

How Cisco IT Uses a SAN to Automate the Legal Discovery Process

Cisco dramatically reduces the cost of legal discovery through better data management.

Cisco IT Case Study/Storage Networking/SAN for Data Management: This case study describes Cisco IT's internal deployment of a storage area network (SAN) within the Cisco network, 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

“Cisco’s goal is to win the litigation we enter into, but to win efficiently and without overspending. To that end, we dramatically streamlined our discovery process by devising an online system that helps us work more effectively with external resources while minimizing disruption to the Cisco employees who are information owners.”

Mark E. Michels, Attorney, Office of the General Counsel, Cisco

BACKGROUND

Large enterprises typically produce, collect, and analyze large volumes of data during day-to-day operations. Corporate legal departments are no exception, particularly those involved in the discovery process for litigation and regulatory disclosure.

In companies such as Cisco, the amount of discovery data for a single legal case can be staggering. Cisco staff must mine many different network storage devices, each of which can hold a small amount of usable information among gigabytes of irrelevant data, and prepare the selected data for legal review. Until the late 1990s, nearly all discovery-related information was paper based. Now, at Cisco, 90 percent of that information is in electronic form, but the

challenges of preparing it for review are very similar to the challenges of preparing paper documents.

CHALLENGE

The Cisco Office of the General Counsel has a lean and active litigation team composed of two attorneys, two paralegals, and one internal IT consultant. The group works with most business units and technical groups throughout the company. At any time, the group manages both active matters or lawsuits and numerous credible threats of litigation.

Conceptually, at least, legal discovery at Cisco is a simple process that consists of collecting data, processing it for review, reviewing it, and delivering it to other parties or government regulators. During data collection, information is gathered from numerous sources, including:

- Structured databases, such as the company's enterprise resource planning (ERP) system
- Pages on internal and external Cisco websites
- Backup storage tapes
- Printed documents
- Individual user data
- Central Cisco document repositories

EXECUTIVE SUMMARY	
BACKGROUND	<ul style="list-style-type: none"> Legal discovery process involves retrieving, storing, and processing huge volumes of electronic data
CHALLENGE	<ul style="list-style-type: none"> Store, organize, and process data related to legal cases Minimize disruption to Cisco employees outside of the legal department Improve productivity of resources used to analyze and process the discovery data
SOLUTION	<ul style="list-style-type: none"> Create a SAN based on a Cisco MDS 9000 Family storage networking switch Cisco Fabric Manager application
RESULTS	<ul style="list-style-type: none"> Cost reduction of 64% in total litigation expenses, and 30% in outside legal reviews Improved utilization of storage resources Ability to handle growing data volumes without hiring additional employees Information can be shared securely with outside counsel and contractors
NEXT STEPS	<ul style="list-style-type: none"> Automate more tasks for legal discovery Improve data access for outside counsel Share information more easily across Cisco

Following the collection process, all retrieved data is copied to Cisco's litigation repository. Once in the repository, files are put into a production process to prepare them for legal review. This process encompasses activities such as the following:

- Email messages are converted into separate files and paired with their attachments
- Application help, read-me, and log files, are removed (a process that can reduce data volume by 70 percent, dramatically reducing the time and cost of legal review)
- The file set is scanned for duplicates, which are tagged and removed
- All files are organized, and metadata tags describing file type, size, and date are added
- Each file is assigned a unique document number for ease of review and retrieval
- The completed data set is moved to the production server

Once the production is complete, the data set is ready for Cisco's outside counsel, who reviews it online and marks it as responsive (relevant to the legal issue at hand) or nonresponsive. Outside counsel also separates proprietary and privileged material for special handling. These marks, the reviewer's identification, the project status, and any reviewer notes are included in the data set.

The data set is output as individual files, converted to TIFF files, burned to CDs or DVDs, or stored in a hard drive. The data is then given to opposing counsel for review.

The biggest challenge facing the Cisco litigation team is the need to store, organize, and process huge volumes of data. A related challenge, once data is assembled, is to retrieve specific information quickly. "Essentially, we build a mountain of information and tear it down," says Mark Michels, co-manager of Cisco's litigation team. "We tend to over-collect, sometimes retrieving gigabytes of material of which very little is actually relevant, so that we only have to do the job once. Our objective is to minimize the disruption to our employees and to get the best possible value from outside counsel and from the contractors who extract the data and prepare it for legal review."

"With our e-discovery process, we've put in place a system that avoids the disorganization and inefficiency that is rampant in discovery," says Jeff Ghielmetti, manager of computer forensics and legal discovery for Cisco's litigation team. "As an analogy, visualize a person who visits his tax accountant with a shoebox crammed to bursting with bills and receipts. Our ultimate goal in building the system was to put each piece of information in the right file folder, in the right order, where it could be easily found."

SOLUTION

Cisco's approach to the challenge posed by discovery revolved around finding a way to build a large storage system that was easy to install, configure, and manage. The litigation group also wanted a system that could expand to accommodate an escalating workload and meet performance expectations well into the future. The group was conscious that, as a corporate cost center, it needed to develop a smart, frugal solution that would minimize impact on the company's bottom line, in part through the use of existing technology to meet these challenges.

The group used a model it called DMAIC (define, measure, analyze, improve, and control) to develop a solution. The

model offered a structured, linear process:

- **Define:** Delineated every aspect of the current discovery process and calculated to the value of fixing each piece
- **Measure:** Quantified Cisco's e-discovery costs per attorney and information owner
- **Analyze:** Examined inefficiencies that could be traced to outside counsel and other legal services vendors
- **Improve:** Suggested the tools and mapped the processes that would be required for true e-discovery
- **Control:** Looked at the way that improvements could be sustained and improved in the future

A Cost-Effective Architecture

Discovery data at Cisco comes from literally thousands of sources: shared data servers, the ERP solution, multiple websites, documents on PC hard drives, databases, and the mailboxes of some 37,000 employees worldwide. The e-discovery system links Cisco, litigation support consultants, and outside counsel through a storage area network

(SAN) built on a Cisco MDS 9000 Family SAN switch and a powerful, but low-cost storage system. Essentially, this e-discovery architecture is a "network of networks" that connects Cisco to outside counsel and other discovery partners (Figure 1).

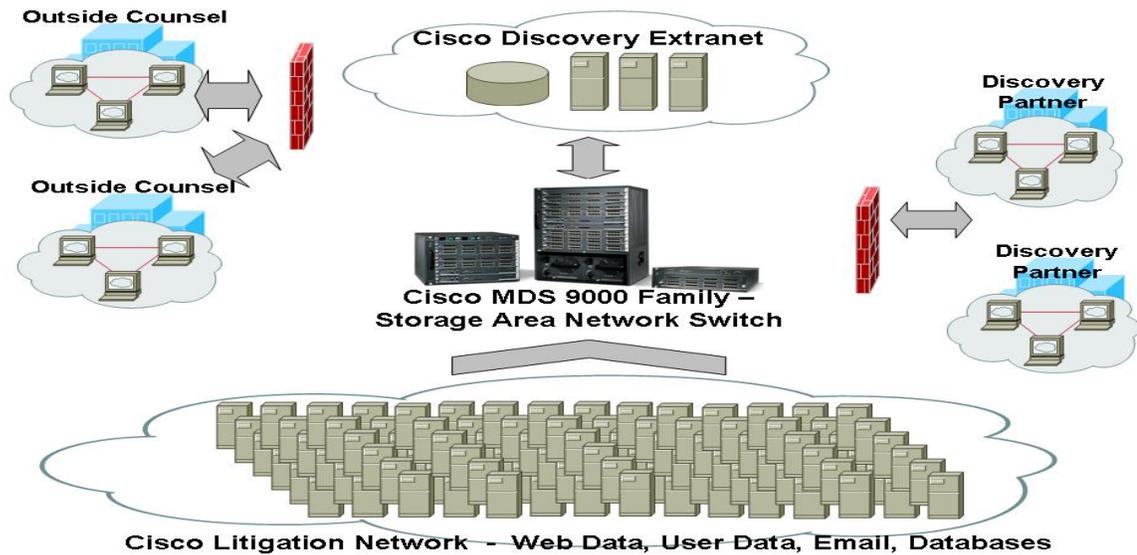
PRODUCT LIST
Storage Networking <ul style="list-style-type: none">• Cisco MDS 9000 Family SAN switch• Cisco Fabric Manager application

An easy-to-use SAN management application virtualizes the disks in the storage array, allowing the litigation team to "see" data as one

large volume rather than many separate disk drives. Virtualization also makes it easier for the team to use the system, and creates operational efficiencies for the e-discovery administrator.

"When we set out to design the system, we quickly saw the tremendous potential of the Cisco MDS switch, which allows us to grow easily in response to future discovery needs," says Ghielmetti. "Traditional network-attached storage was too slow to handle the volumes of data we were expecting to process. And using SCSI-based devices was far too expensive and labor-intensive, requiring one server for every two disk drives. The large, established vendors proposed solutions that ran into millions of dollars. Using a foundation of our own technology, we found a solution with the right price, capabilities, and flexibility. In fact, it let us consolidate data on 27 separate drives into a single massive volume that is much cheaper and easier to manage."

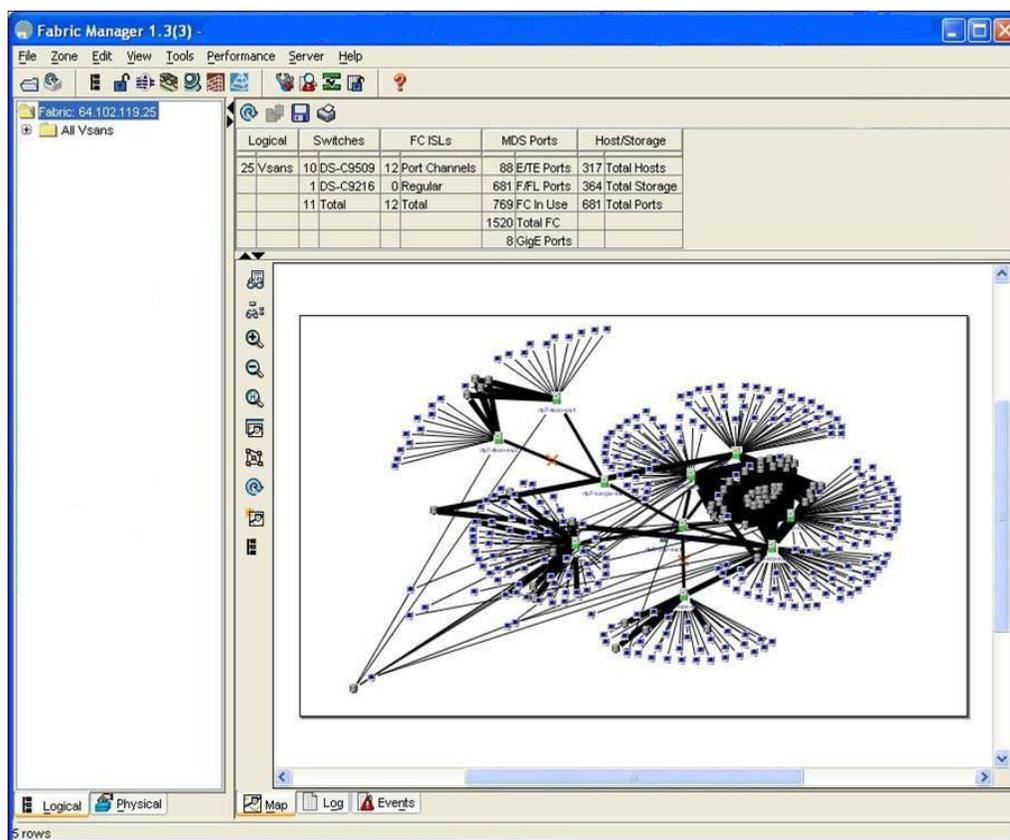
Figure 1. Design of the Litigation Repository SAN at Cisco



Management and Security

Security for a system of this sort is challenging and requires ingenuity in its design. The Cisco Fabric Manager application helps the e-discovery SAN to operate rapidly and efficiently (Figure 2). The entire fabric of the SAN is integrated with Cisco network security and with the firewalls of Cisco's outside counsel.

Highly sensitive business, financial, and technology information must be shared with confidence across the team, so secure transmission is extremely important. The Litigation Lab exists on its own subnet of the Cisco network and is connected to a switch controlled by the Office of the General Counsel, isolating the data so effectively that no one without authorized access knows that the subnet exists. In addition, management applications related to e-discovery incorporate basic user authentication based on Lightweight Directory Access Protocol (LDAP) and Cisco's deployment of Microsoft Active Directory. In daily operation, these security arrangements eliminate the challenge of coordinating security policy through disparate firewalls.

Figure 2. Cisco Fabric Manager Map of the Litigation SAN

RESULTS

Initial results from implementing e-discovery at Cisco have encouraged the Office of General Counsel to continue refining the solution. It is much more efficient to search for data in one location rather than in multiple separate systems, and when litigation team members are more efficient, they can more easily handle a growing workload. “It’s vital that we developed such an effective system to deal with electronic data. The volume is accelerating so rapidly that we need the team to be able process more in the same number of hours,” says Michels. “We had planned for a 30-percent increase in volume per year, but recent cases we’ve received have caused us to revise our estimates upward to 100 percent.”

Cost avoidance figures have also been encouraging. In 2002, discovery comprised 67 percent of Cisco’s total litigation expenses in five significant cases, as benchmarked against industry-standard methods of conducting electronic discovery. As of the writing of this case study, Cisco has been able to reduce its costs of discovery by approximately 97 percent (Table 1), for an overall reduction in litigation expenses of 64 percent.

Table 1. Cisco Cost Savings Due to Automated Litigation Discover Process

Matters	Estimate (* First Year)	Actual Cost	Savings
1	US\$9,500,000	\$571,711	\$8,928,289
2	\$9,500,000	\$71,711	\$9,428,289
3	\$1,800,000*	\$99,211	\$1,700,800
4	\$400,000*	\$18,554	\$381,446
5	\$3,000,000*	\$68,604	\$2,931,396
Total Savings	\$25,200,000	\$1,829,780	\$23,370,220

Cost avoidance figures have also been encouraging. The storage solution cost one-third of a standard enterprise storage array and allows Cisco to cheaply and easily add capacity as needed. At the same time, the litigation team has reduced the expense of outside legal review by 30 percent, and by using the abilities of these outsourced providers, has been able to avoid hiring 20 full-time employees.

NEXT STEPS

With the “network of networks” in place and functioning smoothly, Cisco’s next step is to automate more of the tasks involved in discovery. The litigation team’s goal is to increase efficiency, reduce litigation expense, and continue to build the confidence of the users of the system. Specifically, by using SAN technology, the Office of the General Counsel plans to connect outside counsel to a central data set. Benefits of this network extension will include:

- A 35-percent reduction in onsite support of the remote networks use by outside counsel
- Increased security of information shared with outside counsel
- A common database of document review and productivity metrics
- Multifirm collaboration without synchronization issues

The litigation team also plans to work with other organizations within Cisco to enable the company to more easily share information across the enterprise.

“When we set out to design the system, we quickly saw the tremendous potential of the Cisco MDS switch, which allows us to grow easily in response to future discovery needs.”

Jeff Ghielmetti, Manager of Computer Forensics and Legal Discovery, Cisco

The Promise of Better Data Management

In essence, the more data that Cisco, or any company, has, the harder it is to locate, assemble, and put to effective use. The e-discovery system built by the Office of the General Counsel illustrates the power of SANs and is relevant to the records management needs of any large enterprise or municipality.

“SANs are unequalled in their ability to enable organizations to work with data that was once difficult to manage because of its sheer volume or geographic dispersal,” says Ghielmetti. “In the legal department at Cisco, we have taken the first step toward making the information we rely on accessible, and minimizing the sheer bulk of data that needs to be produced in litigation.”

Companies typically preserve documents for many reasons. They have substantial legal and regulatory compliance responsibilities (the Sarbanes-Oxley Act is a particularly relevant example). They want to preserve software development records or other documentation that contributes to a robust corporate knowledge base. Unfortunately, they may also save documents because they are unsure of how to determine what is useful.

Although data storage costs have declined, there is significant expense in managing data and determining what is important to save. “This is not a new issue, and it will grow as our storage ability grows,” says Michels. “Think of a

library versus a warehouse. A library is a deliberate selection of useful information, but a warehouse can be a collection of unsorted data. Traditionally, data was usually analyzed before it was boxed and labeled. In today's electronic world, every document a company has can be saved relatively easily, which is not always to our benefit. Cisco approaches data mining and management in a creative way that can benefit companies across the board. That's what we've done in the Office of the General Counsel in our dialogue with our document retention team."

A well-designed SAN allows a company to make data access transparent to the user and, equally important, allows managers to apply record retention and other policies uniformly to large repositories of data across the enterprise. With the experience of the Office of the General Counsel behind it, Cisco anticipates continued progress in mining its corporate knowledge base and applying it to daily operations.

FOR MORE INFORMATION

For additional Cisco IT case studies on a variety of business solutions, visit Cisco on Cisco: Inside Cisco IT 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.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

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.

CCDE, CCVP, Cisco Eos, Cisco StadiumVision, the Cisco logo, DCE, and Welcome to the Human Network are trademarks. Changing the Way We Work, Live, Play, and Learn is a service mark; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCI, 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, Collaboration Without Limitation, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Event Center, FastStep, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, IQ Expertise, the iQ logo, IQ Net Readiness Scorecard, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo 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. (0901R)