agent intercepts system calls between applications and the operating system, correlates them, compares the correlated system calls with a set of behavioral rules, and makes an “allow” or “deny” decision based on the comparison results. This process is called INCORE, which stands for “intercept, correlate, rules engine” (see Figure 1).

According to Ted Doty, product manager for Cisco Security Agent, the basic mechanisms behind keeping viruses and worms in check have not changed much over time. “I like to think of it as catching thieves in the bank before they can rob it,” he says. “We’re looking for malicious behavior based on system calls to files, network registry sources, or to dynamic, run-time resources.”

Cisco Security Agent includes a management console that resides on a Microsoft Windows 2000 server and host-based agents deployed on desktops and servers. The agents use HTTP and 128-bit Secure Sockets Layer (SSL) for the management interface and agent-to-management console communications. Running between network applications and operating system kernels, Cisco Security Agent checks applications against their security policies and either allows or denies the operation. Such real-time prevention is based on enforcing security policies combined from distributed firewalls, operating systems, antivirus software, and audit event collection.

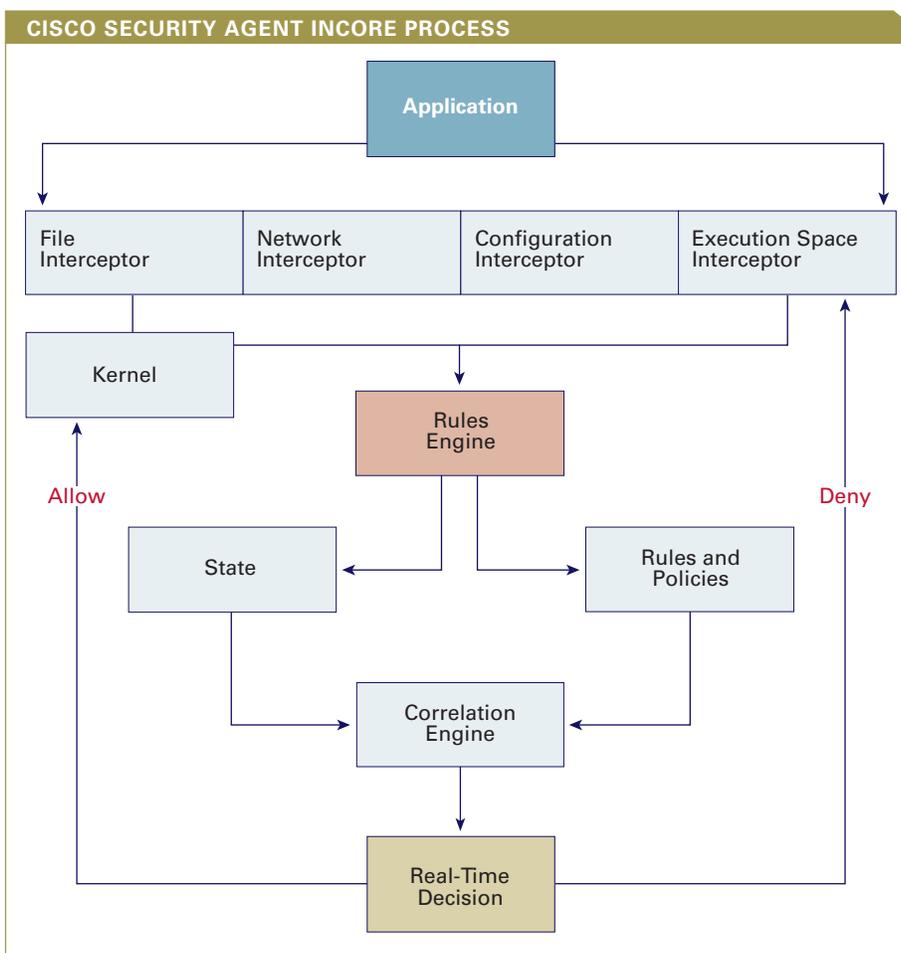


FIGURE 1 Cisco Security Agent applies an “intercept, correlate, rules engine” process—INCORE—that compares correlated system calls with a set of behavioral rules.

“Correlating events with policies to allow or shut down activity is what makes Cisco Security Agent unique on the market,” says Doty. “Other solutions for viruses, worms, and spyware and adware detection rely on applying the latest security patches. From a couple of hundred patches issued per year in the mid-1990s, now there are about 4000 patches a year. Getting them tested and deployed on every server and desktop in a network has been bleeding customers dry in time and resources.”

Defense in Depth for Siemens and IFF

For Kathy Taylor, information security officer at Siemens Energy and Automation in Alpharetta, Georgia, deploying Cisco Security Agent was like acquiring a staff of new security administrators to watch the 250 servers and 7000 desktops on the company’s highly distributed network serving users throughout the US and Mexico.

“We had previously been hit hard by the W32/Blaster Worm in the summer of 2003 and soon after got the approval to install Cisco Security Agent,” says Taylor. “The following spring, there was another global virus outbreak, but this time we had no issues.” Taylor and her colleagues could see viruses trying to attack their computers, but none of these network operations were allowed to proceed.

“Cisco Security Agent gives us time to do the antivirus updates and test the new OS security patches before installation,” she says.

With facilities in 32 countries, International Flavors & Fragrances Inc. (IFF), a creator and manufacturer of flavors and fragrances used in a wide variety of products, had a similar sobering experience before investing in endpoint prevention. The Welchia virus of early 2003 swept across the company’s network globalwide.

“Welchia hit in our offices in China first,” recalls Michael Wasielewski, senior manager for network systems at IFF, which is based in Union Beach, New Jersey. “By the time we realized we were dealing with a virus, two hours later it had spread to Europe and Asia, only because the western world wasn’t yet awake. The antivirus signatures weren’t available for another eight hours.”

Though IFF squelched the Welchia virus without any serious disruptions, Wasielewski says, “We saw the agony other companies went through, and we made the decision to buy an endpoint security system.” There were alternatives to Cisco Security Agent, and Wasielewski researched them. They included devices that would block network access to unpatched systems and others that would inspect systems to determine whether they were at the proper virus patch level.

“We still couldn’t get around the fact that we had to deploy these patches,” Wasielewski says. “And the process of getting them, and testing and deploying them was too slow. The viruses were coming too fast. Back then, Microsoft was patching patches. We decided that we needed ‘Day Zero’ protection, a solution that didn’t depend on catching up to an already-detected new intrusion event.”

Wasielewski and his network colleagues at IFF found Cisco Security Agent to be further ahead in its behavioral approach to preventive security than any other product they researched. They have since deployed Cisco Security Agent on 4500 desktop computers throughout IFF.

“It’s the first product we’ve seen that really delivers this extra layer of endpoint security, which we now see as the first layer of protection even before antivirus or anti-spyware tools,” says Wasielewski.

Thwarting Spyware and Adware

Among the intrusive network behaviors targeted by Cisco Security Agent Version 4.5, the latest release introduced in February, are spyware (programs that install themselves on computers without a user’s consent and read and relay private information, including passwords and credit card numbers) and adware (marketing programs bundled with freeware that sprout pop-up ads and links). Cisco Security Agent 4.5 protects against spyware and adware infections by preventing these programs from initially installing and, if already installed, by preventing them from executing.

Cisco Security Agent is aptly suited to thwarting spyware and adware because these programs are rarely delivered through e-mail, which is subject to antivirus screening. This software is also an improvement over spyware detection

Continued on page 51



FIGURE 2 Cisco Security Agent Version 4.5 detects an attempted keystroke capture and alerts the user with courses of action.



FIGURE 3 A network status “Events” view summary report generated using the CSA MC Web-based interface.



“We decided that we needed ‘Day Zero’ protection, a solution that didn’t depend on catching up to an already-detected new intrusion event.”

Michael Wasielewski, senior manager for network systems, IFF

Endpoint Security, Continued from page 49

ENDPOINT SECURITY AND NETWORK ADMISSION CONTROL

Cisco Security Agent can be considered a first-order dampener to the effects of virus and worm propagation. Making sure endpoints are compliant with OS patches and antivirus software updates is an effective second-order dampener to such propagation. Enter Cisco's Network Admission Control (NAC) program. The NAC industrywide initiative was created to help ensure that every endpoint complies with network security policies before being granted access to curtail damage caused by viruses and worms.

NAC technologies control access by interrogating devices connecting to the network to determine whether they comply with network security policy. For example, NAC can determine if Cisco Security Agent or antivirus software is installed and current, along with the current OS and patch level. NAC uses this information to determine appropriate network admission policy enforcement for every endpoint based on the security state of the OS and associated applications rather than simply on who is requesting access. In addition to controlling access, NAC gives IT administrators the means to automatically quarantine and remediate noncompliant endpoints. Launched in June 2004, NAC is supported on routers running Cisco IOS Software Release 12.3(8)T and higher.

The Cisco NAC program is open to vendors who design and sell third-party client and server applications that incorporate features compatible with the NAC infrastructure. To date, more than 30 vendors are actively integrating their technologies into the network. For more information on this program, visit cisco.com/packet/172_6d1.

tools because, like antivirus and other forms of security software, these tools are passive and reactive, with patches lagging behind new and mutating spyware attacks. Instead, Cisco Security Agent 4.5 hardens the Windows operating system with its behavior correlation engine, preventing spyware from executing.

In Figure 2 (page 49), for example, Cisco Security Agent detects the problem Silent-Log.exe, a "keystroke logger" program that quietly captures all keyboard input and logs it to a file. Spyware often installs such keystroke loggers to capture passwords entered by users.

In response to stealthily downloaded spyware or adware attempting to execute, Cisco Security Agent alerts the user with a message screen and will default to terminating the application unless the user allows the process to continue (by clicking "Yes"). Administrators can configure Cisco Security Agent to automatically stop the application from executing without user intervention. If the spyware attempts to swamp users with repeated requests to download—a form of social

engineering intended to trick or frustrate users into selecting "Yes"—they need only select "Don't ask me again" to stop the requests.

Cisco Security Agent does not require cryptographic analysis of file system contents, so its impact on performance is negligible.

Other Benefits of Cisco Security Agent

Besides detecting, analyzing, and acting on network behavior, Cisco Security Agent can track which applications are installed on a single computer or workgroup; which applications use the network; the identity of all remote IP addresses with whom a server or desktop computer communicates; and the state of all applications on remote systems, including user-specific installation information and whether undesired applications are attempting to run.

Administrators can perform detailed forensics of any application on any computer, collect information about the application's behavior, and create a control policy based on that application's "normal" behavior. All Cisco Security Agent policies are configured and deployed via the Cisco Security Agent Management

Center (CSA MC) Web-based user interface. CSA MC also provides a reporting tool, allowing administrators to generate reports with various views of their network's health and status (see Figure 3, page 49).

Cisco Security Agent Version 4.5 also adds compatibility with international operating systems and expands platform support to include Linux servers and desktops and Windows clusters. It ships at no additional charge with all Cisco IP telephony products, including Cisco CallManager and Cisco Unity.

"Now we're considered the most robust IP telephony solution from a security perspective," says Doty, adding that more than two million desktop PCs and servers have installed Cisco Security Agent since 2001.

Frontline of the Self-Defending Network

"Cisco Security Agent complements Cisco's Self-Defending Network strategy. In addition to providing a first line of real-time intrusion prevention, its presence on endpoints allow them to acquire state information that might not be available at the network edge," says Joshua Huston, a technical marketing engineer in Cisco's VPN and Security Business Unit specializing in Cisco Security Agent marketing. "This capability provides a feedback loop between the endpoints and the network, so the network can readily adapt to emerging threats." (For more on the Self-Defending Network strategy and new security products from Cisco, see "In Self Defense," page 26.)

Cisco Security Agent embodies other attributes of a Self-Defending Network, adds Huston: It's flexible, future-proof, and highly effective whether a user is at work, at home, or on the road. ■

FURTHER READING

- White paper: *Cisco Security Agent—An Enterprise Solution for Protection Against Spyware and Adware*
cisco.com/packet/172_6d2
- Cisco Security Agent home page
cisco.com/go/csa

The RFID-Ready Network

IP-Based Network Connectivity for RFID Deployments

By Roland Saville and Dennis Vogel

The adoption of Radio Frequency Identification (RFID) is steadily growing across industries and applications. It is driven not only by retail and government mandates, but also by organizations that are beginning to recognize the potential for RFID to increase productivity, improve customer satisfaction, and strengthen competitive advantage. As the number and scale of these deployments grows from small internal pilots to larger multisite installations with data shared across companies, it is increasingly important to understand that tags and readers are only a portion of what is required for a working RFID system.

IP networks are critical to fulfilling the promise of RFID, by providing the foundation of services that allow the coexistence and communication between multiple systems required to reliably deliver RFID data when and where needed. The combined network intelligence that secures, stores, and shares Internet information remains the key to enabling real-time event data generated by RFID that will improve efficiencies and automate decision making.

Customers who have started to look beyond their short-term testing are beginning to appreciate the value of their existing IP networks in the context of RFID applications. IP networks are a vital way to help increase efficiency by providing an open, standard framework for real-time product information exchange. Network managers who are thinking about the addition of RFID to their networks are also considering how to best use existing resources, as well as evaluating what additional services might be useful to them.

Role of the IP Network Within the EPC Architecture

While the practices described here can be applied to many wireless identification and tracking systems, this article uses

EPCglobal Network Standards

EPCglobal (www.epcglobalus.com) is a not-for-profit organization made up of industry leaders whose mission is to establish global standards for the development, implementation, and adoption of Electronic Product Code (EPC) and RFID. The architecture designed by EPCglobal enables identification and sharing of information of items in various supply chains. The supporting technical standards developed through EPCglobal are meant to ensure that the various components and technologies of the EPCglobal architecture work together, on a global scale.

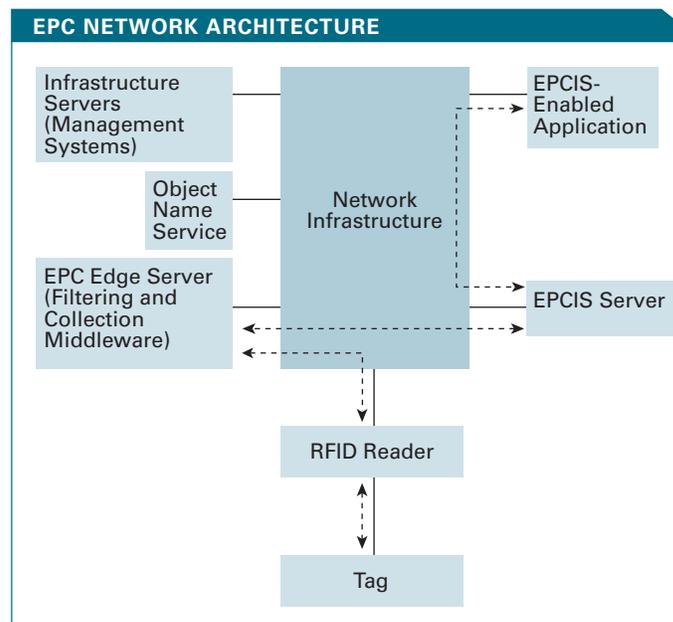


FIGURE 1 The IP network infrastructure is the enabling technology that facilitates the communication flows between the individual architecture components.

EPCglobal passive tag systems as an example. Figure 1 is a simplified view of the EPC network architecture.

In this example, the RFID reader, filtering and collection middleware, EPC Information Service (EPCIS) server, and EPCIS-enabled application are deployed as separate physical entities, with traffic flowing across the IP network infrastructure between these components. In reality, multiple functions can be standalone devices or combined into a single physical device or devices.

RFID-Ready Network Architectural Framework

The Cisco RFID-Ready Network is based on a set of services that the network infrastructure can provide to facilitate the implementation of RFID deployments on top of the existing IP-based network infrastructure. Its purpose is to ensure that EPC and RFID traffic and devices can coexist with other devices and traffic on a single, converged, and open-systems IP-based infrastructure (Figure 2).

The network can provide the following services to facilitate the implementation of RFID deployments:

Connectivity services provide basic wired and wireless Ethernet connectivity for RFID devices on the network. Also included are

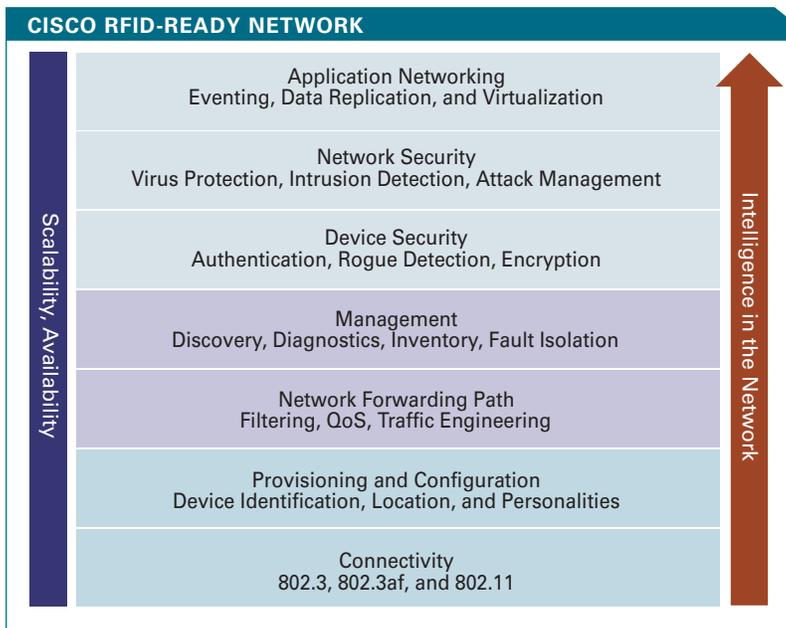


FIGURE 2 The network's built-in intelligence facilitates RFID deployments.

such features as IEEE 802.3af Power over Ethernet (PoE), which can help reduce deployment costs.

Provisioning and configuration services include Dynamic Host Control Protocol (DHCP), Domain Name Service (DNS), Trivial File Transfer Protocol (TFTP), and Network Transfer Protocol (NTP), among others, which assist in the deployment of RFID devices onto the network. This layer also includes services such as Cisco SmartPorts macros, which assist in provisioning the network infrastructure rapidly and more easily to support RFID devices.

Network forwarding path services ensure that RFID traffic has the necessary quality of service (QoS) parameters assigned to it and that the network has provisioned the necessary bandwidth and prioritization of RFID traffic for proper operation and coexistence with other traffic on the network.

Management services assist network operations staff in identifying network outages and correlating such information with the resulting application-level outages.

ROLAND SAVILLE is a technical lead for the Enterprise Systems Engineering group within Cisco's Internet Technologies Division. He has more than nine years experience as a systems engineer, consulting systems engineer, technical marketing engineer, and technical lead at Cisco. He can be reached at rsaville@cisco.com.

DENNIS VOGEL is part of the Multiservice Customer Edge Business Unit (MCEBU) advanced technologies team, focused on RFID product and technology development for Cisco's Chief Development Organization. Previously, he was product line manager of firewall and VPN products as well as an advocate for IPv6 security offerings. He can be reached at dvogel@cisco.com.

Device security services assist network operation staff in ensuring only authorized RFID devices exist on the network, and in determining and locating potential rogue devices on the network.

Network security services mitigate the spread of malicious activity, and limit the potential of denial-of-service attacks, man-in-the-middle attacks, etc., across the network. Such attacks could result in a loss or degradation of service or sacrifice confidentiality of information across the RFID deployment.

Application networking services identify and facilitate the movement of RFID data across the network from an application perspective.

The layers represent increasing intelligence built into the network that facilitates RFID deployments, from connectivity services up through application networking services. In addition, each of these layers must also scale from very small RFID deployments up to large RFID deployments and have the ability to be deployed in nonresilient and resilient designs for those who require high availability. This article focuses on the bottom layer: connectivity services.

Connectivity Services for RFID Deployments

The following recommendations represent guidelines and technology designed to enhance the ability of the network infrastructure to support an RFID implementation. They focus on the connectivity services a LAN infrastructure can provide for an RFID deployment, including physical connectivity of RFID readers and printers, RFID middleware edge servers, and application servers to the network.

Connectivity for Wired RFID Devices

Physical connectivity for wired RFID readers is typically provided by dedicated 10/100-Mbit/s Ethernet connections for each RFID reader or printer. Given the amount of traffic produced or consumed, Gigabit Ethernet connectivity is not expected to be necessary for wired RFID readers and printers in the near future. Shared 10 or 100 Mbit/s hub technology is not considered desirable for an RFID deployment, because every RFID reader and printer must contend with each other for access to the network. In a large deployment, this could easily lead to congestion and lost data within the LAN. Given the real-time nature of cases or pallets moving along a high-speed conveyor or through a dock door portal; any data loss is undesirable.

IEEE 802.3af Power over Ethernet

When deploying hundreds of RFID readers around a large distribution center, eliminating the need to run AC power to each reader can result in significant cost benefits. Therefore, the ability of an RFID reader to support IEEE 802.3af Power over Ethernet (PoE) is considered desirable. A small but growing number of RFID readers support IEEE 802.3af PoE. When upgrading an existing infrastructure or deploying a

new LAN infrastructure to support current and future RFID deployment requirements, select IEEE 802.3af PoE-capable LAN switches. In addition to RFID readers, PoE-capable LAN switches can significantly reduce the costs of 802.11a/b/g access point and IP telephony deployments as well.

For EPC edge servers and application servers, physical connectivity services are typically provided by dedicated 10/100-Mbit/s Ethernet or Gigabit Ethernet switch ports.

Wireless Connectivity

Multiple RFID reader deployment models rely upon wireless connectivity, including the following examples:

- A fixed-position wireless reader that can be used in place of a fixed-position wired reader. Physical connectivity services in such a model are provided to the RFID reader itself.
- Deployment of an RFID reader on a mobile vehicle, such as a forklift, within a distribution center. The RFID reader is serially attached to a mobile industrial PC mounted on a forklift. EPC data, in the form of tag reads, enters the network via the serially attached RFID reader and is converted to application data within the industrial PC. The industrial PC then transmits and receives this information to a local server running inventory management, warehouse management, enterprise resource planning (ERP), or other enterprise applications. Physical connectivity services in such a model are provided to the industrial PC, rather than the RFID reader itself.
- Deployment of an RFID reader in a handheld mobile computer within a store. The existing barcode scanner is either augmented or replaced with an RFID reader. EPC data in the form of a tag reads enters the network via the RFID reader and is converted to application data within the handheld computer. The industrial PC then transmits and receives this information to a local server running inventory management, warehouse management, ERP, or other enterprise applications. Physical connectivity services in such a model are provided to the mobile handheld computer with the RFID reader attachment or PC card.

In each of these models IEEE 802.11g, 802.11b, or 802.11a, wireless connectivity can be provided by Cisco Aironet access points, which are themselves connected to Cisco Catalyst switch ports using 10/100-Mbit/s PoE connections (Figure 3).

Because wireless connectivity is a shared medium, individual RFID readers and printers must contend with each other and other non-RFID devices for access to the network. In large deployments or very busy networks temporary congestion could result in data loss within the LAN, particularly as RFID tagging reaches the item level in the future. These concerns are not as great for

CISCO AIRONET ACCESS POINT PRODUCT LINE		
PLATFORM	RADIO	RAW BANDWIDTH
Cisco Aironet 1100 Series Access Point	802.11b or g	10 Mbit/s or 54 Mbit/s
Cisco Aironet 1130 AG Series Access Point	802.11a and g	54 Mbit/s and 54 Mbit/s
Cisco Aironet 1200 Series Access Point	802.11g	54 Mbit/s
Cisco Aironet 1230 AG Series Access Point	802.11a and g	54 Mbit/s and 54 Mbit/s

mobile wireless RFID devices, because the number of such devices deployed is not expected to be great. However, the number of fixed-position wireless devices around dock doors and conveyors in a large distribution center could be in the hundreds. The network design engineer needs to closely monitor the number of wireless RFID devices per access point and the amount of aggregate traffic generated by those devices. Additional technologies, such as applying QoS to the wireless traffic, might need to be considered over time.

IEEE 802.11b provides a shared-access medium with a maximum throughput of 11 Mbit/s. However, due to communications overhead, the effective maximum throughput is approximately 5.5 Mbit/s. IEEE 802.11b operates in the 2.4-GHz frequency spectrum.

IEEE 802.11a provides a shared-access medium that operates at a raw speed of 54 Mbit/s. However, effective throughput is typically in the mid 20 Mbit/s range. IEEE 802.11a operates in the 5-GHz frequency spectrum and is not backward compatible with IEEE 802.11b.

Similar to IEEE 802.11a, IEEE 802.11g provides a shared-access medium that operates at a raw speed of 54 Mbit/s. Effective maximum throughput is around 24.5 Mbit/s. However, IEEE 802.11g operates in the 2.4-GHz frequency spectrum and is backward compatible with IEEE 802.11b.

RFID systems will extend existing IP networks through the installation of new devices like RFID readers as well as integration with larger supply-chain systems. It is essential that those deploying RFID understand the requirements of this application and how it will coexist with other services. Now is the time for network administrators to plan for RFID installations and learn how to maximize the services available to them in their IP network infrastructures. ■

FIGURE 3 IEEE wireless standards are supported in several Cisco Aironet platforms.

FURTHER READING

- Cisco RFID Website
cisco.com/go/rfid
- EPCglobal
www.epcglobalus.com/Network/how_works.html

Data Center Networking

Designing the Server Farm Access Layer

By Mark Noe and Mauricio Arregoces

The central repository of both computing capacity and data storage, the data center is a highly critical component of today's enterprise network architecture. Because of this, the data center requires the highest levels of resilience, performance, and flexibility to meet business requirements and support accelerated growth. The access layer of the data center provides the port density and connectivity to the enterprise server farm. Determining proper access layer design is critical in providing a flexible, scalable, secure, high-performance architecture that supports a mix of existing application requirements while easily adapting to changing conditions.

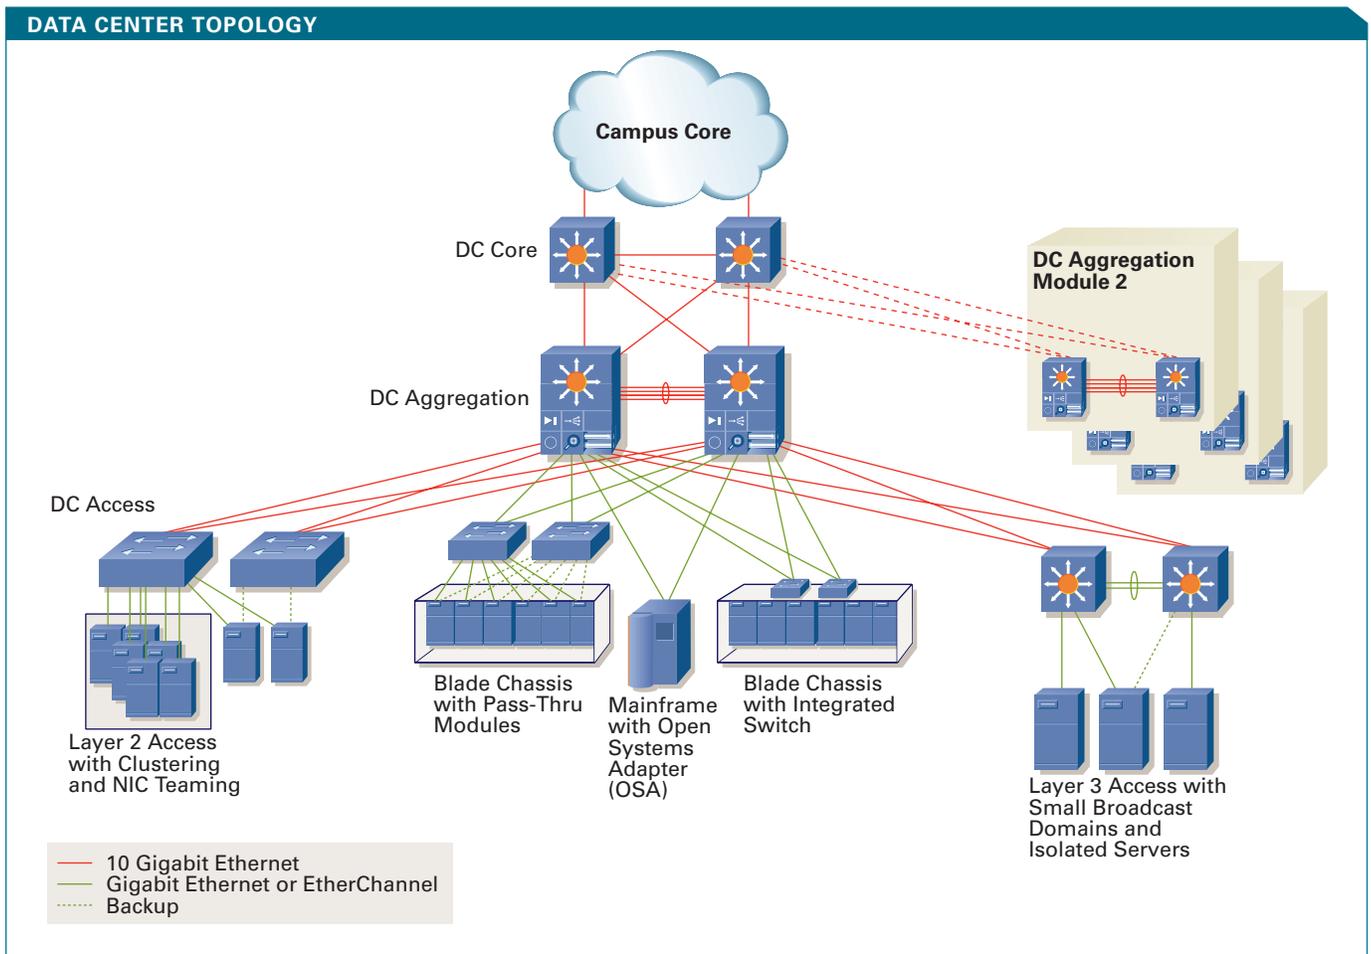
This article helps you prepare for the next-generation server farm by discussing how PCI-X, 10 Gigabit

Ethernet, and port density influence your design considerations. It does not cover other important topics such as data center services, high availability, storage interconnect (that is, Fibre Channel) or server interconnect (that is, Infiniband), as these are the subject of future articles.

Access Layer Considerations

The access layer provides physical connectivity to the server farm. The applications hosted by the server farm have many different requirements; some are business-critical applications requiring dual-homed servers, while others require server clustering for high availability and scalability. Some applications reside on a mainframe that may run a Layer 3 protocol,

FIGURE 1 The enterprise data center topology includes multiple access layer configurations and aggregation modules, which are repeatable for scaling the server farm.



while others are hosted on multiple 1RU servers which may have Layer 2 adjacency requirements (Figure 1).

Layer 2 Access Model Definition

The Layer 2 access model is defined as an access switch that is connected to the aggregation layer through an IEEE 802.1Q trunk. The first point of Layer 3 processing is at the aggregation switch. Layer 3 routing is not performed in the access layer switch. Blade-server chassis that use integrated Layer 2 Cisco switches such as in the IBM BladeCenter and the HP BladeSystem also appear in Figure 1. Because these are Layer 2 switches, they connect directly to the aggregation layer in the same manner as an external access switch. Later, this article describes network interface card (NIC) teaming and clustering, which influence the requirement for the Layer 2 access model.

The Layer 2 model provides significant flexibility by supporting virtual LAN (VLAN) instances through the entire set of access layer switches that are connected to the same aggregation layer. This allows new servers to be “racked” in any available rack yet still reside in the particular subnet (VLAN) in which all other application-related servers are located.

Layer 3 Access Model Definition

The Layer 3 access model is defined as an access switch connected to the aggregation layer through a Layer 3 link (instead of an 802.1Q trunk as in the Layer 2 model) on its own subnet. Layer 3 routing is performed at the access switch. The access switch contains a route processor that provides all required Layer 3 processing functions. Despite the Spanning Tree Protocol being active, there are no blocked ports, thus all uplinks are actively forwarding because this model does not have a looped topology.

The Layer 3 access switch also provides Layer 2 connectivity to the server farm. Unlike the Layer 2 access model, the STP domain is confined strictly to the access switch or to a very small group of access switches, as shown by the Layer 2 links between the Layer 3 access switches in Figure 1. Layer 3 access is used when there is a requirement for multiple active uplinks or to segment groups of servers into smaller broadcast domains to address server stability concerns or to isolate different application environments.

Layer 2 Adjacency Drivers

When Layer 2 adjacency exists between servers the servers are in the same broadcast domain. When servers are Layer 2 adjacent, each server receives all broadcasts and multicast packets from another server. If two servers are in the same VLAN, they are Layer 2 adjacent. However, certain features, such as private VLANs (PVLANS), allow groups of Layer 2 adjacent servers to be isolated from each other but still be in the same subnet.

Frequently, the requirement for Layer 2 adjacency is unexpected or overlooked. High availability clustering, and NIC teaming are primary examples of when Layer 2 adjacency is required.

Clustering

Server clustering implementations vary from high availability clusters such as the Microsoft Windows Server 2003 Cluster Service (MSCS) to high-performance parallel processing clusters as in the Beowulf Linux cluster operating system.

The common goal of clustering is to combine multiple servers to operate as a unified system through specialized software and network interconnections, improving scalability, performance, and resiliency. Clustering technology, initially used in university and scientific environments, is now popular in enterprise data centers primarily for high availability reasons.

Some high availability clusters use mechanisms that require Layer 2 adjacency because the cluster protocol packets are not routable. An example of this is with MSCS cluster environments that use Layer 2 multicast packets to enable all hosts in the cluster to listen to incoming network traffic and pass heartbeats between nodes to determine availability.

NIC Teaming

There are always mission-critical applications that cannot tolerate downtime. To eliminate server and switch single points of failure, servers are dual-homed to two different access switches and use NIC teaming drivers and software for the failover mechanisms. NIC teaming features are provided by a NIC vendor and are becoming widely used in enterprise data centers.

NIC teaming options generally come in three common configurations: Adapter Fault Tolerance (AFT); Switch Fault Tolerance (SFT), also sometimes known as Network Fault Tolerance (NFT); and Adaptive Load Balancing (ALB) (see Figure 2).

The basic objective of NIC teaming is to use two or more Ethernet ports connected to two different access switches. In some cases, two Ethernet ports connect to a single switch using different line cards. The standby NIC port in a server configured for NIC teaming uses the same IP and MAC address of a failed primary server NIC, which results in the requirement for Layer 2 adjacency. An optional signaling protocol is also used between active and standby NIC ports. The protocol

MARK NOE AND MAURICIO ARREGOCES, CCIE No. 1331 and CCIE No. 3285, are members of the Data Center Design and Architecture team at Cisco. They can be reached at marregoc@cisco.com, and mnoe@cisco.com respectively.

heartbeats are used to detect a NIC failure. The frequency of heartbeats is tunable between 1-3 seconds. These heartbeats are multicasted or broadcasted and therefore require Layer 2 adjacency.

Density and Scalability Implications

Enterprise data centers traditionally use a modular access switch with line cards, which can support hundreds of server ports. This type of access switch is strategically placed within or at the end of the server cabinet row and the servers in the row are cabled to it. This scenario permits flexible oversubscription capacity and reduces the ratio of managed network devices to servers.

A simpler option is to place small 1RU access layer switches in each server cabinet with the cabling remaining in the cabinet. These 1RU switch uplinks are cabled directly to the aggregation switches. This simplifies cabling infrastructure by reducing cable bulk and can provide more rack and stack flexibility.

Because there may be different requirements in building a scalable access layer environment, you should carefully consider the connectivity scheme. When selecting which access layer model to use keep in mind the following criteria.

Server density: Consider the most effective way of scaling to the maximum server density and the required number of network connections per server such as client-to-server, server-to-server, server-to-storage, and integrated Lights Out (iLO) out-of-band management.

Network management: How many network devices are managed in the network? Determine the network device-to-server ratio as a guideline to understanding the total cost of ownership (TCO) of the network.

Oversubscription: What is the oversubscription ratio per uplink? Keep in mind the current or future true server capacity (PCI, PCI-X, PCI-Express) and the type of uplinks (Gigabit EtherChannel, 10 Gigabit Ethernet, 10 Gigabit EtherChannel) that can be supported on the access switch. A server that is limited by older PCI bus technology can increase the application traffic dramatically after a platform replacement.

Equipment sparing: Consider how equipment sparing standards decrease skill set requirements and the time to resolution.

Device-level redundancy: Determine the failure exposure level for each application and whether servers require CPU or power redundancy to avoid a single point of failure.

Cabling: Consider the cabling structure design and air flow issues related to cable density through the rack floor entry.

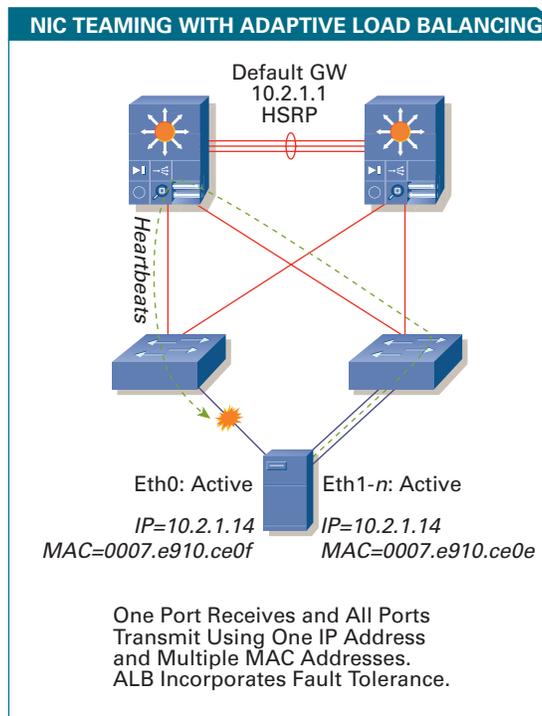


FIGURE 2 Adaptive Load Balancing, a popular NIC teaming solution, permits multiple NIC cards to balance outbound client traffic while providing switch fault tolerance.

Spanning Tree Protocol scalability: Understand how Spanning Tree scales and behaves based on the number of uplinks, VLANs per uplink (logical port limits, number of EtherChannels), and VLANs across switches (STP diameter).

In general, it is important to closely examine the environment and ensure that it has the flexibility to support emerging requirements that can push scalability and performance. The particular areas most affected in the access layer are spanning-tree scalability, cabling, 10 Gigabit Ethernet support, port density, and redundant CPU or power.

Scaling Bandwidth with Gigabit EtherChannel and 10 Gigabit Ethernet Uplinks

Network designers must choose between Gigabit EtherChannel or 10 Gigabit Ethernet uplinks on the access layer switches. EtherChannel technology permits bundling together of up to eight Gigabit Ethernet ports to provide a single logical high-speed uplink. When determining the correct access layer uplink strategy, take into account the following areas.

EtherChannel Hashing Algorithm

EtherChannel provides different hashing algorithms to determine how to balance packets across the ports in the channel group. The most common hashing algorithm used is based on IP source/destination pairs or Layer 4 port number source/destination pairs. The number of servers behind the switch that use the Gigabit EtherChannel uplink and how the IP protocols are used greatly influence how well traffic is balanced across the EtherChannel. If the algorithm

Continued on page 81

Data Centers, Continued from page 59

cannot balance traffic evenly across all ports in the channel, aggregate throughput is degraded. Examine traffic characteristics and determine the proper hashing algorithm to use. If proper load distribution cannot be achieved, 10 Gigabit Ethernet likely provides a better solution than Gigabit EtherChannel.

PCI-X, PCI-Express, and 10 Gigabit Ethernet NICs

When determining the proper oversubscription ratio to use in the access layer, consider the particular technologies that have a dramatic impact on network usage. PCI-X bus-based NICs are common on today's server platforms. Compared to PCI-based NIC cards, PCI-X increases the amount of traffic many times on its interface(s) with less CPU overhead. The next-generation PCI-Express bus interface continues this trend with 4X performance over PCI-X and introduces other improvements, such as Remote Direct Memory Access (RDMA), that reduce CPU overhead. Combined with these bus technology improvements is the ability to tune TCP with larger window sizes, jumbo frames, and TCP offload engines that further improve throughput and lower CPU overhead.

The overall bursting capacity increases as a result of servers using the most current bus architecture (PCI-X or PCI-Express) and newer NIC (10/100/1000 or 10 Gigabit Ethernet) technology. The overall server farm throughput is higher.

10 Gigabit Ethernet NIC cards are available today with driver support across almost all operating systems. These NICs are showing up to a 7X improvement over Gigabit Ethernet with lower latency and less CPU overhead. As the requirements for low latency fabrics increase for clustering, storage, and specialized applications, 10 Gigabit Ethernet NICs are becoming more common. 10 Gigabit Ethernet is also useful when consolidating servers in the data center. By using 10 Gigabit Ethernet NICs in higher-end servers, a reduction in the number of server components can be realized, lowering overall TCO.

As these trends become mainstream, their impact translates into higher strain on access layer uplinks. The ability to migrate from a Gigabit Ethernet EtherChannel uplink to a 10 Gigabit Ethernet or 10 Gigabit Ethernet EtherChannel uplink is an important consideration in choosing the correct access layer platform.

Environmental Concerns

An industry trend is to decrease the amount of rack space used by server and network components by compacting them into smaller platforms. Examples of this are blade-server technology, 1RU servers, and higher density switches. This trend is placing stress on other critical areas in the data center, including cooling, power, weight, and cabling, so be aware of the following general environmental concerns.

Cabling: Examine the network connection requirements of the server farm (iLO, front end, back end, storage) cable routing and the cable bulk at the entry to each cabinet.

Cooling: Consider the cooling capacity per rack and the heat dissipation of the components in each rack. Although rack space may be conserved with denser platforms, cooling capacity limits the density achieved. Cable bulk can also block air flow to the cabinet.

Weight: Consider that the floor tile and subfloor construction might not have the weight support rating that racks loaded with newer high density products are now reaching.

Power: Check the power requirements of newer high-density products and determine whether the current amount of power provided to the cabinet is enough. For example, you might need to retrofit the existing power to include support for 220V or 30A service.

When building a new data center or scaling an existing one, your key considerations should include whether to use the Layer 2 or Layer 3 access models based on either adjacency requirements or a need to create smaller broadcast domains. You also want to consider server density, oversubscription, and management factors when determining whether to use a modular or 1RU access switch. Other considerations include understanding the impact of new server technology such as PCI-X and PCI-Express on the data network architecture and specifically the access layer. ■

FURTHER READING

- Data Center Fundamentals by Mauricio Arregoces and Maurizio Portolani (Cisco Press, ISBN: 1587050234)
ciscopress.com/datacenterfundamentals
- Build the Best Data Center Facility for Your Business, by Douglas Alger (Cisco Press, ISBN: 1587051826)
ciscopress.com/title/1587051826
- PCI Bus Technology
pcisig.com/specifications/pciexpress
- Data Center Best Practices
cisco.com/go/datacenter

Safe Metro Aggregation

Innovative Catalyst switch and QoS features bring security, reliability, resilience, and high performance to the metro aggregation layer.

By Rupa Kaur

The Metro Ethernet market is undergoing explosive growth, and security is extremely important at the metro network's entry location. Modular Cisco Catalyst switches are used to terminate multiple DSL access multiplexers (DSLAMs) at the aggregation layer. DSLAMs are intelligent devices and support multicast Internet Group Management Protocol (IGMP) snooping for triple-play voice, video, and data services. They also support Dynamic Host Control Protocol (DHCP) interface tracking (Option 82) and isolation for end subscribers.

DSLAMs, however, do not offer security features such as dynamic protection from man-in-the-middle attacks, IP spoofing, and DHCP denial-of-service (DoS) attacks. These functions with innovative quality-of-service (QoS) features are performed at the metro aggregation switching layer. Generally, in this topology each DSLAM connects to a Cisco Catalyst switch using two Gigabit Ethernet interfaces (see Figure 2, page 62).

Catalyst 4500 with Supervisor Engine V-10GE

A wire-speed 10 Gigabit Ethernet-enabled Cisco Catalyst switch is a de facto choice for service providers, because it allows them to offer high-bandwidth, rich services that will satisfy customers and keep the service providers competitive. Performance of up to 136-Gbit/s switch capacity and 102 million packets per second (pps) of wire-speed forwarding are supported on a single Cisco Catalyst 4500 Series Switch with a Supervisor Engine V-10GE. Modular supervisors support a full range of 4096 active virtual LANs (VLANs) in accordance with IEEE 802.1q. In addition, none of the services suffer a performance penalty, because they are performed in hardware—allowing providers to offer a greater number of Metro Ethernet point-to-point or point-to-multipoint Ethernet services.

The bidirectional Ethernet (100BaseBX, 1000BaseBX) interfaces in the Catalyst 4500 Series Switch implement full duplex, wire-speed, Fast or Gigabit Ethernet point-to-point services on a single fiber cable. The GLC-BX-U (upstream, customer end) and GLC-BX-D (downstream, service provider) Small Form-factor Pluggable (SFP) interfaces are supported on the switch's Gigabit Ethernet ports. These interfaces provide additional return on investment (ROI), decreasing the cost of underground dark fiber by half. The new bidirectional SFPs are installed in pairs (blue + purple on each end), and each SFP carries a different wavelength. Common deployments include the bidirectional SFPs terminating subscribers on switch downlink connections and 2x10GE line rate uplinks to the Cisco 7600 Series Router in the Metro Ethernet core.

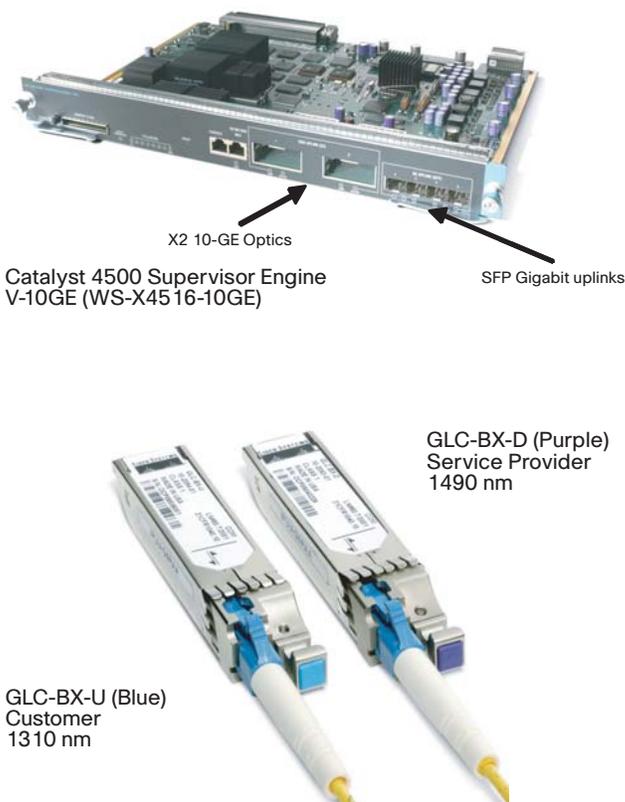


FIGURE 1 New security and quality-of-service features supported by the Cisco Catalyst 4500 Series Switch bring greater performance and protection to metro aggregation deployments. Bidirectional Gigabit Ethernet interfaces can lower the cost of underground dark fiber by half.

New Security and QoS Features in Catalyst Switch

Several new security and QoS features for the modular Cisco Catalyst switches bring a comprehensive security portfolio to metro aggregation deployments and allow network managers to dynamically control security threats at their inception. These tightly integrated software and hardware features work together and can be simultaneously deployed on a switch.

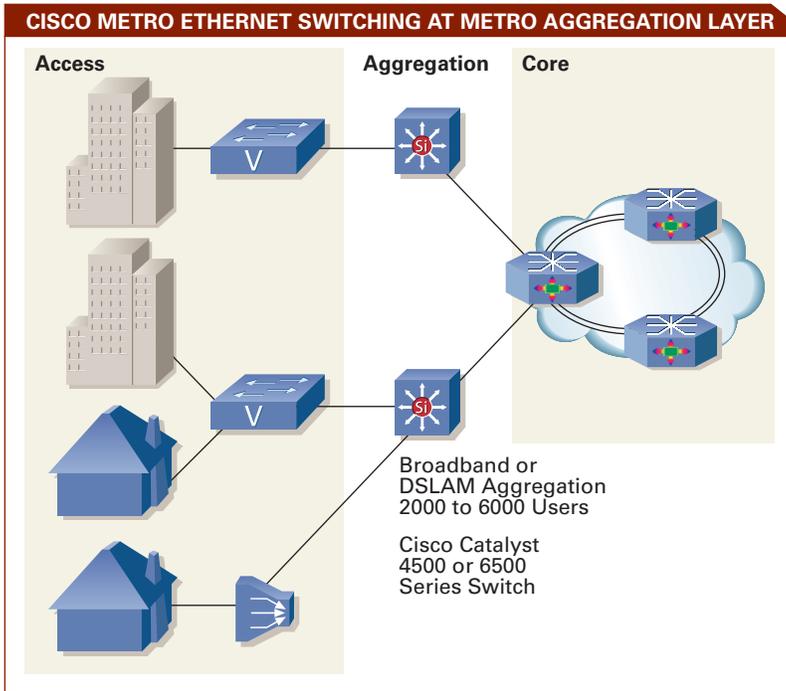


FIGURE 2 Security features—such as dynamic protection from man-in-the-middle attacks, IP spoofing, and DHCP DoS attacks—are performed by the Catalyst switch at the metro aggregation layer.

- **Private VLAN (PVLAN) trunk ports** allow content and media (data, voice, video) distribution to homes from different service providers over the same infrastructure. The trunk ports can carry multiple isolated and regular VLANs and also provide isolation between different ports on downlink connections. This feature simplifies IP address management by keeping all clients on the same IP subnet. PVLAN trunk ports enhance security by providing isolation between the ports and eliminating spoofing attempts of services across subscribers.
- **Promiscuous PVLAN trunks** carry multiple VLANs on uplink trunk ports connected to the router. This security feature also maintains isolation between subscribers carried over the same trunk and simplifies network implementation across the wirespeed 10 Gigabit Ethernet uplink ports.
- **Trunk port security** mitigates MAC spoofing attempts on an inter-switch link (ISL) or 802.1q trunk port. A Catalyst switch can be configured to limit MAC addresses with a per port per VLAN emphasis. This approach prevents various MAC table exhaustion attacks.
- **Per port per VLAN QoS (PVQoS)** is a new feature for input and output QoS. Prior to this feature, a Catalyst switch could only be used for either port level or VLAN level QoS, but not both. This feature allows a metro service provider to customize its own granular QoS service policy per VLAN on any port to better differentiate service offer levels. Multiple service policies for each VLAN are also supported on any given port.

- **8000 input and 8000 output policers** are supported for concurrent input and output policing with PVQoS. This feature allows a service provider to fine tune traffic traversing ingress and egress ports to thousands of subscribers.
- **Aggregation DHCP Snooping, Dynamic ARP Inspection, and IP Source Guard** to dynamically prevent DHCP, man-in-the-middle, and IP spoofing attacks, respectively. These re-engineered security features now allow providers to enable security policies on a DHCP packet even when DHCP interface tracking (Option 82) is performed at the DSLAM level. Prior to this enhancement, these dynamic Layer 2 security features could only be used for metro access deployments, e.g., in a building basement without DSLAMs.
- **Hardware Ternary CAM 3 (TCAM)** is used to look up one or more matching bits in the incoming packets and classify them for different features, such as security ACLs and QoS classification. The set of bits and its value to match is programmed in a TCAM entry, and the set of bits to be considered for matching is programmed in a TCAM mask. In TCAM3, there is one-to-one correspondence between a TCAM entry and TCAM mask, whereas in earlier versions, there are eight TCAM entries for a given mask. Because the new TCAM has more TCAM entries utilization than previous TCAMs, it allows for more security and QoS classification rules.

In addition, with the TCAM3 hardware interface, packet lookup is performed at wirespeed by the switching engine ASIC. These TCAMs also make it possible for Catalyst switches to process security services on any range of IP address in hardware. Because of the one-to-one correspondence between a TCAM entry and its mask, the TCAM3 is amply equipped to meet the future needs of metro aggregation security and QoS features. Even when classifying flows from 6000 different IP addresses and all dynamic aggregation security features, the TCAM entry utilization is only 10 percent.

Figure 3 shows the applicability of these new features within a metro aggregation network. The topology is included to the extent that specific security features should be requested to mitigate the effects of certain attacks.



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

```

Policy-map P31_QoS // A 200 Mbps policer definition
Class RT
Police 200m 16k conform transmit exceed drop // Up to 8K in & 8K out policers

interface range GigabitEthernet3/1-48 // Sample Downlink ports
switchport trunk encapsulation dot1q
switchport private-vlan trunk native vlan 401
switchport private-vlan association trunk 200 201 // PVLANS secondaries as services
switchport private-vlan association trunk 300 301 // PVLANS secondaries as services
switchport mode private-vlan trunk // Private vlan isolated trunk
switchport port-security // Enable port security
vlan-range 201 // PVQoS and trunk port security
port-security maximum 3
  service-policy input P31_QoS // Ingress PVQoS for VLAN 201 (includes policing)
  service-policy output P31_QoS // Egress PVQoS for VLAN 201 (includes policing)
vlan range 202
port-security maximum 3
  service-policy input P32_QoS // Ingress PVQoS for VLAN 202 (includes policing)
  service-policy output P32_QoS // Egress PVQoS for VLAN 202 (includes policing)
spanning-tree portfast trunk

interface range tengigabitethernet1/1-2 // Uplink ports
switchport mode private-vlan trunk promiscuous // PVLAN promiscuous trunks

```

FIGURE 4 Sample configuration for ingress/egress policing, trusting DSCP, and giving precedence to voice packets on a Cisco Catalyst 4500 Series Switch.

PVLANS are extremely useful in a Metro Ethernet environment because they automatically provide isolation between multiple DSLAMs. PVLAN isolated trunks are used to multiplex several VLANs on the same port while still maintaining isolation between subscribers. The feature also allows a customer to subscribe to multiple ISPs with transparent networks. The PVLAN promiscuous trunks (2HCY2005) are used to carry services for thousands of subscribers on the switch's uplink ports. Prior to promiscuous trunks, a Catalyst switch could only carry one VLAN on a promiscuous port, thus requiring a greater number of physical ports. With this new feature, many primary PVLANS can be multiplexed onto one or both (resilient and load sharing) 10 Gigabit Ethernet or Gigabit Ethernet uplinks. The PVLAN promiscuous trunks on the Catalyst 4500 connect to the Cisco 7600 Series Router where IPv4, IPv6, or Multiprotocol Label Switching (MPLS) services are performed.

Trunk port security is also supported on PVLAN trunk ports. It restricts the allowed MAC addresses or the maximum number of MAC addresses to individual VLANs on a trunk port. It restricts the trunk port to configured MAC addresses so no other MAC address can join the network. When a trunk port security violation occurs, the trunk port is shut down and a Simple Network Management Protocol

(SNMP) trap might be generated. Trunk port security can be used when a Catalyst switch has an 802.1q or ISL trunk attached to a neighboring Layer 2 switch or DSLAM.

Per port per VLAN QoS allows network managers to create their own service policy per VLAN. This policy, performed in hardware, might consist of ingress and

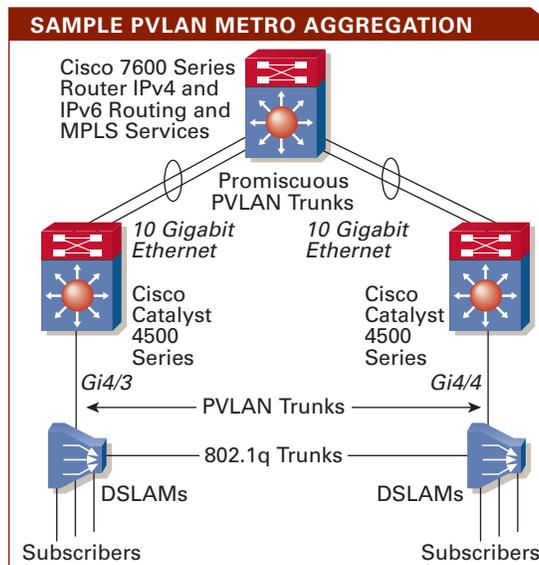


FIGURE 3 Security features, such as PVLAN, on Cisco Catalyst modular switches allow network managers to dynamically control security threats at their inception.

FIGURE 5 Sample configuration for DHCP Snooping on a Cisco Catalyst 4500 Series Switch.

```
Switch#show ip dhcp snooping binding interface Gi4/1
```

MacAddress	IpAddress	Lease(sec)	Type	VLAN	Interface
00:09:6B:50:B8:28	10.33.235.45	131585	dhcp-snooping	201	GigabitEthernet4/1
00:02:B9:A7:55:A5	10.33.232.47	439124	dhcp-snooping	200	GigabitEthernet4/1

FIGURE 6 Sample configuration for Dynamic ARP Inspection and IP Source Guard on a Cisco Catalyst 4500 Series Switch.

```
ip dhcp snooping allow-untrusted // Option 82 Passthrough
ip dhcp snooping // Enabling dhcp snooping and activating it on vlans 2-10
ip dhcp snooping vlan 2-10
ip arp inspection vlan 2-10 // Enabling Dynamic ARP Inspection on vlans 2-10
interface range gi2/1 - 48 // DHCP and ARP DoS attack rate limiters (in pps)
ip dhcp snooping limit rate 200
ip arp inspection limit rate 200
ip verify source vlan dhcp-snooping port-security // IP source guard
```

egress policing, trusting Differentiated Services Code Point (DSCP), or giving precedence to voice packets over data. Figure 4 shows a sample configuration for these three features on a Cisco Catalyst 4500 Series Switch.

DHCP interface tracking, or Option 82, satisfies the legal requirements of many countries, which stipulate that DSLAMs constantly track the DHCP offers and releases. DHCP interface tracking only provides a tracking path for the DHCP packet but does not enforce security. The DSLAM inserts information about itself in the DHCP request packet traversing from a client to a server. When a Catalyst switch is used in aggregation mode, it cannot change the Option 82 coming from DSLAMs, forcing the port to be trusted. Trusting the port rendered the industry-leading DHCP Snooping, Dynamic ARP Inspection, and IP Source Guard security features to be unavailable for Metro Ethernet. Today, Cisco has re-engineered the Catalyst switch for Option 82 passthrough, which means that the switch transparently passes Option 82, enabling deeper inspection of the DHCP packets.

DHCP Snooping combats rogue DHCP servers while protecting the network from DoS attacks. It achieves this by rate limiting the incoming DHCP packets and limiting client-facing ports for sending DHCP request and renew traffic only. The edge ports, for example, cannot offer DHCP lease, which is a function for the DHCP server. DHCP Snooping also forms the basis for other security features such as IP Source Guard and Dynamic ARP Inspection. This feature allows the switch to “snoop” the switching traffic for DHCP packets and create a dynamic binding (see Figure 5).

Dynamic ARP Inspection uses the DHCP Snooping bindings to prevent ARP spoofing and man-in-the-middle attacks for both static and dynamic IP addresses. Any violating hosts can be logged and the ports error-disabled until an administrative action is taken. IP Source Guard mitigates IP address spoofing by dynamically maintaining per port VLAN ACLs. IP Source Guard adds security to IP source address using DHCP Snooping table. The feature automatically locks an IP and MAC address to a given port. The dynamic ACL is removed when the user releases the IP address, for example, with “ipconfig /release.” Figure 6 shows a sample configuration for these features.

All the Cisco Catalyst security and QoS features discussed in this article build on one another—much like a set of security “stairs” upon which all services are deployed concurrently. These security features empower service providers with resilient, high-performance, reliable, and secure metro Layer 2+ networks. The switches are not only built for the needs of today’s networks but are well equipped to meet the demands and challenges of tomorrow. ■

FURTHER READING

- Cisco Catalyst 4500 Series Supervisor Engine V-10GE cisco.com/packet/172_8a1
- Cisco Metro Ethernet Switching Solution for Service Providers cisco.com/packet/172_8a2

The Service Exchange Framework

Mastering services requires comprehending and controlling every packet and policy in your network. Here's how to do it.

By Janet Kreiling

Next-generation networks (NGN) are rife with potential: personalized services, interoperating applications, keeping a connection from home to office and from PDA to PC to phone. But to a large extent, this potential has yet to be realized with intelligent networks that provide the subscriber and application awareness required to effectively converge the plethora of residential and business services available today. Consumers have access to a wide array of communications, entertainment, and online services, usually from a wide variety of broadband service providers. These services traverse multiple network types, each with their own unique capabilities, and they originate and terminate on many different devices. Most service providers do not yet have the capability to converge these services, but a few are beginning to implement new levels of network intelligence that will dramatically change the way we work, play, and learn.

For example, Sprint is using presence and location information from the Cisco Call-Session Control Platform (CSCP) to give its subscribers “push-to-talk” service. The CSCP tracks which people on the user’s talk-to list are present on the network. Plala Networks, a subsidiary of NTT East Corp., is using the Cisco Service Control Engine (SCE) to monitor and manage bandwidth-gobbling peer-to-peer traffic so other users can receive acceptable service. The Cisco SCE detects peer-to-peer file transfers; then policy servers limit the bandwidth available to this type of traffic.

These uses are just the beginning. For a next step, cellular providers might evolve push-to-talk services so subscribers could push to access a voice-enabled portal offering applications such as Mapquest. Directions could be spoken or displayed. A provider monitoring and limiting bandwidth could determine, for example, that demand is less during certain hours of the day and offer a managed service package for small businesses with guaranteed bandwidth and quality of service (QoS).

Service Exchange Framework Building Blocks

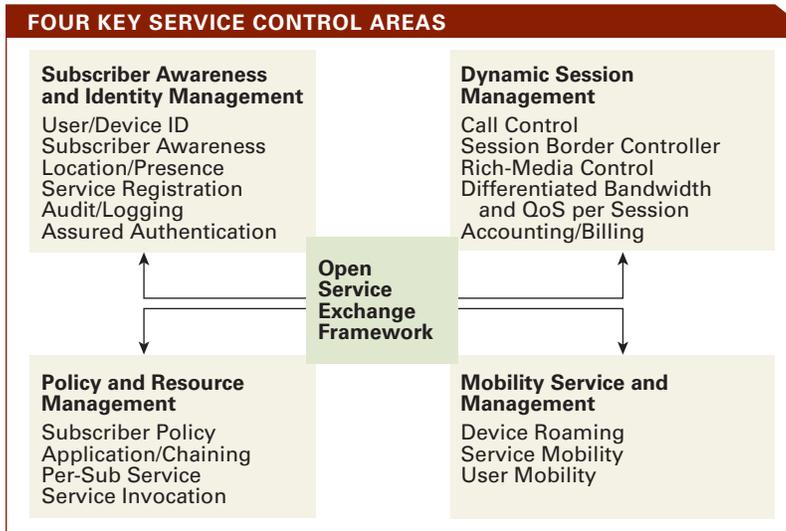
Imaginative, wide-ranging, useful services such as these are possible now. Every required network element—hardware and software, core and edge—is available. And only Cisco offers them all, in the *Service Exchange Framework (SEF)*, which can be customized

to support a provider’s unique next-generation business model. The SEF provides the intelligence that enables service convergence—a fundamental element of IP-based next-generation networks (see *Packet* First Quarter 2005, cisco.com/packet/172_8b1).

“The SEF enables service providers to analyze, optimize, secure, and meter application- and content-based services in their existing IP networks,” says Thomas Barnett, Jr., senior service provider marketing manager at Cisco. “That’s every capability they need to maximize and deliver services and applications in an IP-based next-generation network.” Right now, he adds, “Most providers have very little of the detailed information or control they need to deliver NGN services. They may not know what device a subscriber is using, where the subscriber is located, what applications he or she should have access to, or what policies such as content or bandwidth control, should be imposed. So the carrier’s ability to deliver services and applications, such as e-mail, instant messaging, voice, gaming, and video, with the desired quality of service and class of service [CoS]—in a way that maximizes network efficiency and the user experience—is limited.”

The following four capabilities are crucial to delivering next-generation service control (see figure):

- *Subscriber awareness and identity management*—The network can identify users and their devices, determine a person’s location, and establish presence for that person, including sharing his or her status (on or off network) with other subscribers. With these abilities, providers can deploy presence-based services such as push-to-talk, instant messaging, voice, and video, call routing and screening, and 3G+ mobile applications such as streaming audio and video and interactive gaming.
- *Policy and resource management*—The provider gains programmable session and policy control including authentication, authorization and accounting (AAA), QoS, and VPN routing and forwarding for all services the customer subscribes to, along with one very tangible subscriber benefit: signing on once with only one password for all applications.



NGN SERVICE CONTROL

Four areas of service control help providers optimize network resources, deliver customized services, and ensure reliable service delivery and performance.

- *Dynamic session management*—End-to-end attributes of individual sessions can be managed in real time using techniques such as queuing, policing, traffic shaping, and packet marking, which ensures contracted-for QoS and accurate billing. Internetwork signaling with border control allows services to follow a subscriber from a wireline to a wireless network, a circuit-switched to an IP network, or an enterprise to a public network.
- *Mobility service and management*—Presence, name space, subscriber data, and service integration—the four information elements central to a mobile call—can all be managed centrally. When tasks such as storing customer data, AAA, provisioning, and service interaction are consolidated, it becomes possible to create new services for mobile networks or for calls moving onto mobile networks quickly with maximum re-use of assets and minimal expense.

Says Barnett, “These capabilities help service providers optimize network resources, deliver customized services, enhance the subscriber experience, and ensure reliable service delivery and performance.”

Cisco Solutions for the Service Exchange Framework

Cisco provides a variety of solutions to deliver these subscriber- and service-aware capabilities for next-generation service control. Because these are stand-alone, inline hardware and software solutions that do not duplicate or interfere with any essential networking functions (such as access or aggregation), service providers can deploy this critical service layer in a phased approach that meets their business and service needs.

The Cisco Mobile Exchange (CMX) is an open platform that provides an intelligent enforcement layer

within the operator’s network and easily interfaces with all of the control elements in the IP network. CMX has proven interoperability with all major radio access network (RAN), AAA, content billing, and content filtering and compression vendors. It provides three primary functions—access and service control, seamless mobility, and deep-packet inspection. CMX also leverages Cisco IOS Software and enables mobile operators to deliver secure, profitable mobile services. The cornerstone of the CMX framework is the Cisco 7600 Series Router. Modular blades deployed in the Cisco 7600 Series provide a variety of CMX capabilities, including packet gateways, mobile services, load balancing, network management, and operations. Together, these components successfully solve the many challenges that face mobile operators seeking profitability from their second-generation (2G), 2.5G, 3G, or 4G mobile packet infrastructures and IEEE 802.11 wireless LANs.

The aforementioned CSCP is designed for broadband operators interested in deploying a robust, carrier-grade, next-generation services delivery environment for providing multimedia applications and services over their IP-based networks. The CSCP consists of the Cisco Service Engine, Cisco Edge Proxy, and the Cisco Name Resolution Server—software solutions all based on Session Initiation Protocol (SIP), the Internet Engineering Task Force (IETF) standard that defines peer-to-peer, multimedia signaling. With the CSCP, operators can offer a differentiated multimedia communication experience to their subscribers with services integrating voice, video, push to talk, presence, geolocation, “buddy” lists, and more.

The Cisco Service Control solution allows providers to take full advantage of their existing IP network infrastructure to differentiate between services such as voice over IP (VoIP), Web browsing, music downloads, video streaming, and peer-to-peer traffic. The Cisco SCE adds a programmable service layer to networks, allowing operators to identify subscribers, classify applications, guarantee service performance, and charge for multiple IP services. Performing Layer 7 stateful, deep-packet inspection at multigigabit speeds, the Cisco SCE is transport- and content-independent, fully extensible and programmable, and easily integrates into existing network fabrics. This network element resides “in traffic” behind an IP aggregation point and enables the network to differentiate between individual services—enabling providers to manage and bill for premium services running on common transport.

Phased Approach to Service Control and Network Intelligence

The starting point for an IP-based NGN is the provider’s existing network. “A provider can start with a specific service, add the solution or solutions required for it, and then build more services onto that

system or add other solutions for other services,” Barnett says. “This is a custom process, driven by each provider’s business plan and customer needs. Providers can deploy solutions to offer services that will differentiate them in their own market.”

SEF solutions from Cisco give service providers a clear path for cost-effective, low-risk rollouts of multimedia applications and help to address operations, administration, management, and provisioning (OAM&P) challenges, adds Barnett. Triple-play (voice, video, and data) services, voice over broadband, and presence-enabled communications are a few examples. A multitude of services can be initiated and expanded using SEF solutions.

Copel Telecommunications, based in Brazil, uses the Cisco SCE deep-packet inspection capability to identify and classify applications and guarantee performance of offerings such as VoIP, Internet access, and multimedia services to enterprise clients. Copel will be able to implement traffic-line intelligence that addresses new standards, protocols, billing techniques, and classification of content type in real time, and can thus customize services to satisfy the needs of individual customers. A provider with such a wide-ranging application of the Cisco Service Control Engine’s abilities might next enable tiered services by adding a policy control solution.

Swisscom Mobile is already using CMX for its “Mobility Unlimited” solution, which allows subscribers to access data from the fastest transmission technology available wherever they are and change to another transmission mode seamlessly while on the move.

An evolution plan for a provider with this starting point might include adding deep-packet inspection to monitor and manage data applications, or policy control to ensure the security of a user’s corporate Intranet and applications while employees travel. With SEF solutions, providers could offer customers a “turbo button” for real-time bandwidth changes, or enable them to cross from a mobile to a fixed network while maintaining the same call. Such a capability would make the provider’s service very attractive to third-party developers of high-bandwidth multimedia applications, Barnett says, and create opportunities for partnering.

Service Optimization

In addition to gracefully evolving to next-gen offerings, providers can use the SEF to finally understand down to the packet what’s going on in their networks—how much bandwidth is consumed by given services, applications, and users. They can

determine how much bandwidth and what QoS to devote to a given user and application. They can see where the inefficiencies are in how their resources are used, and where efficiencies can lead to new service and revenue opportunities.

For example, providers can reduce costs by monitoring and regulating how much capacity is used by bandwidth-hungry applications, and limit them or permit dynamic bandwidth allocation through policy control. “Taking the guesswork out of capacity planning and detailing subscriber demographics helps operators uncover hidden operational costs and new revenue potential in wireline, mobile, and cable networks,” Barnett points out.

In a truly intelligent network, subscribers will be able to sign themselves up for services and service attributes. Examples include subscribing to a cell phone plan, VoIP, or pay-per-view; ordering bandwidth or QoS that varies by time of day; and imposing parental or employer controls. Each such interaction is one fewer that is handled by a customer service representative, for further cost savings. Subscribers have proven willing to pay more for a more personalized experience—getting any service they want, when they want it, over any device—so the IP-based NGN also offers income potential even beyond revenues from services and ventures with application developers.

As providers continue to face competitive pressures to cut costs and provide more flexible services, they are “already creating multiservice converged networks by eliminating networks dedicated to specific types of traffic and applications, or reducing layers within networks,” says Barnett. “IP is the basis for the sweeping transformation we’re seeing in networks, and the SEF is the way to master profitable, innovative service delivery over the IP next-generation network.” ■

Read the full white paper, *The Service Exchange Framework: Providing Greater Control for Cisco IP Next-Generation Networks*, at cisco.com/packet/172_8b2.

FURTHER READING

- White paper: *The Cisco Call Session Control Platform*
cisco.com/packet/172_8b3
- White paper: *Bridging the Infrastructure Gap: The Importance of Service Control in Broadband Networks*
cisco.com/packet/172_8b4
- White paper: *Cisco and the Service Provider IP Next-Generation Network Journey*
cisco.com/packet/172_8b5

IP/MPLS Interprovider

Extending Network Infrastructures and Services Beyond Administrative Boundaries

By Santiago Alvarez

Service provider customers increasingly demand networking services that meet the geographical requirements of their business. They expect a single point of contact for these services instead of approaching multiple providers for a solution. In many cases, service providers find it more cost effective to extend their footprint through agreements with other providers than to invest in new network infrastructure. In other cases, regulatory or political reasons might require service providers to rely on other providers to meet customer requirements. Some providers using Multiprotocol Label Switching (MPLS) have implemented these agreements for some time; however, most service providers now have come to realize the inevitable need for such accords to compete effectively in their markets.

Cisco MPLS offers service providers the technology to extend their MPLS networks and the coverage of their services. This technology includes support for implementing MPLS VPN, multicast VPN, MPLS traffic engineering (TE), and Layer 2 VPN (L2VPN) across multiple autonomous systems (AS). An AS typically defines a service provider's administrative boundary. However, some providers might use multiple autonomous systems in their networks. In those cases, the *Cisco IP/MPLS Interprovider solution* also

helps expand services throughout multiple networks belonging to a single service provider.

MPLS VPN Inter-AS and Carrier Supporting Carrier

Cisco introduced support for MPLS VPN Inter-AS and MPLS VPN Carrier Supporting Carrier (CSC) in 2001. MPLS Inter-AS is an extension of MPLS VPN that enables peering agreements between two or more autonomous systems to offer IP VPN services. MPLS VPN CSC is an MPLS feature that allows a service provider to act as a transport network for other MPLS networks. The transport network offers an MPLS service that can carry multiple services (e.g., Internet, IP VPN, L2VPN). The presence of multiple providers is transparent to the end customer for both MPLS VPN Inter-AS and CSC.

Cisco has made numerous enhancements to these solutions since their introduction. The most recent improvement extends load balancing capabilities allowing AS boundary routers (ASBRs) to maintain a single external Border Gateway Protocol (eBGP) session using loop-back peering across multiple physical links.

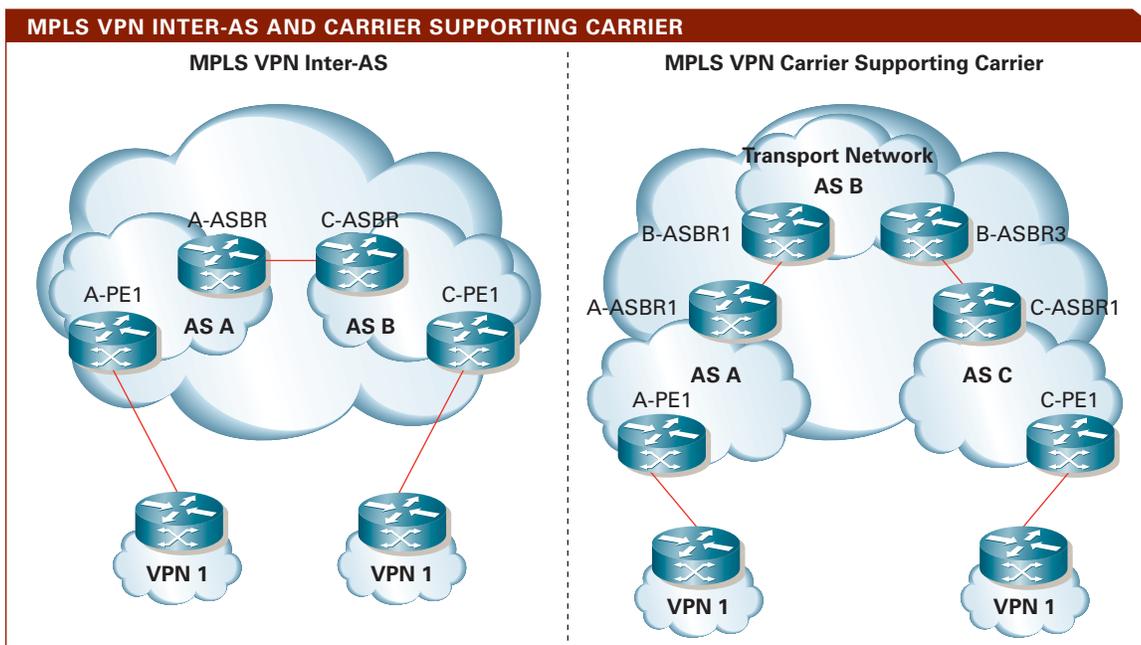


FIGURE 1 Cisco MPLS VPN Inter-AS and CSC can be used across IP and MPLS networks, providing greater flexibility for interprovider service agreements or for service providers that have multiple networks.

FIGURE 2 With Cisco Inter-AS mVPN, two multicast distribution trees are created between two separate autonomous systems.

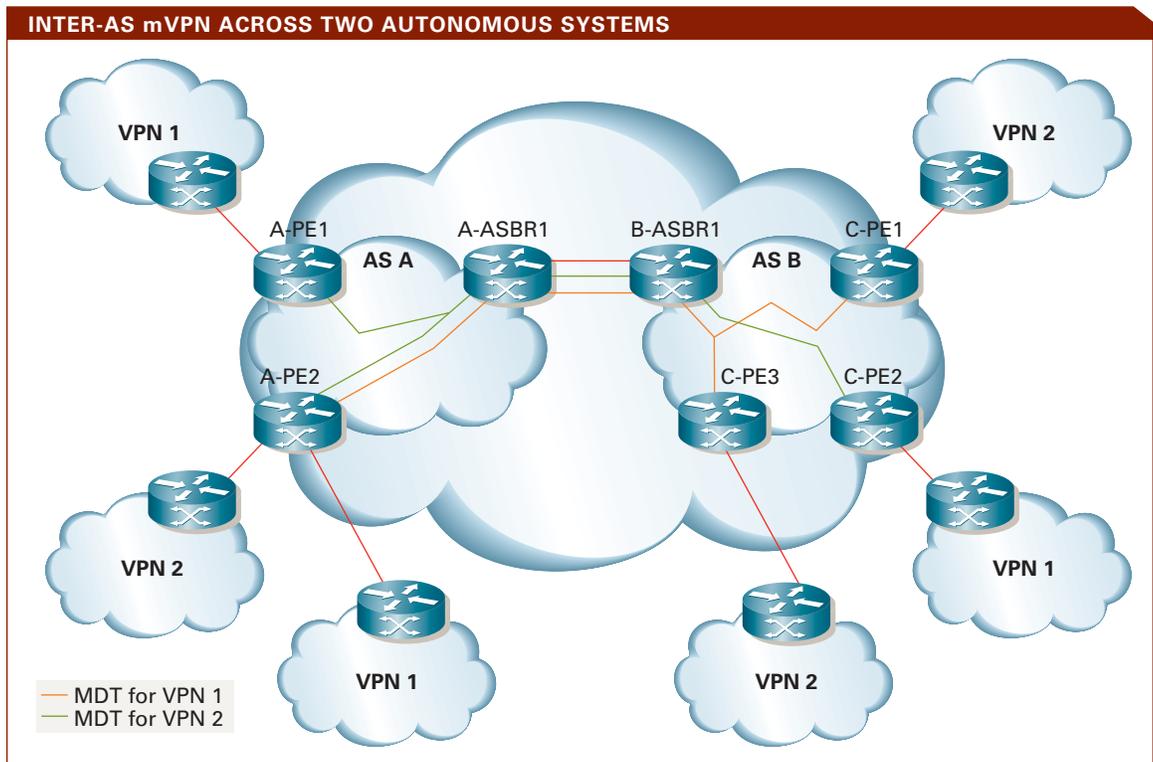


Figure 1 illustrates the configuration of MPLS VPN Inter-AS and CSC. In the case of MPLS VPN Inter-AS, two autonomous systems A and B peer to provide VPN service to two separate VPN sites. This technology offers three different peering alternatives that provide different levels of control and scalability. In the other scenario, AS B provides an MPLS transport service to AS A and AS C. CSC defines a hierarchical VPN model where B sees A and C as VPN sites with a single signaling relationship regardless of the number of VPN instances and VPN sites in A and C.

Interprovider MPLS VPN over IP

Interprovider MPLS VPN over IP is an extension of the inter-AS and CSC functions in MPLS VPN to include IP networks that are not MPLS-enabled. In 2004, Cisco introduced support for MPLS VPN over IP using Layer 2 Tunneling Protocol version 3 (L2TPv3) encapsulation. This encapsulation provides intrinsic spoofing protection for VPN traffic. The Cisco MPLS VPN Inter-AS and CSC features can now be used across both IP and MPLS networks, providing greater flexibility for interprovider service agreements or for providers that have multiple networks



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

(some of which are not enabled with MPLS). MPLS VPN Inter-AS and CSC with IP networks was introduced in Cisco IOS Software Release 12.0(30)S.

In Figure 1, one or even both of the autonomous systems in the inter-AS solution could use an IP backbone. In the CSC scenario shown in Figure 1, any network can be an IP or MPLS network as long as the customer autonomous systems A and C are of the same type.

Inter-AS Multicast VPN

Service providers can now establish agreements to expand their multicast VPN (mVPN) services. Multicast distribution trees (MDTs) can be created between two autonomous systems without having to exchange additional unicast routing information. Provider edge (PE) routers can create trees triggered by Protocol Independent Multicast (PIM) joins across AS boundaries. Intermediate hops can perform the Reverse Path Forwarding (RPF) check even if unicast reachability cannot be verified directly across autonomous systems. Ultimately, the VPN customer receives a service experience as if a single AS or service provider were present.

Cisco Inter-AS mVPN introduces an extension to PIM and a new address family for BGP. A new PIM join message encoding includes the exit point (called an *RPF vector*) to the other AS together with the source. The PE router that originates the join message has learned the exit point and MDT group address using BGP.

Cisco Inter-AS mVPN uses a new BGP sub-address family identifier (SAFI) to distribute MDT information.

Intermediate routers that run BGP select an RPF interface by conducting a direct lookup in a special BGP MDT table. Intermediate routers that do not run BGP use the RPF vector in the PIM message to find the RPF interface. Figure 2 shows two MDTs created between two separate autonomous systems. The Inter-AS mVPN feature was introduced in Cisco IOS Software Release 12.0(30)S (including PIM RPF vector).

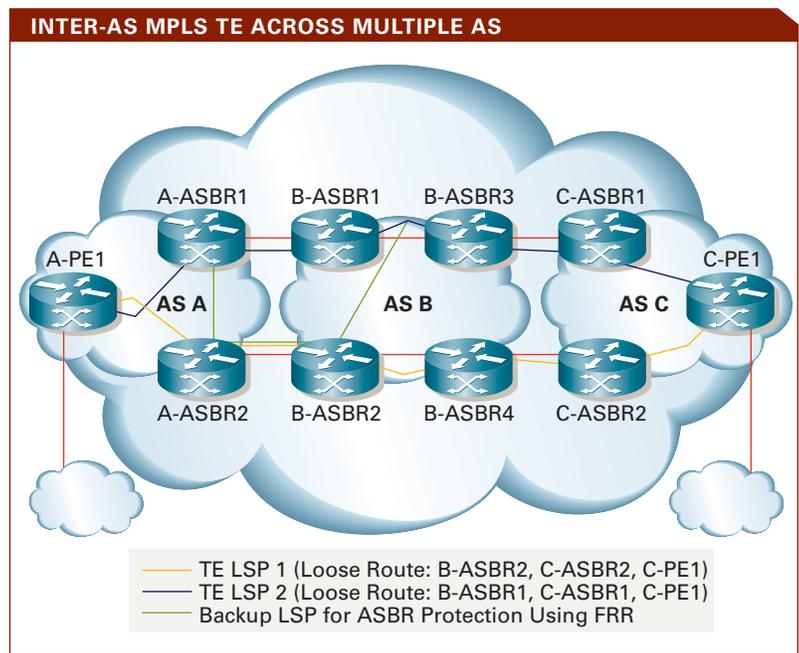
Inter-AS Traffic Engineering

Inter-AS TE is an extension of MPLS TE that enables the creation of traffic-engineered label switched paths (LSPs) across AS boundaries. Historically, MPLS TE technology has been used within a single AS, in part because it is generally deployed closely tied to a link-state Interior Gateway Protocol (IGP). With Inter-AS TE, each AS can have its own IGP, and TE LSPs can be created across AS boundaries and still take advantage of the constrained-based routing, admission control, explicit routing, and protection capabilities of MPLS TE. These TE LSPs can be used to extend MPLS VPN and L2VPN services.

Figure 3 shows a sample implementation of Inter-AS TE. AS A signals two TE LSPs toward AS C. The LSP headend relies on loose routing for end-to-end path computation. A list of loose path entries specifies ingress points to all autonomous systems. Those devices are responsible for path computation within their AS and signaling to the next entry in the path. Using the Cisco Inter-AS TE feature, ASBR devices can enable a secure interface using Resource Reservation Protocol (RSVP) authentication and local policy control to prevent unauthorized access. Network availability is also enhanced in this scenario using MPLS TE Fast Reroute (FRR). A backup LSP between AS A and AS B protects against the failure of one of the ASBR devices in AS B. FRR could be deployed at different points to provide node, link, and shared-risk link group protection. The Inter-AS TE feature was introduced in Cisco IOS Software Release 12.0(29)S.

Interprovider L2VPN

Service providers can extend L2VPN services across multiple autonomous systems. Cisco L2VPN pseudowire switching allows network administrators to establish a pseudowire to a peering device that interfaces with other autonomous systems. This model increases scalability, reduces peering relationships, and can be extended to support disparate L2VPN control planes (e.g., MPLS and L2TPv3). The previous alternative approach for interprovider L2VPN relies on a single end-to-end pseudowire being signaled across multiple autonomous systems. In that case, an LSP connects the two PE devices, and they must be able to establish a directed LDP session to set up the pseudowire. MPLS VPN Inter-AS CSC and Inter-AS TE provide configurations that satisfy these requirements. Cisco IOS Software Release 12.0(31)S introduced support for L2VPN pseudowire switching.



Future Work

QoS is an important area of future work in inter-provider environments. The current MPLS QoS capabilities enable end-to-end QoS across multiple networks. However, there are several aspects of interprovider QoS that can be enhanced. First, service providers can benefit from an agreement framework for traffic classes, their syntax and semantics. Second, better understanding is needed on how to budget application service-level agreements (SLAs) across service provider networks. In addition, control plane enhancements might be made to allow QoS signaling across AS boundaries. Some of these interprovider QoS components require the definition of frameworks with guidelines or best practices. Others might require protocol enhancements. The new developments should simplify the implementation of QoS peering agreements.

FIGURE 3 With Cisco Inter-AS Traffic Engineering, ASBR devices can enable a secure interface using RSVP authentication and local policy control to prevent unauthorized access.

With the Cisco IP/MPLS Interprovider solution, service providers can build agreements with other providers to extend their MPLS VPN, multicast VPN, TE, and L2VPN services. These interprovider capabilities are increasingly important as MPLS adoption grows worldwide and more customers expect a single point of contact for their networking services.

Cisco leads the industry in implementing and standardizing IP/MPLS interprovider technologies. Interprovider L2VPN and QoS are two areas of active discussion at standard bodies and forums where Cisco participates. ■

FURTHER READING

- Cisco IP/MPLS Interprovider Solution cisco.com/packet/172_8c1

Integrated Services Routers in the Small Office

Cisco extends integrated services routers with new models and integrated wireless across the portfolio.



PERFORMANCE PUNCH New models in the Cisco integrated services router line bring concurrent services, such as security and wireless LAN, at wire speed to small offices and enterprise teleworker sites.

By David Barry

Following up on the successful launch in September 2004 of the integrated services router line (see *Packet* magazine, Fourth Quarter 2004), Cisco recently announced availability of new models and capabilities to these routers that extend its powerful integration into small and remote offices. These models satisfy a market demand for platforms that deliver greater performance for deploying services such as security and wireless LAN (WLAN) capabilities to the enterprise branch and small and mid-sized businesses—platforms that are easy to deploy and cost effective to manage.

The new Cisco 800 and 1800 Series integrated services routers offer concurrent services including firewall, virtual private networks (VPNs), and WLANs at an attractive price point for small offices. These Cisco IOS-based platforms also deliver centralized management features that make them ideal for small office or teleworker sites as part of an enterprise or service provider network.

“These new models extend the benefits of the [integrated services router] line to small offices for both enterprise and small and medium-sized business customers,” says Marc Bresniker, product manager in the Premises Communications Business Unit at Cisco. “Their strong performance allows businesses to layer on new services, such as security, QoS [quality of service] for voice, or wireless LANs, while taking full advantage of DSL and cable broadband speeds. And for those customers who want to use the newer and faster DSL standards, such as ADSL2+ [more than 20-Mbit/s downstream speeds] and multipair symmetric DSL [G.SHDSL], the DSL models in the portfolio will support these new standards.”

Cisco 1800 Series: Greater Performance and Services Integration
The most highly integrated of the fixed-configuration integrated services routers is the Cisco 1800 Series. These models include a full suite of advanced security features: firewall, IP Security (IPSec) VPNs, and support for intrusion prevention and Cisco

WLAN Capabilities Added to Integrated Services Routers

Cisco's integrated services routers now offer WLAN capabilities across the entire portfolio. The recently launched fixed-configuration 800 and 1800 Series integrated services routers include 11 factory-configured wireless models with antennas. Also included in this launch is a high-speed wireless interface card (HWIC) for the modular integrated services router platforms. Installing the HWIC into a slot on the Cisco 1841 or the Cisco 2800 or 3800 Series routers enables businesses to integrate a wireless access point onto their access router.

"Most exciting about the new wireless capabilities is that they further build on the compelling premise of the integrated services router line—delivering secure data, voice, and video services to wired and wireless users to maximize productivity," says Sunny Mahant, product marketing manager in the Multiservice Customer Edge Business Unit at Cisco. "Not only can these models run multiple services such as security, voice, and VPNs without degrading broadband connections, now they can also run wireless, all from one integrated platform."

Offices that need to support survivable IEEE 802.1X local authentication can combine a modular integrated services router with a wireless HWIC or with several Cisco Aironet access points. "This allows the router to act as a local authentication server to authenticate wireless clients when the AAA [authentication, authorization, and accounting] server is not available," says Mahant.

Fixed and Modular Access Points

With support for 802.11 Wi-Fi Certified Access on both the fixed and modular integrated services routers, Cisco provides a low-cost entry point for companies that want to add WLAN connectivity to their branch or small office. These routers eliminate the requirement for dedicated wireless appliances at each site when only one access point is needed—simplifying wireless access deployment and management. No changes are required to the existing wired infrastructure.

Network Admission Control (NAC) for security policy control and protection against viruses and worms. Also included are models with options for fully integrated ISDN BRI, an analog modem, and dual Fast Ethernet ports for redundant WAN links and load balancing.

With the increased availability and affordability of broadband DSL and cable, some companies are looking to use dual broadband WAN ports; for example, they are contracting with separate broadband providers to ensure automatic failover if either service experiences congestion or failure. An integrated internal power supply on the Cisco 1800 Series also makes it easy to deploy with fewer cords to set up.

The 1800 Series models with an integrated 8-port switch are targeted for small offices. With support for advanced QoS and multiple virtual LANs (VLANs), businesses can configure and segment their network for application performance and security.

Further integration of the Cisco 1800 Series is achieved with an option for Power over Ethernet (PoE) support. PoE is especially beneficial for companies that deploy a Cisco 1800 Series model with

IP phones or external wireless access points and want to eliminate the need for separate power supplies for those devices.

The Cisco 1800 Series models include an option for integrated wireless access points, providing secure WLAN services in a single device—and helping businesses reduce their total cost of ownership with simplified WLAN deployment and management capabilities while maintaining network security. The integrated wireless access point can support IEEE 802.11b/g and 802.11a simultaneously to provide added flexibility in high-speed wireless applications. The removable, replaceable antennas allow choices for mounting in different locations to place wireless coverage where needed. For instance, a retail store that deploys a Cisco 1800 Series Integrated Services Router behind the front counter or in a utility room can mount the antennas elsewhere for better wireless coverage.

Cisco 800 Series: Small but Powerful

The Cisco 800 Series has several models targeted for small remote offices and teleworkers, each providing a cost-effective solution for delivering secure WAN

Both the fixed and modular integrated access points deliver robust, predictable 802.11 wireless coverage with strong radio sensitivity and superior performance, notes Mahant. They support Wi-Fi Protected Access (WPA) for per-user IEEE 802.1X mutual authentication with an Extensible Authentication Protocol (EAP) such as Cisco LEAP. They also support 802.11e for QoS, Wi-Fi Multimedia (WMM), VLANs, and multiple service set identifiers (SSID).

Scaling Wireless and Deploying Advanced Services with Aironet Access Points

In situations where only a single access point is needed, businesses gain the full benefits of integration and cost effectiveness by choosing a fixed-configuration integrated services router with a built-in access point or by installing the HWIC into a modular model. Depending on the integrated services router and whether it operates in single mode (802.11 b/g) or dual mode (802.11 a/b/g), the router will support up to 20 or 50 users, respectively.

For companies or sites that require more than one access point either immediately or in the future, Cisco Aironet access points are recommended, says Mahant. The modular integrated services router platforms presently support either the Cisco Aironet access points or the modular platform access points with HWIC, not both options.

Cisco Aironet access points deliver high security with WPA2 and high-capacity wireless access for offices and challenging RF environments. These robust access points are perfect for single access point deployments that require flexible, secure installation options, or for enterprise deployments that require more than one access point.

To learn more about new wireless capabilities of the Cisco integrated services routers, visit cisco.com/go/isr.

connectivity with optional integrated IEEE 802.11b/g for WLANs in a single device (see sidebar, “WLAN Capabilities Added to Integrated Services Routers”). In addition, the Cisco 800 Series is easy to set up and deploy using the Web-based configuration tool, Cisco Router and Security Device Manager (SDM)—ideal for small offices with minimal local technical resources.

The Cisco 870 Series includes hardware-assisted encryption for VPNs. Integrated security features are further enhanced with support for intrusion prevention and Cisco NAC for security policy control and virus and worm protection. Each model also has 802.11b/g WLAN capabilities with removable, replaceable dual diversity antennas.

The Cisco 870 Series offers advanced QoS support which, along with its increased performance for encryption, makes it ideal for teleworker or remote call agent applications. Users can connect an IP phone to the router’s switch port to act as an enterprise extension and give voice traffic precedence over data applications.

The Cisco 850 Series, with four 10/100 Mbit/s ports and 10/100 Fast Ethernet or ADSL connections, supports up to 10 users and offers a basic set of security features, including stateful inspection firewall and

VPN encryption. Each model has an option for integrated wireless, the Cisco 851W and Cisco 857W, and come equipped with a single, fixed antenna and 802.11b/g WLAN support.



Affordable broadband access is changing the way businesses communicate with customers, suppliers, and employees. WLANs can further extend the effectiveness of business applications. To take advantage of these high-speed connections, small offices must have the same level of security enjoyed by their larger counterparts. With its new line of fixed-configuration integrated services routers, Cisco is delivering the right combination of integrated services with the performance punch small offices need. ■

FURTHER READING

- Cisco Integrated Services Routers home page cisco.com/go/isr

Buying Strategies

Initially designed as a way to keep hazardous materials out of landfills, global environmental regulations—including the Restriction on the Use of Certain Hazardous Substances and the Waste Electrical and Electronic Equipment Directive—have sparked a significant trend: More than ever, used hardware products are being refurbished and resold to other businesses, often at a fraction of their original cost.

Major PC vendors such as IBM, HP, and Dell are providing environmentally sound disposal of electronic devices and offering an array of solutions to resell legacy equipment, according to Michael Burlison, an analyst with META Group.

Buying refurbished gear is a good way for SMBs to save money, but there are some important issues to consider if you want to ensure that you won't end up paying more in the long run. The quality and condition of used products available for purchase on the secondary market varies considerably and is unpredictable, so be sure you deal with a reputable vendor—ideally, the

company that manufactured the products in the first place, or one of its authorized resellers. Products are available from other sources, too, but purchasing equipment from these providers involves a number of limitations and risks, including the following:

Many used products turn out to be “gray market” merchandise—new merchandise that is being distributed illegally, without the authorization of the manufacturer. Other products turn out to be counterfeit.

Used products are typically not eligible for a service contract without a vendor-authorized inspection to verify their condition.

Software licenses are generally non-transferable and must be purchased from the appropriate vendors.

Learn about the Cisco Authorized Refurbished Equipment program at cisco.com/go/iq-refurb. ■

—David Baum

Digital Security

Financial institutions manage risk and regulatory compliance proactively, with Cisco Self-Defending Networks.

By Rhonda Raider

When financial institution CIOs lose sleep, information security is usually to blame. Peter Simonsen, CIO of Arizona State Savings and Credit Union, the largest state-chartered federally insured credit union in the state, explains why: “Only financial institutions convert their assets into zeroes and ones and store them on a hard drive, which is why digital security issues have risen to the top of the list of boardroom concerns.”

And money is not the only asset at risk from network security breaches. “At the highest level, security is about banks and credit unions protecting their most valuable asset, customer trust—their stock in trade,” adds Rune Olslund, a financial services market manager at Cisco. “If that trust is ever compromised or violated, the bank will have a hard time restoring its reputation.”

New Factors in Financial Information Security

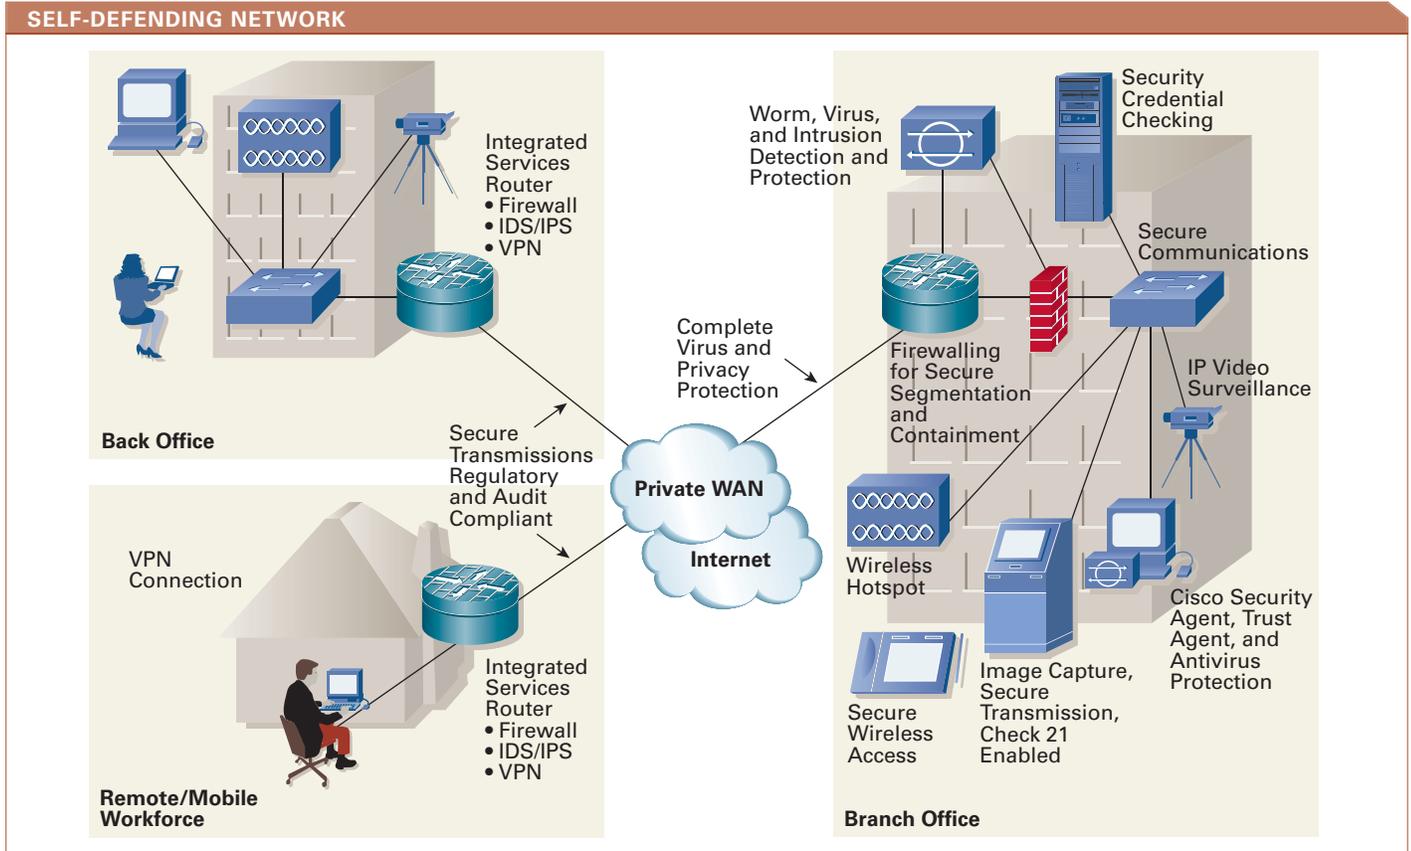
Several factors have converged to catapult information security to the top of priority lists for financial institution IT groups.

Changed Regulatory Landscape

Prudent information security and technology risk management are no longer just industry best practices, they’re mandated by a spate of security and privacy regulations including the Gramm-Leach-Bliley Act Data Protection regulation, Sarbanes-Oxley Act, BASEL II-Operational Risk, and the USA Patriot Act. “Financial institutions deserve credit for having monitored security and privacy without regulatory oversight for many years,” says Paul Reymann, co-author of the Gramm-Leach-Bliley Data Protection regulation and CEO of ReymannGroup, Inc. (reymanngroup.com), a team of leading subject matter experts on finance, healthcare, retail, and

DEFENSE IN DEPTH

Financial institutions can manage risk proactively through a technology infrastructure with an integrated suite of solutions for secure connectivity, threat defense, and trust and identity management.



manufacturing. “Many of today’s compliance requirements evolved from what had been industry best practices. Now, rather than being followed by a few leading organizations, they’re mandated for everybody.”

Greater Reliance on Internet Channels

Not so long ago, network security was intended to keep out outsiders. “Now financial institutions offer a slew of self-service products that allow customers behind the firewall, including Internet banking, online bill pay, and electronic statements,” says Simonsen. “Electronic delivery channels are required to succeed in the marketplace, which makes digital security the forefront issue for CIOs of financial institutions.”

Widespread Publicity

Recent high-profile privacy breaches have sharpened public interest in information security. “Our members are very concerned about identity theft,” says Ray Carsey, assistant vice president of technology for Mountain America Credit Union, Utah’s second-largest credit union, with branches in four states. “We do everything in our power to protect our members’ information. We know that if we don’t have a member’s trust, that member is going to go somewhere else.”

Greater Board Oversight

“Until a few years ago, financial institution board members would redirect all security concerns to the technology staff and say, ‘You deal with this,’” says Olslund. Now boards take a more hands-on approach because they’re held more accountable under many security and privacy regulations.

Recognition of Financial Institutions’ Role in National Security

The US government recognizes financial institutions as one of eight vertical industries critical to sustaining the way of life in the event of a doomsday scenario, according to Simonsen. “Imagine if you couldn’t get to your money,” he says. “That’s why there’s no such thing as too much attention to security.”

What’s Needed for Compliance and Security

To satisfy the new demands, financial institution IT groups face the following information security challenges:

- Protect the security and confidentiality of customers’ nonpublic personal information, the focus of the Gramm-Leach Bliley Data Privacy regulation.
- Institute administrative, technical, and physical safeguards against internal and external threats.
- Protect against viruses, worms, and Distributed Denial-of-Service (DDoS) attacks. “DDoS and hacker attacks are becoming more of a concern for small and medium financial institutions because more customers conduct business using Websites,” says Michael Payne, financial services enterprise manager at Cisco.

Regulatory Compliance

In the late 1990s, regulators had lengthy discussions about the dynamic threat environment. By the time the ink was dry, the threats would be changing. Therefore, we decided to write the rules to address the risks as we know them today, but also to create an environment where financial institutions would need to go through a discovery process to find a compliant solution rather than buying pre-packaged solutions that might quickly become obsolete. Today, the rules keep up with the changing environment because they’re written to be supplemented by new guidance from the Federal Financial Institutions Examination Council (FFIEC).

Common threads in most regulations are protecting customers’ nonpublic personal information; continuous risk management; and monitoring, auditing, and adjusting.

—Paul Reymann, CEO, ReymannGroup, Inc.

- Implement digital security defense-in-depth to address the issues of technology, people and processes. “You need to be able to prevent, protect, prosecute, and ultimately recover successfully from potential external attacks” says Simonsen.
- Establish continuous, risk-based information security policies with board oversight. This contrasts with the traditional project-based approach to risk management. “The infrastructure must continuously manage and monitor risk, and the technology needs to be proactive rather than reactive,” says Reymann.

Self-Defending Networks

While Cisco does not promote regulation of IT security, the company recognizes that regulatory compliance is an urgent business concern for its financial institution customers. The Cisco Self-Defending Network approach gives financial institutions the capabilities they need to meet the security and regulatory challenges they face in guarding against network invasions.

The premise of the Self-Defending Network is that with more threats, more advanced networks, and more regulations, financial institutions can no longer rely on a reactive approach to threat defense. Instead, they need a way to proactively manage regulatory compliance and risk. “New threats are being introduced in such high volumes and such a fast pace that ‘reactive mode’ is no longer practical,” says Reymann. “Financial institutions need to take a proactive stance, and many find they need help from partners—not just vendors, but partners—who are trusted and proven.”

Self-Defending Network Strategy

BANK NEED	STRATEGY	CISCO TECHNOLOGIES
SECURE CONNECTIVITY		
Protect data and voice confidentiality	Secure the transport	Virtual private networks (VPNs) based on IPSec and Secure Sockets Layer (SSL)
THREAT DEFENSE		
Detect and prevent external attacks	Defend the network edge	Router-integrated security Firewalls Intrusion detection systems (IDS) Intrusion prevention systems (IPS) AutoSecure
Protect against internal attacks	Protect the interior	Cisco Catalyst switch-integrated security Firewalls IDS Content engines
Protect hosts against infection	Guard endpoints	Cisco Security Agent
TRUST AND IDENTITY MANAGEMENT		
Control access to network by individuals and devices	Verify user and device against access policies	Cisco Identity-Based Networking System (IBNS) Cisco Secure Access Control Server (ACS) Cisco Network Access Control (NAC)

Developing a Self-Defending Network requires attention to technology, processes, and people (see sidebar on page 78). Cisco provides the technology infrastructure with an integrated suite of solutions for secure connectivity, threat defense, and trust and identity management (see table on this page).

Integrated Security

Unlike point security solutions, the Cisco Self-Defending Network relies on integrated security, applied on multiple layers throughout the network. “In the old paradigm, an intrusion detection system would spot traffic that shouldn’t be there, but then a human had to do something about it,” says Olslund. “Because Cisco solutions are integrated, components can communicate with each other to take action much faster and more effectively.” The result is a more in-depth defense.

“The network has to be designed with integrated security—not security as a point solution or afterthought,” says Simonsen. “Defense in-depth has to be

comprehensive so that if one system is compromised, it doesn’t affect the rest of your business.”

Case in point: Mountain America Credit Union integrates its internal and external Cisco IDS sensors with CiscoWorks VPN/Security Management Solution (VMS), and Cisco Threat Response technology.

“The Cisco IDS sensors report attempted intrusions to the CiscoWorks VMS management console, which immediately alerts my staff,” says Carsey. “Cisco Threat Response, which is integrated with the Cisco VMS console, collates the IDS alerts so that instead of receiving hundreds of messages about an IP address, we’ll receive one message saying, ‘This person is launching the attacks.’” The proactive angle? Carsey has configured the Cisco IDS to automatically stop any perceived attack with low, medium, or high severity. “We can blackhole the attacker using either the Cisco IDS sensor or the Cisco PIX Firewall,” he says. “Later we can check to see if we’ve been too aggressive, but in the meantime, we’ve protected our customers and our business.”

Mountain America has additional layers of security, as well, including Cisco Security Agent installed on critical servers. “On the outermost layer, we have IDS sensors on our external network that protect against broad types of attacks,” Carsey explains. “On the innermost layer, Cisco Security Agent guards against specific attacks targeted to Web or SQL servers.”

As another example of the value of integrated security compared to point solutions, suppose an employee takes home a laptop that becomes infected with viruses and spyware, and then brings it back to work. When the employee attempts to connect, Cisco Trust Agent on the laptop communicates with Cisco Network Access Control to direct the laptop to a remediation service and to third parties that will provide current patches and signatures. The result: The financial institution network proactively evades infection through integration of two Cisco solutions and third-party services.

Finding the Overlap in Regulations

Small and midsize financial institutions, in particular, need to handle network security cost effectively, and one way to do that is to invest in solutions that satisfy multiple regulatory requirements (see sidebar, “Regulatory Compliance”). “Most regulations were not written with the others in mind,” says Olslund. “And with so many mandates, smaller banks and credit unions are struggling to identify areas of overlap.”

Cisco helps banks and credit unions identify and address risk management elements common to multiple regulations. “Cisco helps smoothly navigate the journey to regulatory compliance,” says Jim Bright, Cisco US industry marketing manager for financial services. “For example, if a financial institution is subject to a regulation regarding safeguarding voice calls or e-mail transactions, we might suggest encrypting traffic over a VPN as one way to address the requirement.”

Making It Manageable

The Cisco Self-Defending Network helps financial institutions manage risks to information and voice traffic in an easier and more manageable environment, according to Bright. “It also helps them provide proof of compliance to regulatory bodies,” he says. “Proving that they’re taking the appropriate steps in proactive fashion puts financial institutions ahead of the game.”

The Self-Defending Network suits the cost-consciousness of small and medium banks and credit unions because it automates manual security processes, avoiding the need to add staff. “Without the Self-Defending Network, I’d need at least one full-time person just to manage our firewall logs and sensor alerts,” says Carsey. “We’ve eliminated that need by tying the threat response into the Cisco VMS console. If the Cisco IDS sensor sees an attack on a

particular host, it checks the host’s operating system version and security patches. If the host is not vulnerable, I’ve specified that the system should not alert me, so I have one less thing to look at.”

A Competitive Advantage

Not only does a Self-Defending Network safeguard customer trust and help meet regulatory requirements, it can help create a competitive advantage. One way is by ensuring service continuity. “If the ATMs for the bank down the street are down because of a security breach and yours are up and running, you stand to gain new customers,” says Olslund.

A secure network also enables financial institutions to offer new, competitive services. For example, with a secure wireless network, financial institutions can improve their customer service with “wireless concierge” service. Equipped with a tablet PC and wireless printer, a teller can provide noncash teller services to customers in line, making service faster.

A Self-Defending Network can cut operational costs, as well, by providing secure network access to teleworkers. Mountain America Credit Union reduced contact center costs by enabling agents to work from home, using the Cisco Secure Access Control Server for access control.

Aligning with a Partner

Simonsen emphasizes the role of technical support in the Self-Defending Network. “In a 24 x 7 business like banking, it’s critical that the business partner you align with offers support whenever you need it. Cisco SMARTnet and Cisco SASU [Software Application Support plus Upgrades] ensure that we have up-to-date files and responsive support around the clock. Cisco matches our level of urgency in getting security right before something goes wrong. That kind of commitment to the customer’s success is what differentiates a business partner from a vendor.” ■

FURTHER READING

- Self-Defending Networks for Financial Institutions
cisco.com/go/sdn_for_finance

Cluster and Grid Computing – Low Cost, High Power

By Joel Krauska and Drew Pletcher

With the recent appearance of new cluster and grid products and services, cluster and grid technologies are receiving considerable press attention. What's behind all the hype?

Think about the fastest computers in the world. These high-performance “supercomputers” model hurricanes, simulate nuclear explosions, and analyze the human genome. Companies such as Cray Research and Silicon Graphics dominated this market for years. Their custom “Big Iron” machines cost millions and were the only options for solving complex computing tasks.

In the past five years, significant changes have occurred. Clusters are replacing “Big Iron” machines. Clusters can be built from standard PC server and networking hardware with price tags a hundred times cheaper than custom supercomputers. Clusters make up half the systems on the “Top 500 list,” (top500.org), which ranks the 500 fastest computers in the world. Users of high-performance computing are changing as well. Corporations now harness the same computing power that only academic and research organizations had access to previously. More than half of the current “Top 500” machines are in commercial industry.

Consider how this applies to the smaller cousins of supercomputers, the proprietary SMP enterprise systems such as those sold by HP or Sun. These systems also can be replaced by clusters. Recently, Microsoft and Apple announced enhancements to their operating systems that allow them to form clusters. Large databases such as Oracle and DB2 now support clusters as well.

Plugging in the “Grid”

A less mature technology than clusters, “grid computing” usually implies a looser collection of computers than a traditional cluster. Grids typically span administrative domains to include different departments, buildings, or even different cities and regions of the world. The SETI@home project (setiathome.ssl.berkeley.edu) is an extreme example of this concept. Hundreds of computers owned by individuals worldwide are being used together to analyze radio telescope data.

The term grid also is often used to describe the concept of an on-demand utility computing model where computing resources can be purchased and allocated to an application as needed. The term draws parallels to normal utility grids like power or natural gas. Utility computing grids can be rented through “computing service providers.”

JOEL KRAUSKA AND DREW PLETCHER are technical marketing engineers in Cisco's Internet Systems Business Unit, and are responsible for high-performance computing cluster testing and analysis. Both are experts in data center switching and routing. They can be reached at jkrauska@cisco.com and drew@cisco.com.

Cluster Applications

In a cluster application, a single problem or query is broken down into multiple smaller pieces. Each piece is distributed to processing nodes in the cluster using a scheduling mechanism and job control. Many different cluster infrastructure tools exist to help accomplish this task.

Cluster applications vary widely, including weather simulation, fluid dynamics, biotech and genetic research, defense and energy research, aerospace and automotive design, graphics and video rendering, and financial analysis.

Three characteristics that can influence cluster application performance are message latency, throughput, and CPU utilization. A graphics rendering application may not care about message latencies. A database query may not be CPU bound. Choosing the appropriate cluster interconnect to support a given application is a common dilemma.

The hidden internal hardware connecting processors inside supercomputers and SMP systems is replaced with standard interconnect switch fabrics on clusters. The ubiquity of Ethernet, and the fact that most modern PC server systems come with a built-in Gigabit Ethernet network interface card (NIC), have made Ethernet the “interconnect of choice” for most cluster implementations. In the future, 10 Gigabit Ethernet will likely be used.

In addition to Ethernet, other cluster interconnect technologies exist for applications requiring higher throughput or lower end-to-end message latencies. InfiniBand, Myrinet, and Quadrics are common examples.

As a lower latency alternative to standard Ethernet, several vendors are developing RDMA-enabled NICs, also known as RNICs, which decrease Ethernet application latencies by bypassing operating system protocol overhead.

Cisco's Role

Cisco has many products to provide the interconnect infrastructure for both cluster and grid computing. Cisco has designed and built multi-thousand-node clusters for customers using Catalyst 6500 Series switches. Several Catalyst 6500 switches can be combined to create a nonblocking Layer 3 fabric, supporting up to 3600 nodes.

Cisco's storage switches can be used to build storage subsystems supporting clusters. Cisco's metro and WAN products facilitate grid implementations. Cisco security products can be used to protect grids and isolate users in the grid utility computing model. Also, Cisco partners with blade server manufacturers to build switches into the backplanes of these devices. ■

SPOTLIGHT ON:**Cisco Catalyst 4948-10GE Switch**

A challenge in any data center is how to serve more application traffic without



becoming overburdened by additional equipment. The new Cisco Catalyst 4948-10GE Switch overcomes this challenge for enterprises that need a device for single-rack-unit, multilayer aggregation of high-performance servers and workstations.

Based on the proven Cisco Catalyst 4500 Series hardware and software architecture, the Cisco Catalyst 4948-10GE offers exceptional Layer 2/3/4 switching performance, bandwidth, and reliability. This fixed-configuration switch delivers wire-speed throughput with low latency for data-intensive applications using a 126-Gbit/s switching fabric with a 102 million packets per second (pps) forwarding rate in hardware. It includes 48 ports of wire-speed 10/100/1000BASE-T Ethernet with two ports of wire-speed 10 Gigabit Ethernet (using X2 optics) for rack-optimized server switching applications.

Optional internal AC or DC 1 + 1 hot-swappable power supplies and a hot-swappable fan tray with redundant fans deliver the high reliability and serviceability required for server switching.

Among the new software features supported on the Cisco Catalyst 4948-10GE Switch are per port per virtual LAN (VLAN) quality of service (QoS) for differentiated QoS to individual VLANs on a trunk or access port; trunk port security; IEEE 802.1X private VLAN assignment and 802.1X private guest VLAN; 802.1X RADIUS-supplied session timeout, which enables the switch to determine the duration of a session and the action to take when the session's timer expires.

cisco.com/packet/172_npd1

Edge Routing, Access, and Aggregation

Cisco 1800 Series Integrated Services Routers: Fixed-Configuration Models

Designed for small offices, the new fixed-configuration Cisco 1800 Series Integrated Services Routers embed data, security, and wireless technology into a single system with wire-speed performance. Among the capabilities of the Cisco 1801, 1802, 1803, 1811, and 1812 fixed-configuration routers are secure broadband access; integrated ISDN basic rate interface (BRI), analog modem, or Ethernet backup ports for redundant WAN links and load balancing; secure wireless LAN (WLAN) for IEEE 802.11a and 802.11b/g operation with use of multiple antennas; advanced security including stateful inspection firewall, IP Security (IPSec) virtual private networks (VPNs), intrusion prevention, antivirus support; and 8-port 10/100 managed switch with virtual LAN (VLAN) and optional Power over Ethernet (PoE). For more on the new Cisco 1800 Series models, see page 73.

cisco.com/go/isr

Cisco 870 Series Integrated Services Routers

The new Cisco 870 Series Integrated Services Routers make it possible for small offices to run secure, concurrent services—including firewall, VPNs, and WLANs—at broadband speeds. The fixed-configuration Cisco 870 Series provides advanced security including stateful inspection firewall; IPSec VPNs; intrusion prevention and antivirus support; 4-port 10/100 managed switch with VLAN support; and secure WLAN 802.11b/g option with use of multiple antennas. Easy to deploy and manage centrally via Web-based configuration tools and Cisco IOS Software, these routers are ideal for deployment in small offices or teleworker sites as part of an enterprise network, and small to mid-sized businesses for secure WAN and WLAN connectivity. For more on the Cisco 870 Series, see page 73.

cisco.com/go/isr

Cisco 850 Series Integrated Services Routers

Designed for small businesses and small remote sites with up to 10 users, the new fixed-configuration Cisco 850 Series Integrated Services Routers provide secure connectivity with stateful inspection firewall and IPSec VPN; 4-port 10/100 switch; secure WLAN 802.11b/g option with a single fixed antenna; and easy setup, deployment, and remote management capabilities through Web-based tools and Cisco IOS Software. For more on the Cisco 850 Series, see page 73.

cisco.com/go/isr

Security and VPNs Cisco Security Auditor

Cisco Security Auditor software helps customers cost effectively audit their network infrastructure to assess compliance with corporate security policies and industry best practices. Deployed in a network operations center, the automated auditing and reporting capabilities of this software reduce audit time and eliminate costly manual auditing operations for large-scale networks. Cisco Security Auditor coverage includes Cisco PIX firewalls, virtual private network (VPN) devices, routers, switches, services modules, and the new Cisco ASA 5500 Series Adaptive Security Appliances (see below). For more on Cisco Security Auditor and other new security products from Cisco, see page 26. cisco.com/packet/172_npd2

Cisco Security Monitoring, Analysis, and Response System

The Cisco Security Monitoring, Analysis, and Response System (Cisco Security MARS) appliances provide a solution for security threat management, monitoring, and mitigation. Product features include network intelligence, context correlation, vector analysis, anomaly detection, hotspot identification, and automated attack mitigation. These features help administrators identify, manage, and eliminate network attacks as well as maintain compliance with security policies. With five Cisco Security MARS models to choose from, they can process up to 10,000 events or 300,000 flows per second and are designed for use in a network operations center. For more on Cisco Security MARS and other new security products from Cisco, see page 26.

cisco.com/packet/172_npd3

Cisco PIX Security Appliance Software Version 7.0

The Cisco PIX Security Appliance Version 7.0 software delivers an extensive list of new features. Key enhancements include capabilities for inspection and control of a broad range of HTTP, voice, and IP-based applications. In addition, a highly flexible security policy framework enables administrators to implement fine-grain control over individual user-to-application flows. For more on Cisco PIX Version 7.0 software and other new security products from Cisco, see page 26.

cisco.com/go/pix

Cisco Intrusion Prevention System Software Version 5.0

Cisco Intrusion Prevention System (IPS) Version 5.0 software enhances inline protection against threats such as spyware, adware, worms, viruses, and anomalous activity through detailed inspection of Layer 2 through 7 traffic. It also supports better targeting of attack-prevention actions, and offers more flexible IPS deployment and collaboration options for improved security control. IPS Version 5.0 is supported by Cisco IPS 4200 Series appliances, and the Cisco Catalyst 6500 Series Switch and 7600 Series Router Intrusion Detection System Module (DSM-2) through a software upgrade. For more on new IPS enhancements and security products from Cisco, see page 26.

cisco.com/go/ips

Cisco VPN 3000 Series Concentrator Version 4.7

The Cisco VPN 3000 Series Concentrator Version 4.7 software provides a new Secure Sockets Layer (SSL) VPN client, support for Citrix-based applications, and integrated support for Network Admission Control (NAC) to enforce security on devices using the Cisco IPSec client. The Cisco Secure Desktop features improve endpoint security by creating a virtual desktop that protects and eliminates sensitive session information. For more on the Cisco VPN 3000 Series Concentrator and other new security products from Cisco, see page 26.

cisco.com/go/vpn3000

Cisco Security Agent Version 4.5

Key enhancements in Cisco Security Agent Version 4.5 software include protection for server and desktop computing systems, or “endpoints.” Rather than relying on signature matching, Cisco Security Agent identifies and prevents malicious behavior before it can occur, including preventive protection against entire classes of attacks, such as port scans, buffer overflows, Trojan Horses, malformed packets, and e-mail worms. It includes support for 100,000 agents per management server, NAC, application inventory and use tracking, and automated checking for current virus definition files. User-based and location-based security policies can be applied based on the physical or logical location of a user’s computer. Cisco Security Agent Version 4.5 is covered in greater detail on page 47.

cisco.com/go/csa

Cisco ASA 5500 Series Adaptive Security Appliances

The new Cisco ASA 5500 Series Adaptive Security Appliances combine best-of-breed security and VPN services with an innovative Adaptive Identification and Mitigation (AIM) services architecture. Designed for small and mid-sized businesses and enterprise networks, the multi-function Cisco ASA 5500 Series provides proactive threat defense that stops attacks before they spread through the network, controls network activity and application traffic, and delivers flexible VPN connectivity. Security capabilities include firewall, intrusion prevention, network antivirus support, and IPSec/SSL VPN. The AIM services architecture allows businesses to adapt and extend the security services profile of the Cisco ASA 5500 Series through highly customizable, flow-specific security policies that tailor security needs to application requirements while providing performance and security service extensibility via user-installable security services modules. For more on the Cisco ASA 5500 Series and other new security products from Cisco, see page 26.

cisco.com/go/asa 



Cisco IOS Firewall: Application Inspection and Control

The Cisco IOS Firewall has been enhanced with advanced application inspection and control capabilities that are delivered through HTTP and e-mail inspection engines. The HTTP Inspection Engine enforces protocol conformance and prevents network access by malicious or unauthorized behavior such as port 80 tunneling, malformed packets, Trojan Horses, and instant messaging traffic that is probing for vulnerabilities. The E-mail Inspection Engine detects misuse of e-mail connectivity and prevents protocol masquerading activity. For more on this and other new security products from Cisco, see page 26.

cisco.com/packet/172_npd5

Cisco IOS Software: IPSec Virtual Tunnel Interfaces

Cisco IPSec virtual tunnel interfaces (VTIs) configure IPSec VPNs between site-to-site devices. These tunnels provide a designated pathway across the shared WAN and encapsulate traffic with new packet headers, which helps to ensure traffic privacy and delivery to specific destinations. The IPSec VTIs were introduced in Cisco IOS Software Release 12.3(14)T. New Cisco security products and enhancements are covered in greater detail on page 26.

cisco.com/packet/172_npd6

Cisco IOS Software: Enhanced Inline IPS Functionality

Enhanced inline IPS capabilities introduced in Cisco IOS Software Release 12.3(14)T increases protection against new classes of threats such as spyware, network antivirus, and malware associated with Instant Messaging (IM) applications that significantly improves the ability to prevent and mitigate damage from worm and virus attacks. The new IPS functionality also allows users to create custom signatures to address newly discovered threats for broader protection. New Cisco security products and enhancements are covered in greater detail on page 26.

cisco.com/packet/172_npd7

Cisco VPN Acceleration Module 2+

The Cisco VPN Acceleration Module 2+ (VAM2+) for Cisco 7200 Series and Cisco 7301 routers provides high-performance encryption, compression, and key-generation services for IP Security (IPSec) VPN applications. Designed for enterprise and service provider network environments, the VAM2+ supports all Cisco VAM2 features, but adds hardware acceleration for 192-bit and 256-bit Advanced Encryption Standard (AES) keys.

cisco.com/packet/172_npd8

Wireless

Cisco 1000 Series Lightweight Access Point

The Cisco 1000 Series Lightweight Access Point is designed for enterprise deployments that require coverage flexibility. These devices handle important IEEE 802.11 a/b/g radio functions within a Cisco wireless LAN, including radio transmit and receive, client probe requests, and air monitoring. The Cisco 1000 Series Lightweight Access Point also handle time-sensitive functions, such as Layer 2 encryption, that enable Cisco wireless LANs to securely support voice, video, and data applications. A unique model, the Cisco 1030 Remote-Edge Access Point communicates with Cisco Wireless LAN Controllers (see below) via most standard WAN technologies. The Cisco 1000 Series Lightweight Access Point is a result of the recently acquired Airespace product portfolio.

cisco.com/go/securewireless 



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

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

Cisco Wireless LAN Controllers

Cisco Wireless LAN Controllers perform system-wide wireless LAN functions such as security policies, intrusion prevention, RF management, quality of service (QoS), and mobility. The controllers work with Cisco Wireless Control System software and Cisco 1000 Series lightweight access points. For small and midsize enterprise facilities, such as branch offices, the Cisco 2000 Series Wireless LAN Controller supports up to six lightweight access points and coverage up to 60,000 square feet. In larger enterprise environments, three models of the Cisco 4100 Series Wireless LAN Controller (the 4112, 4124, and 4136) support 12, 24, or 36 lightweight access points and dual Gigabit Ethernet uplinks for connection to the wired LAN. The Cisco Wireless LAN Controllers are a result of the recently acquired Airespace product portfolio.

cisco.com/go/securewireless

Cisco Wireless Control System

Cisco Wireless Control System (WCS) software enables managers to centrally design, control, and monitor Cisco wireless networks that encompass hundreds of Cisco Wireless LAN Controllers and thousands of Cisco 1000 Series lightweight access points. The graphical software includes functions for RF prediction, policy provisioning, network optimization, troubleshooting, user tracking, security monitoring, trending and analysis reports, and wireless LAN systems management. The Cisco WCS is a result of the recently acquired Airespace product portfolio.

cisco.com/go/securewireless

Cisco 3800, 2800, and 1800 Series Integrated Services Routers: New Wireless LAN Interface Cards

New wireless LAN interface cards in the high-speed WAN interface card (HWIC) form factor provide integrated, secure IEEE 802.11 access point functionality for the Cisco 3800, 2800, and 1800 (modular models) Series integrated services routers. Designed for enterprise branch offices and small and midsized businesses, the Cisco HWIC-AP 802.11b/g and HWIC-AP 802.11a/b/g interface cards deliver single-band 802.11b/g or dual-band 802.11a/b/g radios; support for various external dipole or dual-mode antennas; and extensive wireless LAN security. These new capabilities for Cisco integrated services routers are covered in greater detail on page 73.

cisco.com/go/isr

Storage Networking Cisco MDS 9000 Series Storage Services Module

The Cisco MDS 9000 Series Storage Services Module (SSM) is an open, standards-based module designed specifically to support intelligent fabric applications from multiple Cisco partners including EMC, IBM, and Veritas. The SSM can be used with any Cisco MDS 9500 Series Director or Cisco MDS 9200 Series Fabric Switch and provides 32 Fibre Channel ports. Embedded ASICs in the module deliver performance improvements for third-party intelligent fabric applications that are based on open standards such as Fabric Application Interface. For an article

on securing storage area networks (SANs) using the Cisco MDS 9000 Series switches, see page 42.

cisco.com/packet/172_npd9

Networked Home Linksys Wireless-G Media Link

The Linksys Wireless-G Media Link supports Digital Transmission Content Protection over IP (DTCP-IP), an industry standard that will help consumers enjoy premium and high-definition video and music services on televisions and stereos around the home. The Wireless-G Media Link connects to TVs and stereos, then to a home network by Wireless-G (802.11g) networking or standard 10/100 Ethernet cabling. The media link can also connect directly to a PC for transferring content to the entertainment center.

cisco.com/packet/172_npd10

Linksys Wireless A/G Game Adapter

The Linksys Wireless A/G Game Adapter gives wireless connection capabilities to any wired, Ethernet-equipped game console. This connection supports lag-free, head-to-head, or Internet gaming at up to 54 Mbit/s over a Wireless-A, -B, or -G home network.

cisco.com/packet/172_npd11

Linksys Wireless-G Router with SRX and Wireless-G PC Card with SRX

The Linksys Wireless-G Router with Speed and Range eXpansion (SRX) and Wireless-G PC Card with SRX deliver faster throughput, fewer dead spots, and increased range over standard Wireless-G networks. SRX is based on Multiple In, Multiple Out (MIMO) technology, a key component in the upcoming Wireless-N (802.11n) standard, which uses smart radio and technology antenna on the wireless router or client adapter.

cisco.com/packet/172_npd12

Linksys Compact Wireless-G Router

The Linksys Compact Wireless-G Router provides an integrated device for home networking. The product includes a wireless access point using Wireless-G (802.11g) or Wireless-B (802.11b); a 4-port 10/100 switch to connect wired Ethernet devices, and a router for Internet access over a high-speed cable or DSL Internet connection.

The router's compact size and built-in antenna make it suitable for placement almost anywhere in the home.

cisco.com/packet/172_npd13

Linksys Power over Ethernet Adapter Kit

The Linksys Power over Ethernet (PoE) Adapter Kit supplies 12-volt AC power directly to wall-mounted or ceiling-mounted devices, including wireless access points, routers, and bridges. The kit includes an injector and splitter, which enable the electrical power and data to travel on one Category 5 cable.

cisco.com/packet/172_npd14 



Cisco IOS Software Cisco IP/MPLS Interprovider Solution

The Cisco IP/MPLS Interprovider solution enables the implementation of Multiprotocol Label Switching (MPLS) services across service provider boundaries—allowing service providers to partner with other providers to extend the coverage of their existing services and introduce new offerings that go beyond their network. The Cisco IP/MPLS Interprovider solution includes five technologies: Inter-Autonomous System (AS) multicast VPN, Inter-AS Traffic Engineering, Interprovider MPLS VPN over IP, MPLS VPN Inter-AS and Carrier Supporting Carrier (CSC) load balancing, and interprovider network management. To manage these new interprovider technologies, Cisco has made enhancements to Cisco IP Solution Center, Cisco Info Center, CiscoWorks LAN Management Solution, and Cisco IOS NetFlow with Reporting and IP service-level agreement (SLA) capability. The Cisco IP/MPLS Interprovider solution is covered in greater detail on page 69.

cisco.com/packet/172_npd15

Configuring and Troubleshooting Dial-Related Issues

The Cisco Networking Professionals Connection is an online gathering place to share questions, suggestions, and information about networking solutions, products, and technologies with Cisco experts and networking colleagues. Following are excerpts from a recent “Ask the Expert” forum, “Configuring and Troubleshooting All Dial-Related Issues,” moderated by Cisco’s Tejal Patel. To view the full discussion, visit cisco.com/packet/172_10a1. To participate in other live online discussions, visit cisco.com/discuss/networking.

Q: I have a Cisco AS5800 and AS5400 and am getting low call success rate. What is the recommended modemcap that can be used to improve the CSR [connection success rate] on the MICA modems and SPEs [service processing elements]?

A: Here is the link to a document called “Recommended Modemcaps for Internal Digital and Analog Modems on Cisco Access Servers,” which talks about the recommended modemcap for MICA and NextPort-based platforms to improve the CSR: cisco.com/packet/172_10a2.

Q: Can we mix the T1 and E1 cards on AS5800 and AS5850?

A: No, mixing T1 and E1 cards are not supported or recommended on the AS5800 and AS5850.

Q: I have a Cisco 2621 Router with VWIC-1MFT-T1 with PRI line for voice calls. Can I support data calls on the same PRI?

A: The VWIC-1MFT-T1 card does not support ISDN data connections. Here is the link to a document called “Cisco Digital 1-Port and 2-Port T1 Multi-Flex Voice WICs,” which talks about the features of that card: cisco.com/packet/172_10a3. When the card is set up to use ISDN PRI signaling, ISDN data connection is not supported. The card is unable to terminate the ISDN 64K or 56K data connection. It only supports voice call termination when using ISDN PRI signaling. Also, the Multi-Flex Trunk, with or without the accompanying voice-enabling hardware, is unable to terminate a modem connection on the router in a traditional NAS dial scenario.

Q: I want to connect to the Internet using two GPRS [General Packet Radio Service] connections via two serial interfaces. The service provider will provide me with two different IP addresses on two serial interfaces (no PPP multilink is available). Is there any method to do the load sharing?

A: Multilink is the best option, which also buys you a fragmentation. But because no multilink is available, you need to rely on routing protocol to queue the packets on two equal cost links for load balancing. You can use Enhanced Interior Gateway Routing Protocol (EIGRP) or Open Shortest Path First (OSPF) on the

router. Here are links to the documents called “How Does Load Balancing Work?” (cisco.com/packet/172_10a4) and “Troubleshooting Load Balancing over Parallel Links Using Cisco Express Forwarding” (cisco.com/packet/172_10a5), which discuss the process in detail.

Q: Using a VWIC-2MFT-T1 in a Cisco 2801 Router, can one port be set to “User” and the other port set to “Network”? If all I do is pass calls from port 0 to port 1 on the VWIC, will the calls require any DSP [digital signal processing] resources? Will there be an effect on modem and fax calls going through the Cisco 2801 to the Norstar?

A: If you use T1 crossover cable between those two ports, you can make calls across with one port as “Network” side and the other as “User” side. DSPs are required to interface the telephony side with the IP side. So, if you want to pass those calls to the IP side, you need to have DSP. Also, it should not affect the operation of fax modem calls.

Q: I have a configuration on a Cisco 1760 Router that runs an IPSec [IP Security] VPN with another site across the Internet. When the VPN goes down, I route my packets out via a dialer interface that uses DDR [dial-on-demand routing] and an access control list to keep uninteresting traffic from initiating a call keeping it up. I would like to run IPSec across the dial link. At what point is a packet being sent across a dialer checked for being interesting?

A: You can configure a router to initiate a DDR session triggered by IPSec. Incoming traffic to a router that needs to go over a DDR link and that matches the crypto map definition will be considered “interesting.” In that case, the IPSec tunnel is created, where the destination IP address is the remote IPSec peer. Check out the document called “Setting Up IPSec on a DDR Link” (cisco.com/packet/172_10a6), which discusses interesting and uninteresting traffic over the IPSec tunnel for DDR in detail.

Do you have a question about configuring or troubleshooting dial-related issues? Ask the NetPro Expert. Send your question to packet-netpro@cisco.com, with the subject line “Dial-Related Issues.” ■



TEJAL PATEL, CCIE No. 6619, is an internetworking engineer in the Cisco Technical Assistance Center, San Jose, California. He specializes in cable and voice solutions. He can be reached at tepatel@cisco.com.

Wireless Security, Continued from page 41

- **Broad and detailed network visibility.** Infrastructure access points alone might not be able to “see” into far corners or stairwells of a building. All Cisco infrastructure access points running Cisco IOS Software act as both access points and intrusion detection sensors. In addition, via the Cisco Compatible Extensions program, laptops and other client devices serve as sensors that report the existence of rogues, allowing intrusion detection to stretch farther for the distributed solution.
- **Multifrequency scanning.** In monitoring the air space, companies are basically taking continual inventory of the airwaves. Sensors should scan 802.11a, b, and g networks, regardless of whether a given organization is actually running traffic on each wireless band. Otherwise, the enterprise risks security threats from ad-hoc and rogue access points operating in the other bands.
- **Location tracking.** Cisco’s ability to apply location tracking to rogue device detection assists in eradicating the source of a number of types of attacks.

In addition to taking appropriate counter measures to spurn the effects of an attack, advanced tracking locates the actual source of undesirable activity to expedite problem resolution.

24/7 Requirements

The use of wireless LANs has ramped up to the point where a mere occasional check on the state of the airwaves for unusual behavior or performance problems is no longer sufficient.

Rather, all network layers must work in real-time to continually protect data both in corporate servers and on end-user devices from the unique susceptibilities of the wireless medium.

Wireless is especially prone to attack because radio is an open medium, shared among anyone transmitting and receiving in the same unlicensed RF spectrum, and because of the self-associating nature of wireless clients and access points. As volumes of 802.11-based traffic grow, ad-hoc associations will only increase, with more potential for mischief. Fortunately, the

protection safeguards are available for diffusing wireless service disruptions and hijacks, whether intentional or not. ■

FURTHER READING

- Expansion of Cisco Wireless Networking cisco.com/packet/172_7c1
- White Paper “Wireless Intrusion Detection & Prevention with Lightweight Access Points” cisco.com/packet/172_7c2
- Cisco Structured Wireless-Aware Network Solution cisco.com/go/swan

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	4
BellSouth Business	www.bellsouth.com/business/netvpn	OBC
Boson Software	www.boson.com	A
Cisco Press	www.ciscopress.com	B
Cisco Systems	www.cisco.com/poweredby	38/68
Cisco Systems-Networkers	www.cisco.com/go/nwpacket	F
Cisco Systems	www.cisco.com/go/isr	56
Corvil	www.corvil.com	6
eiQ Networks	www.eiqnetworks.com	13
extraxi	www.extraxi.com/packet	14
Fluke Networks	www.flukenetworks.com/packet	72
GL Communications	www.gl.com	76
Liebert Corporation	IP.Liebert.com	25
NetScout	www.netscout.com/ad/cii	52
Network General	https://networkgeneral.mnl.com/c1	82
New Edge Networks	www.newedgenetworks.com/products	22
OPNET Technologies	www.opnet.com	60
Panduit	www.panduit.com/dp08	IBC
Solsoft	www.solsoft.com/packet	8
SurfControl	www.surfcontrol.com/go/cisco	46
Trend Micro	www.trendmicro.com/cisco	30/31
Websense	www.websense.com/patch1	50

CACHE FILE

Snippets of Wisdom from Out on the Net

CYBER QUOTE

“What goes up must come down. Ask any system administrator.”

—Anonymous

RFID Market Growth

The worldwide Radio Frequency Identification (RFID) tag market will grow nearly tenfold from US\$300 million in 2004 to US\$2.8 billion in 2009, according to an In-Stat report. That would make RFID tags the most prevalent wireless technology since the cell phone. The key drivers of the projected growth will be shipping cartons and other supply chain elements, expected to account for 35.1 percent of the RFID market by 2009, up from 4.9 percent in 2004. The second largest market for RFID in the forecast's later years will be consumer products, currently one of the most privacy-sensitive verticals.

Wireless Gaining Subscribers Worldwide

The untethered life holds great appeal for Internet users worldwide as the number of subscribers to wireless applications continues to grow. The Yankee Group predicts that wireless users will grow nearly 9 percent from 2002 to exceed 1.75 billion in 2007, and In-Stat/MDR expects the number of worldwide wireless Internet subscribers will have risen from 74 million at the end of 2001 to more than 320 million by the end of 2006. The Asia-Pacific region alone will account for a significant portion of the global usage, reports the Yankee Group.

Net Lingo

Pixel dust—Slang for the thin coat of dirt on your computer screen. (netlingo.com)

Traveler's First Trip Is Often the Internet

Whether it's used for research or for purchase, the Internet is an increasingly valuable tool for travelers. According to an April 2003 My AvantGo survey of more than 1,000 individuals, 52 percent purchased more than half of their travel needs online, with 29 percent indicating that they made all their travel arrangements on the Web. As an added bonus, 30 percent plan to increase their online travel purchases over the coming year.

Online Safety: Users Talk the Talk, Don't Walk the Walk

When it comes to online security, perception is definitely not reality. According to a joint AOL/National Cyber Security Alliance (NCSA) Online Safety Study, users believe themselves to be generally “safe,” but their computers indicate otherwise. The survey found 77 percent of users believe their home computers are either very safe (28 percent) or somewhat safe (49 percent) from online threats. When pollsters asked how safe users feel they are against viruses, the combined percentage dropped to 73 percent. When asked about safety from hackers, the numbers drop further still, to only 60 percent. The majority of participants (67 percent) either hadn't updated their virus protection in the past week, or had no antivirus protection at all.

THE 5TH WAVE



“Our security program responds to three things. A false access code, an inappropriate file request, or sometimes a crazy hunch that maybe you're just another slime-ball with misappropriation of secured data on his mind.”

©The 5th Wave, www.the5thwave.com