

# How Cisco Multilayer Director Switch Operates with Other SAN Switches During SAN Migration

Cisco Multilayer Director Switch offers safe migration path from other vendor SAN switches.

**Cisco IT Case Study / Storage Networking / SAN Switch Interoperability in Migration:** This case study describes how Cisco Systems® upgraded its engineering storage area network (SAN) from McDATA switches to a more scalable environment with Cisco MDS 9509 multilayer director switches at the core and McDATA 3032 switches at the edge, running in interoperability mode. After demonstrating flawless performance in interoperability mode for several months, Cisco replaced the McDATA edge switches with Cisco MDS 9120 multilayer fabric edge switches to gain additional benefits, including virtual SAN (VSAN) capabilities. The Cisco global network is a leading-edge enterprise environment that is one of the largest and most complex in the world. Cisco customers can draw on Cisco IT's real-world experience in this area to help support similar enterprise needs. A tribute to thorough planning, the transition to a mixed Cisco MDS and McDATA SAN was achieved without service interruption.

“When we took out one core switch, the host didn't blink and the clients didn't even notice.”

– Dave Angulo, Program Manager, Cisco Systems

## BACKGROUND

Through the 1990s, many applications at Cisco relied on direct-attached storage (DAS). “DAS is fine when you're small,” says Jesse Adam, system administrator for Cisco engineering. “But as Cisco grew and acquired other companies, DAS became difficult to manage because of its lack of scalability.”

Specifically, moves or upgrades required that IT staff physically visit each device. Adding additional storage required shutting down the server. The proliferation of different types of storage by Cisco organizations worldwide made it difficult to identify the source of performance problems; for example, slow disks, different firmware versions, or different storage types.

DAS also prevented Cisco from optimizing utilization. If an application needs additional storage capacity with DAS, the only option is to add an entire new disk, even if only a fraction of its capacity is needed. “We wanted a new approach to storage that would let us grow in flexible increments,” says Adam. “We also wanted to be able to consolidate storage so that we could allocate it without being constrained by the physical location of the application or storage device.” Utilization, in fact, was a chief incentive for shifting away from DAS. In 2001, just 1.35 percent of Cisco servers operated at or above 85 percent utilization, while 42.4 percent operated at or under 20 percent utilization.<sup>1</sup> Storage frames, too, were often underutilized, usually because of port density issues at the disk adapter, host bus adapter (HBA), or switch port level.<sup>2</sup>

<sup>1</sup> Chiefly, Cisco uses Sun Microsystems servers running the Solaris operating system (OS) and HP servers running HP-UX, in addition to a few servers running Linux or Microsoft Windows 2000.

<sup>2</sup> Most storage frames at Cisco are re EMC Symmetrix 8530 and 8830 or HP XP512.

At the end of 2001, Cisco began transitioning from its DAS model, in which individual business units owned storage devices that were directly connected to the host machine, to a shared SAN environment. By September 2002, Cisco IT had deployed more than 55 SAN switches from McDATA and Brocade in three large business data centers, for a total of 1400 ports. Buoyed by the success of the DAS-to-SAN migration in business data centers, engineering data center administrators within Cisco began migrating their storage to a SAN environment, using the McDATA Enterprise Director 6064 switch at the core and McDATA Sphereon 3032 switches at the edge. “The total cost of ownership for the SAN was at least 12 percent lower than for DAS,” says Lance Perry, vice president of IT infrastructure. Ultimately, Cisco deployed approximately 100 third-party SAN switches in more than 10 business and engineering data centers.

## CHALLENGE

When Cisco introduced the Cisco MDS 9000 Series Multilayer Switch, the company wanted to take advantage of its better scalability, higher performance, and virtual SAN (VSAN) capabilities. Cisco MDS 9509 multilayer director switches offer up to 224 Fibre Channel ports on a single chassis, compared to 64 ports on the largest McDATA switch, the 6064. With VSANs, Cisco could support multiple SANs from a single switch, eventually making the physical location of hosts and storage irrelevant.

Instead of replacing all its McDATA switches with Cisco MDS 9000 Series multilayer switches at once, Cisco engineering first replaced only McDATA 6064 core switches with the more powerful Cisco MDS 9509 multilayer director switches. The engineering SAN would continue to use McDATA 3032 switches at the edge, operating with Cisco MDS 9509 multilayer director switches in interoperability mode. “Many companies started out with lower-end McDATA switches and now need a higher-end switch like the Cisco MDS 9509,” says Adam. “We realize that some of them might not want to make the capital outlay to replace all their McDATA switches at once. We decided to run our engineering SAN in interoperability mode for several months so that we could learn the ins and outs for our customers who would be making the same migration.”

The plan: In the summer of 2003, in a Cisco engineering building in San Jose, California, IT would install two Cisco MDS 9509 multilayer director switches at the core and connect them to the existing McDATA 3032 edge switches and the storage arrays. After six months of running the SAN in operability mode, Cisco would replace the McDATA 3032 edge switches with Cisco MDS 9120 multilayer fabric switches to gain additional benefits. The technical challenge would be to manage the transition without interruption to Cisco business-critical applications that rely on the SAN.

## SOLUTION

Cisco wanted to test operability on a SAN that supports a mission-critical application with high performance demands. Its choice: IBM Rational ClearCase, which thousands of Cisco developers rely on daily as their source-code library to develop Cisco IOS® Software. “ClearCase is a demanding application with thousands of transactions per second, low latency requirements, and critical availability requirements,” says Dave Angulo, Cisco program manager for engineering storage.

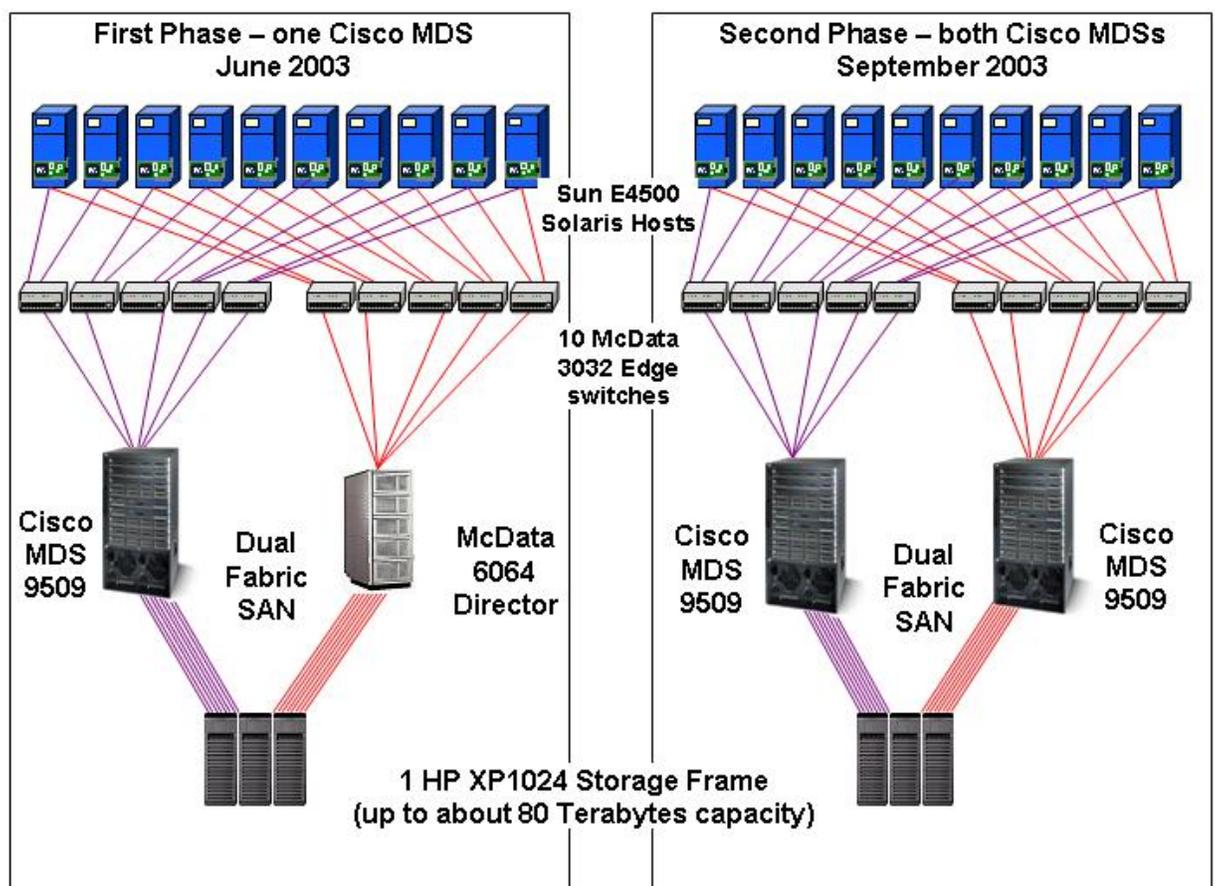
To ensure business continuity in the event of hardware, software, or network failure, the Cisco IT group designed the SAN with complete redundancy (see Figure 1). Each of the two Cisco MDS 9509 multilayer director switches would connect to two different HBA cards on the ClearCase host using two existing McDATA 3032 edge switches. Each Cisco MDS 9509 Multilayer Director Switch would have multiple paths to the Hewlett-Packard Development Company, L.P. (HP) XP 1024 disk array. The Cisco MDS 9509 multilayer director switches have dual backplanes, dual power supplies, and dual supervisor engine modules, with stateful process failover and restart. “We designed the SAN with complete redundancy so we had no single point of failure—from the host to the edge and from the core to the disk array,” says Adam.

Cisco migrated the first core switch in June 2003 and the second in September 2003.

### Cutover Process

The cutover required configuring the Cisco MDS 9509 multilayer director switches, disconnecting the McDATA core switches from the storage array and edge switches, and finally reconnecting the storage array and edge switches to the new core Cisco MDS switches. Cisco accomplished the transition in two phases—one Cisco MDS 9509 Multilayer Director Switch at a time. The phased migration ensured that the host continued operating without interruption during the transition (see Figure 1).

**Figure 1.** Migrating from McDATA 6064 Core Switches to Cisco MDS 9509 Multilayer Director Switches, in Two Phases. Configuring the Cisco MDS Switch



To configure the Cisco MDS 9509 multilayer director switches for the production environment required establishing logins and security settings and configuring the firmware. “Cisco IT storage administrators built an operating template that saves time during configuration,” says Hagen Finley, an HP consultant who worked on the project. “We started with that and then cut and pasted a lot of the functionality in terms of logins, ‘phone home,’ and other parameters.”

When it was time to deploy the second Cisco MDS 9509 Multilayer Director Switch in September 2003, three months after the first core switch, Finley took advantage of a unique feature of Cisco MDS 9000 Series multilayer switches—the ability to export configuration information to a text file, edit it for the second switch on a personal computer (PC), and then import the edited configuration file back to the switch. “We saved a lot of time,” Finley says. In fact, configuring both switches took one person approximately eight hours, including planning.

## Disconnecting the McDATA Switch

The next step was to bring down the first network fabric. Like most SAN storage environments, the Cisco environment consists of a dual fabric. Each host on the SAN has two HBAs, connected to two separate McDATA edge switches in different fabrics. Each edge switch, in turn, is connected to one of the core switches, which connects to a storage array. Therefore, shutting down one fabric does not affect host access to storage. “We sent out notification that we would be shutting down one of the fabrics, and avoided peak times,” says Finley. The fabric was down for less than six hours while Finley and his team accomplished the following procedures:

- HP disconnected the McDATA Intrepid 6064 Director switch.
- HP connected the Cisco MDS 9509 Multilayer Director Switch to the HP XP 1024 storage array.
- HP changed the modality of the McDATA 3032 edge switches from “McDATA fabric” to “Open Fabric.” This configures the McDATA switches for interoperability.
- HP validated the domain IDs of the McDATA edge switches. Only a subset of domain IDs work in interoperability mode because the McDATA switch supports fewer domain IDs than the Cisco MDS switch.
- HP checked for zoning conflicts. Edge switches have zoning tables that cannot easily be deleted. When those edge switches are connected to the Cisco MDS 9509 multilayer director switch, it propagates its own zones. “We had to validate that the proper MDS zone had been propagated and that the old zone definition wasn’t interfering,” says Finley.

With these steps complete, the McDATA 3032 edge switches were prepared to interoperate with the Cisco MDS 9509 Multilayer Director Switch.

## Connecting McDATA Edge Switches to the Cisco MDS Switch

Finley connected the McDATA edge switches to the Cisco MDS switch one at a time, using Inter-Switch Links (ISLs), which are physical fiber connections. The team confirmed the connection in two ways. One was to check that the Cisco MDS 9509 Multilayer Director Switch acknowledged the connection. For extra assurance, a team member logged onto a UNIX server on the storage path to confirm that the host could see its storage. “The first switch worked perfectly the first time,” says Finley. Initially, the second switch did not work, and the team quickly identified the problem—a higher domain ID than McDATA switches permit<sup>3</sup>. A change in domain IDs immediately solved the problem.

A tribute to thorough planning, the transition to a mixed Cisco MDS and McDATA SAN was achieved without service interruption. “When we took out one core switch, the host didn’t blink and the clients didn’t even notice,” says Angulo.

## RESULTS

The Cisco engineering SAN ran flawlessly in interoperability mode from June 2003 to January 2004. At that point, Cisco declared the interoperability demonstration a success and replaced its McDATA 3032 edge switches with Cisco MDS 9120 multilayer fabric edge switches to gain the benefits of a SAN based exclusively on Cisco MDS switches. During the interoperability phase, Cisco improved scalability and utilization and achieved its availability goals. Soon it will take advantage of the powerful VSAN capability of MDS 9000 Series multilayer switches, as described in “Next Steps.”

## Scalability and Utilization

The higher port density of the Cisco MDS 9509 Multilayer Director Switch—32 ports per blade and up to 224 Fibre Channel—is reducing costs, improving utilization, and freeing data center floor space. In fact, Cisco engineering

---

<sup>3</sup> Companies that use dynamic domain IDs instead of static domain IDs avoid this problem. Cisco IT chose to use static domain IDs to maintain unique domain IDs throughout the company. Presently, it is necessary only to confirm that all domain IDs within a single data center SAN are unique. When Cisco eventually interconnects all its data centers and SANs, today’s decision to maintain unique domain IDs will prevent potential domain ID conflicts.

plans to add 50 additional Cisco MDS 9216 multilayer fabric edge switches and will be able to support all 60 edge switches with only two Cisco MDS 9509 multilayer director switches, rather than the four to six McDATA 6064 switches that would be needed.

The scalability of the MDS 9509 Multilayer Director Switch ensures optimal resource utilization, as well, according to Angulo. “Suppose the ERP environment runs out of ports but the engineering environment has free ports,” he says. “In the past we couldn’t utilize the available resource. But with the MDS 9509 Multilayer Director Switch, we can simply plug the ERP hosts into the free ports and configure them as part of the ERP VSAN.” To participate in a VSAN, the Cisco MDS 9509 Multilayer Director Switch can reside anywhere in the Cisco campus metropolitan-area network (MAN), including a different data center.

### **High Availability**

The two Cisco MDS 9509 multilayer director switches have remained available without interruption since they were deployed. Their high availability is a result of redundant supervisor engines, fully stateful supervisor engine failover, redundant crossbars, hitless software upgrades, individual process restart ability, and process isolation within VSANs. The ClearCase host experienced no service interruption during migration, and no application outages since then have been attributable to the Cisco MDS 9000 Series multilayer switches. “Everything stabilized almost immediately after we brought up the second fabric,” says Finley.

Finley notes that the high availability of Cisco MDS 9509 multilayer director switches is especially critical when companies transition from a DAS to a SAN environment. In a DAS environment, losing a Fibre Channel connection between a host and storage typically affects only one set of applications, for one business group. In contrast, in a SAN environment, losing connectivity between hosts and storage can take down applications and business groups throughout the data center.

### **VSANs**

Presently, all hosts and storage in the Cisco engineering SAN reside on the same VSAN, and Cisco is not yet taking full advantage of VSAN capabilities. If administrative traffic volume begins to affect performance, Cisco can solve the problem by creating multiple VSANs on the same switch, each with its own broadcast domain. For Cisco customers, an advantage of VSANs in a mixed MDS-McDATA environment is that a single Cisco MDS 9509 Multilayer Director Switch chassis can support some VSANs in interoperability mode while other VSANs retain normal functionality. This is important because a few features of the Cisco MDS 9509 Multilayer Director Switch do not function in interoperability mode; for example trunking, which involves aggregating I/O so that multiple cables appear as one. With the McDATA switch, in contrast, interoperability mode applies to the entire switch or not at all.

## **LESSONS LEARNED**

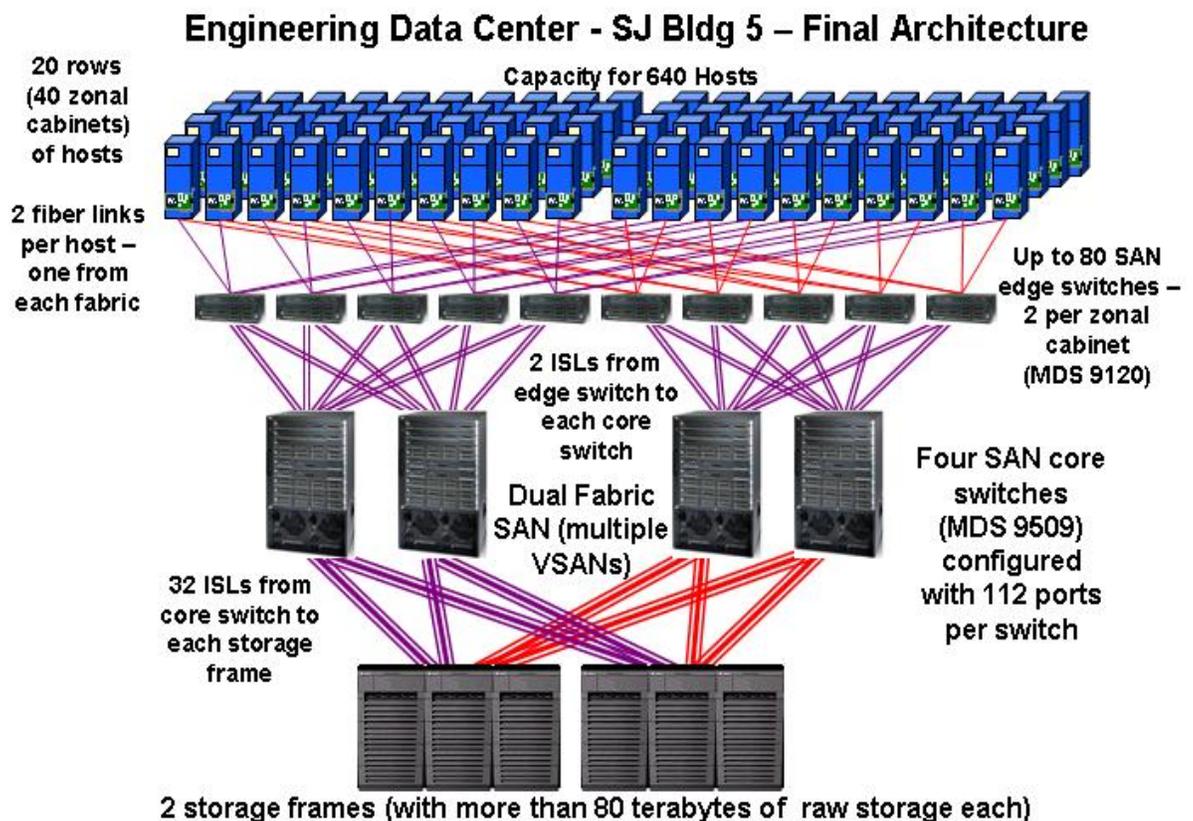
“It’s extremely important to plan the cutover, because if something breaks, you affect many hosts,” says Adam. For example, Cisco planned the fiber network topology so that the team could replace the two core switches one after the other instead of both at once—good insurance for a smooth transition, even if one of the fabrics should experience a problem.

Finley emphasizes that it is important to check that the McDATA domain IDs are in the allowed range for interoperability mode and to check for zoning conflicts between the Cisco MDS switches and McDATA switches. It is also important to realize that when Cisco MDS 9000 Series multilayer switches run in interoperability mode with switches from other vendors, features are limited to the intersection of the feature sets of the two switches. Notably, the sophisticated VSAN and network management features available on Cisco MDS switches do not work in an interoperability environment. After swapping out its remaining McDATA edge switches for Cisco MDS 9120 multilayer fabric edge switches in January 2004, Cisco gained unrestricted VSAN capabilities, greater scalability, and superior network management.

## NEXT STEPS

Cisco IT plans to continue expanding the SAN in the engineering data center in San Jose building 5, adding additional hosts and storage frames as needed. The capacity is unprecedented. Each pair of Cisco MDS 9120 multilayer fabric edge switches can support 16 hosts. With 40 half-rows in the data center, that is 640 hosts in the same SAN (see Figure 2). “The enabling technology for this massive scalability is the VSAN support in Cisco MDS switches,” says Angulo. “With VSAN capability, the administrative traffic from only a fraction of the hosts would overwhelm the SAN.” Each storage frame can house 80 terabytes, and with two storage frames the SAN can serve 180 terabytes, or more than 10 percent of all Cisco storage.

**Figure 2.** Planned Architecture for Engineering Data Center: Capacity for 640 Hosts



When Cisco begins consolidating its storage and hosts across business units—to reduce capital costs, improve resource utilization, and simplify management—VSANs will be invaluable in separating traffic from different business units. That is, VSANs will enable Cisco to isolate its business logically instead of physically. A single Cisco MDS 9509 Multilayer Director Switch can support multiple VSANs—for instance, one for tape backups, another for engineering, and another for enterprise resource planning (ERP). Isolation will be increasingly important when Cisco combines business verticals into one infrastructure, because Cisco does not want ERP systems to see traffic associated with the engineering system, both to eliminate unnecessary administrative traffic and to isolate the respective systems from problems the others might experience. “Traditionally, the engineering SAN would need a physical connection to access the tape backup devices,” says Angulo. “But with this setup, you have connectivity at a virtual abstract layer. Our eventual goal is to have massive storage, in any location, that any host can provision as needed.”

Other upcoming steps include installing Cisco MDS-based SANs in engineering data centers in Bangalore, India; South Netanya, Israel; Research Triangle Park, North Carolina, USA; and Boxborough, Massachusetts, USA. SANs are near completion in Cisco business data centers in other locations. It is a decision made confidently, based on success. "This isn't a lab," says Angulo. "It's a real-world, P1 environment. If something didn't work, thousands of engineers at Cisco headquarters in San Jose would stop working and development of Cisco IOS Software would grind to a halt. The 100 percent redundant deployment gave us the confidence we needed to migrate."

## Contacts

For more information about the integration of Cisco MDS and other vendor SAN equipment, or other Cisco production deployments of Storage Area Networking in the Cisco IT Data Center environment, start a conversation with the IT experts in the Storage Networking forum at:

[http://forums.cisco.com/iframe/servlet/SCom?page=scom&folderID=caw&CommCmd=MB?cmd=display\\_messages&mode=new&location=.ee77368](http://forums.cisco.com/iframe/servlet/SCom?page=scom&folderID=caw&CommCmd=MB?cmd=display_messages&mode=new&location=.ee77368)

## FOR MORE INFORMATION

To read the entire case study or for additional Cisco IT case studies on a variety of business solutions, visit Cisco on Cisco: Inside Cisco IT [www.cisco.com/go/ciscoit](http://www.cisco.com/go/ciscoit)

## NOTE

This publication describes how Cisco has benefited from the deployment of its own products. Many factors may have contributed to the results and benefits described; Cisco does not guarantee comparable results elsewhere.

CISCO PROVIDES THIS PUBLICATION AS IS WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some jurisdictions do not allow disclaimer of express or implied warranties, therefore this disclaimer may not apply to you.



**Americas Headquarters**  
Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
[www.cisco.com](http://www.cisco.com)  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 527-0883

**Asia Pacific Headquarters**  
Cisco Systems, Inc.  
168 Robinson Road  
#28-01 Capital Tower  
Singapore 068912  
[www.cisco.com](http://www.cisco.com)  
Tel: +65 6317 7777  
Fax: +65 6317 7799

**Europe Headquarters**  
Cisco Systems International BV  
Haarlerbergpark  
Haarlerbergweg 13-19  
1101 CH Amsterdam  
The Netherlands  
[www-europe.cisco.com](http://www-europe.cisco.com)  
Tel: +31 0 800 020 0791  
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).

©2007 Cisco Systems, Inc. All rights reserved. CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, Packet, PIX, ProConnect, ScriptShare, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0705R)