SYS-CON

ISSJ PREMIERING THIS FALL!

information STORAGE+ SECURITY journal

www.ISSJournal.com

PREVIEW 2004 VOL.1 ISSUE 1

Coming Up: in the Premier Issue!

- Virtualization
- Logging
- Data Centers
- Identity Theft
- DLM
- Authentication
- SAN/NAS Convergence

Implementing RSA's SecurID for Microsoft Windows

SAN NAS

Emerging technology trends and market maneuvers

COMPLIMENTARY COPY • COMPLIMENTARY COPY • COMPLIMENTARY COPY • COMPLIMENTARY COPY •

SPECIAL SNEAK PREVIEW OF ISSJ

$5.99US $6.99CAN

07>

0 71486 01793 6

INFITECH

X5 NAS

empower your data network

High Performance Rack Mount Servers and Storage Solutions

> Simplify your network: X5 NAS will replace your file servers for Microsoft, UNIX and Apple clients. Manage a single network storage box vs. three legacy file servers. When more storage is required, simply plug another X5 NAS to an open network port.

> Remote, secured management: X5 NAS can be configured, maintained and monitored from anywhere in the world, as long as you have connection to the Internet. Use secured, HTTP(S) access for protection against unauthorized access.

> Faster access, more simultaneous clients: X5 NAS has proven to be faster and more responsive. Due to its optimized embedded OS, X5 NAS will outperform traditional file servers exponentially. Faster means more simultaneous users and getting jobs done quicker.

> Robust & highly available: Embedded OS, high quality hardware components, continuous on-going reliability test makes X5 NAS extremely reliable. Furthermore, its true server-to-server mirroring and real-time fail-over, makes X5 NAS the most highly available storage solution.

> Server to Server Fail-Over & Mirroring
> Snap Shot Data Recovery
> Embedded OS
> RAID 0,1,5,10, and JBOD
> SATA, PATA and SCSI HDD Support
> Hot Swap HDD and PSU
> SCSI/Fibre Channel Subsystem Support
> PDC/ADS/NIS/Host IP Blocking
> Dual Gigabit NIC with Fail-Over
> Up to 3TB in 3U
> 64bit, PCI-X for I/O

Powered by NetEngine

Visit Us www.infi-tech.com

or Call 1-800-560-6550

to Find Out More

information STORAGE+ SECURITY journal

President and CEO
Fuat Kircaali fuat@sys-con.com

Vice President, Business Development
Grisha Davida grisha@sys-con.com

Group Publisher
Jeremy Geelan jeremy@sys-con.com

Advertising

Senior Vice President, Sales and Marketing
Carmen Gonzalez carmen@sys-con.com

Vice President, Sales and Marketing
Miles Silverman miles@sys-con.com

Advertising Sales Director
Robyn Forma robyn@sys-con.com

Director, Sales and Marketing
Megan Ring megan@sys-con.com

Advertising Sales Managers
Kristin Kuhnle kristin@sys-con.com
Beth Jones beth@sys-con.com

Editorial

Executive Editor
Gail Schultz gail@sys-con.com

Associate Editors
Nancy Valentine nancy@sys-con.com
Jamie Matusow jamie@sys-con.com
Jennifer Van Winckel jennifer@sys-con.com

Online Editor
Lin Goetz lin@sys-con.com

Production

Production Consultant
Jim Morgan jim@sys-con.com

Art Director
Alex Botero alex@sys-con.com

Associate Art Directors
Louis F. Cuffari louis@sys-con.com
Richard Silverberg richards@sys-con.com
Tami Beatty tami@sys-con.com

Web Services

Vice President, Information Systems
Robert Diamond robert@sys-con.com

Web Designers
Stephen Kilmurray stephen@sys-con.com
Christopher Croce chris@sys-con.com

Accounting

Financial Analyst
Joan LaRose joan@sys-con.com

Accounts Receivable
Charlotte Lopez charlotte@sys-con.com

Accounts Payable
Betty White betty@sys-con.com

Customer Relations

Circulation Service Coordinators
Shelia Dickerson shelia@sys-con.com
Edna Earle Russell edna@sys-con.com
Linda Lipton linda@sys-con.com

Editorial Offices

SYS-CON Media,135 Chestnut Ridge Rd.
Montvale, NJ 07645
Telephone: 201 802-3000 Fax: 201 782-9638

Worldwide Newsstand Distribution
Curtis Circulation Company, New Milford, NJ

For List Rental Information:

Kevin Collopy: 845 731-2684
kevin.collopy@edithroman.com

Frank Cipolla: 845 731-3832
frank.cipolla@epostdirect.com

Newsstand Distribution Consultant
Brian J. Gregory/Gregory Associates/W.R.D.S.
732 607-9941, BJGAssociates@cs.com

SYS-CON MEDIA
The World's Leading i-Technology Publisher

From the Group Publisher

A Necessary Marriage of Technologies

BY JEREMY GEELAN

ANYONE WHO HAS worked with data-intensive computing and storage environments over the past five or so years has seen the quantity of sensitive data that organizations and enterprises carry on their servers and storage devices spiral upwards. Exponential growth in storage capacity, coupled with emerging regulatory requirements, has led to greater emphasis on storage network vulnerabilities.

It has become the goal of many enterprises to achieve efficiencies and cost reductions by making back-end data available to its staff anytime, anyplace, and on any device. Yet the Internet, precisely because it is ubiquitous and flexible, is replete with security concerns.

It is high time, many IT professionals have said, that storage and security issues were dealt with simultaneously. The quality of thinking and writing on this topic that you will see in ***Information Storage & Security Journal*** is evident in this preview issue.

As Diana Kelley from Computer Associates says: "Storing data without taking into consideration the security requirements and potential threats is not sufficient in today's enterprise. Legal requirements, audit needs, and shareholder interest all demand that corporations not only protect live data, but log, archive, and store critical, historical data in a safe and retrievable manner. Storage and security are intimately linked." Oracle's Chief Security Officer, Mary Ann Davidson, reinforces the synergy: "A recent report from PricewaterhouseCoopers," she writes, "confirmed that most security breaches occur in stored data."

Yet the prospects aren't all negative. IT professionals – and IT these days might just as well mean Infrastructure Technology – will be encouraged, we hope, by the article from Mark Griffiths, VeriSign's VP of Authentication Services, on how identity theft could soon become a relic of a bygone era. This would be a significant breakthrough, given that the 2002 Federal Trade Commission's annual study on consumer complaints cited ID theft as the most frequent reason individuals contacted consumer protection authorities.

John Worrall peeks over the horizon from his vantage point at RSA Security and tells us that not only present technologies like two-factor authentication and smart cards, but also emerging ones like mass-market biometrics, need to become standard. SGI's Laura Shepard answers the questions that we believe many ***ISSJ*** readers should and will be asking: What is information lifecycle management and where is it in the evolution of migration solutions? She also does a great job in explaining how it differs from data lifecycle management (DLM).

Andrew Bulkley, of GE Security, reminds us of the role that standards play – particularly in how they are essential in helping to make access control agree with a company's disparate systems.

"Proactive security" is the main focus of Eric Vishria's article on a new breed of technology – IT automation software. Working "from the inside out," comprehensive automation systems "can take into consideration the people, processes, and technology that can turn even the most complex environments into truly impenetrable targets." In other words, automation software complements perimeter defense systems by reducing the chance for human error and keeping systems up-to-date automatically.

Whatever your position in the industry, hold on to your hat: storage and security are becoming more and more enmeshed, and ***ISSJ*** will be there to help deliver storage subject matter in context with popular security applications, and vice versa. Our aim is to guide, motivate, and inspire senior IT and business management leaders in the planning, development, deployment, and management of successful enterprise-wide security and storage solutions.

As the next generation of enterprise networks arrives, and as the protection and management of data in heterogeneous environments becomes increasingly important, from the Fortune 500 to small and medium-sized businesses, SYS-CON Media is pleased to bring its decade of print and online content excellence to this expanding field.

About the Author

Jeremy Geelan is group publisher of SYS-CON Media, and is responsible for the development of new titles and technology portals for the firm. He regularly represents SYS-CON at conferences and trade shows, speaking to technology audiences both in North America and overseas.
jeremy@sys-con.com

Integrating Access Control with Other Systems

THE NEW NECESSITY

BY ANDREW BULKLEY

LET'S REVIEW THE TYPICAL ACCESS control system in use today. Not only are the various components disconnected but they are from different manufacturers and do not and will not integrate with each other. Some use incompatible hardware, or proprietary, unsynchronized databases, or completely inconsistent user interfaces that compete for space and attention. This system is inefficient and requires just too many people to manage it. It is not the kind of system that will make people in management very happy.

They know that such systems increase employee and training costs, foster unnecessary equipment expense, cause security and safety breaches, and produce mission-critical downtime. Since their budgets and management are beginning to dictate what will be used, access control systems must conform.

Today, although companies need to integrate all of their security and facility systems under one control system, they also have special integration requirements due to the size and deployment of their personnel. Basically, access control systems need to be linked to personnel (Human Relations – HR) systems to control which employees are currently employed by the company. The linkage of these systems ensures that as employees are terminated or re-assigned, the access control is completely synchronized with the personnel moves, without manual intervention.

For instance, Noridian Administrative Services LLC is a regional claims contractor for the U.S. Medicare program, processing Medicare claims for the states of North Dakota, Minnesota, South Dakota, Iowa, Wyoming, Colorado, Alaska, Oregon, Washington, Nevada, Arizona, and Hawaii from its Fargo, North Dakota, headquarters.

Noridian has put together a world-class integration system using the GE Security Secure Perfect 4.O Enterprise as its security platform, which integrates into the organization's PeopleSoft system used for human resources.

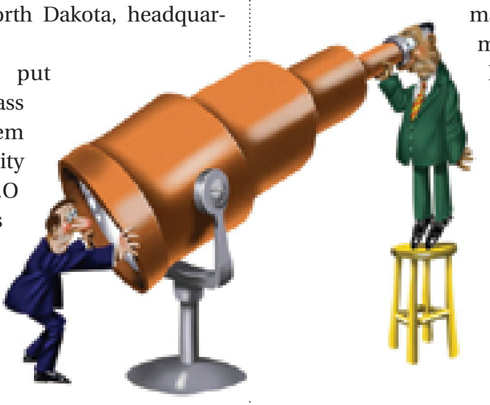

In this integrated system, Secure Perfect pulls down certain fields, such as first name/last name/employee ID number/employee status, from PeopleSoft, not the access control system, so there are no variances

Leveraging such technology breakthroughs and foreseeing a need for increased security, companies will also begin to rapidly adapt smart cards, biometrics, and intelligent video into both their physical and logical access control systems. As a result, both security and IT managers will be faced with greater system complexity and forced along the pathway of integrated business solutions. These have become inarguable facts.

Creating a Command and Integration Platform

An integration platform is needed to bring all of these systems together because most companies have a wide variety of manufacturers' equipment installed. Different sites standardized on different manufacturers. What's needed is a complete command and control integration platform that integrates all aspects of security and facility management within a single screen.

Such a platform must provide a completely open architecture with published APIs, plug-and-play compatibility, cross-platform support, adherence to industry standards, and the ability to seamlessly create a modular facility environment. With it, you would have a single, intuitive, integrated console that lets you protect and manage your business.

Defining the Platform

First, the platform would need to be tightly integrated with the security management system, offering advanced access control, alarm monitoring, intrusion detection, fire alarm, intercom and personal safety/duress systems, credential production, and employee and visitor management functionalities. Additionally, though, the platform must address and enhance security management system capabilities by integrating digital video from multiple manufacturers as well as integration and support for fire, intrusion, personnel, and complete facility management.

A command and integration platform needs to provide a single window on the enterprise. Today's managers must employ a centralized, consistent user interface for managing security and facility alarms and events across the entire company. That's not to say there can't be delegation. A company most certainly might want multiple, separate security systems for administration purposes but still maintain centralized control. For example, an Asian-based company may want system hosts in North America, Europe, and Asia.

Nonetheless, the security director for each region can be delegated the task of configuring doors, managing employee access, integrating with specific alarms and other tasks within their regions. Multiple synchronized, geographically dispersed servers are also important for alarm monitoring. Each host or region can do its own alarm monitoring while, from the main server, the head security management team can monitor all regions from a single screen. Yet, the master control in Asia will still have ultimate control of all three servers.

Seamless integration would mean the physical access control departments, as well as other groups in the enterprise, would have the freedom to select different technology vendors, relying on the command and integration control platform to handle the integration.

A New Respect for Standards

In this new world proprietary is a bad word. Multivendor support is only achievable through the use of IT industry standards such as XML, TCP/IP, SNMP, LDAP, and SMTP. The platform must support commercial off-the-shelf operating systems such as Red Hat Linux and Microsoft Windows in its many flavors; database platforms such as Microsoft SQL Server, MSDE, Informix, IBM DB2 Universal Server, and Oracle Server; user directories such as LDAP and MS Active Directory; networks such as Ethernet; report generators such as Crystal Reports; and common administrative utilities for system backups and fault tolerance. Likewise, it must seamlessly integrate with

Card Swipe: We're entering a whole new world of integrated intelligent access sytems.

Picture Perfect: Digital video from multiple manufacturers needs integrating external applications, such as time and attendance systems, and peripheral devices such as printers.

Only then will enterprises be able to achieve real time, bidirectional data exchange and actions between security systems and other infrastructure and applications, including HR and ERP systems. Management of people's access rights will be streamlined with policy-based management across physical and logical security. With one step, an enterprise can set up or delete a complete set of access rights for any employee.

We recently introduced such a platform. The GE Security Facility Commander integrates security and facility management into one system. All applications, presently disjointed, can be viewed within a single, common, easy-to-use interface. All hardware, even video, alarm and printing equipment, works seamlessly within its framework. The ability to enter security and facilities data just once, and have the framework synchronize with existing legacy systems automatically, is now a reality.

"We're entering a whole new world of integrated, intelligent access control systems"

How It Works

A standards-based command and control integration platform, Facility Commander lets organizations integrate multiple aspects of their security and facility management within a single screen. This single, intuitive interface provides one console for all access control, video surveillance, and alarm management functions. Built-in drivers support GE's Picture Perfect and Secure Perfect access control systems as well as digital CCTV, analog CCTV switchers, intrusion, and intercom systems.

Users and integrators can use its included drivers for multiple access control, digital video surveillance, alarm, and other security and facility management systems. Or, they can use the Facility Commander System Developers Kit (SDK) and open APIs to develop plug-and-play drivers for their existing digital video equipment and software.

- **Access Control:** Facility Commander is closely integrated with GE Security's Picture Perfect and Secure Perfect security management systems. It integrates events and actions between access control and other systems. For instance, users can lock or unlock a door from a graphical map of door locations. If an access alarm is triggered, Facility Commander can map the location and direct the surveillance system to begin recording.
- **Digital Video:** Facility Commander works with video surveillance systems from multiple manufacturers. From the console, users can view live images from surveillance cameras, control pan/tilt/zoom cameras, or search for video clips stored on digital video recorders (DVR) by time, date, event, event type, camera, or DVR. When an event or alarm is triggered, Facility Commander 2.0 can tell the DVR to begin recording, display live video from a linked camera at the location, map the alarm location, and send an e-mail to the security director.
- **Analog CCTV Switcher:** Even if a user's present video system employs analog equipment, Facility Commander will work with it by automating camera call-up on specific monitors when events and alarms occur.
- **Intrusion System:** From the Facility Commander console, users can arm or disarm an intrusion area. When an event or alarm is triggered, Facility Commander's pop-up alarm function displays a map of the alarm location and links it to the digital video system to begin recording at that location. Users can retrieve video clips by associating them with alarms to investigate and resolve incidents more quickly.
- **Intercom System:** When a call comes in from an intercom, Facility Commander can automatically trigger the Event Action Mapping function to display the intercom call station on the console. The intercom can be connected to the video system to show live video from that call station. Users can also link the intercom to the access control system to unlock or lock a door at that location.

Third-party security vendors will find Facility Commander easy to use. It features an open architecture based on industry standards. It runs on commercial off-the-shelf operating systems including Windows, Linux, and AIX. It supports popular databases such as SQL, Informix, DB2, and Oracle. With its SDK and open APIs, vendors can create their own drivers. Indeed, they could even interface it with Mr. Meyers' Alliance Platform and give him extra eyes for his jewelry store.

Integration Is No Longer a Luxury; It's a New Necessity

With the convergence of physical access control and other security and IT systems, new open system architectures are providing smaller users as well as global enterprises with the solutions they need. We're entering a whole new world of integrated, intelligent access control systems. ■

About the Author

Andrew (Andy) Bulkley is senior director of product strategy for GE Security, Enterprise Solutions. He is a veteran of the security industry.
andrew.bulkley@gesecurity.com

DATA CENTER ABOUT READY TO BLOW?

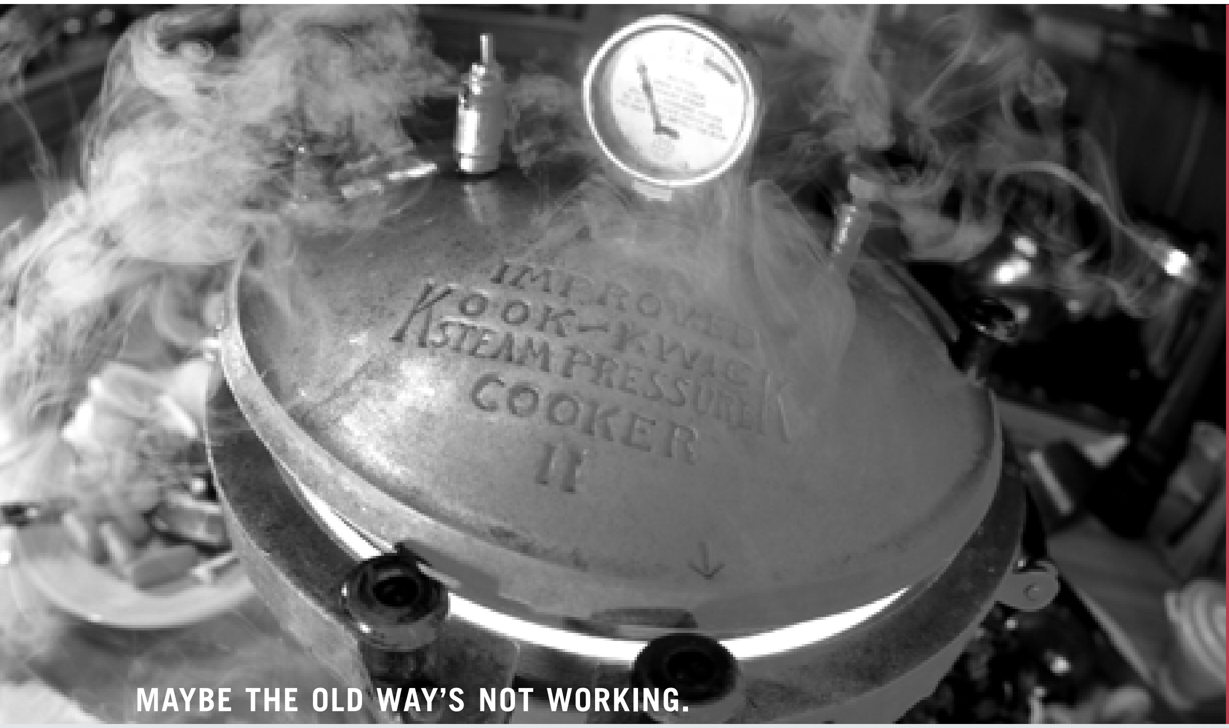

Rackable Systems servers—now with Distributed DC Power—improve thermal management, reliability and power efficiency.

GREATER RELIABILITY & MINIMIZED MAINTENANCE

Now with the option of Distributed DC Power, Rackable Systems server clusters are more reliable, flexible and efficient than the competition. By distributing redundant DC Power to each server—and replacing the standard AC power supply with a 70 times more reliable DC card—availability is significantly improved and servers are easier to maintain. They also run cooler, leveraging the low power consumption and extraordinary performance of the AMD Opteron™ 200-Series processors to provide the flexibility and scalability required for demanding enterprise-class applications.

THERE IS A BETTER WAY

As pioneers of large-scale data center solutions, Rackable Systems continues to lead the industry with innovative server designs for Electronic Arts, Google, Lawrence Livermore Labs, nVidia, SLAC, Tellme Networks, Toshiba America Electronic Components, Webex and Yahoo!. Our open architecture configurations are built-to-order to ensure maximum flexibility and ease of serviceability. Our innovative cabinet design features twice the density at a fraction of the cost, while optimizing thermal efficiency. And with Rackable Systems' new Distributed DC Power technology, you've got a solution that can't be beat.

Rackable® systems

Flexible Configurations

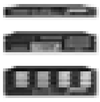

- Customized & built-to-order
- Open architecture
- Superior remote management

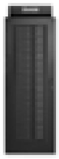

Improved Reliability

- Redundant power & drives available
- Thermally efficient & optimized
- Reduced data center costs

Easy Deployment

- Fully configured & tested
- Pre-cabled, pre-racked
- Quick ship

Learn more at **www.rackable.com/dcpower** or call 1.866.548.2527.

AMD 64 Opteron™

IN WITH THE NEW

Rackable® systems

Storage and Security Management for Logging and Archiving

BEING A PACK RAT IS NO LONGER AN OPTION

BY DIANA KELLEY

THERE'S NO STEMMING THE TIDE of information; with more users and more servers and more connectivity than ever before, the task of logging, storing, and archiving all of that activity is astounding. The temptation may be to simply, save it all.

Recent legislation has placed a demand on security professionals to log and archive massive amounts of data. The default plan for being prepared when audit and forensics investigators come knocking is to have everything logged and backed up – somewhere, somehow. But keeping a copy of every single event, every file, every document, may not be feasible. Storage has certainly become cheaper, but it's not free. And management of an overloaded SAN can introduce inefficiencies and potential security vulnerabilities into the process. In this article I'll take a look at the synergy between security and storage as they contribute to keeping an organizations logs and archives in hand and on-demand.

Introduction

Who among us doesn't have a bit of pack-rat mentality in them? But the reality is that the chaos and confusion resulting from so much storage doesn't decrease the risk, it simply makes for ineffective clutter. What should be kept? What is the value? What are the threats?

By answering these questions organizations can begin to understand how to balance security and storage requirements, especially as it relates to critical log data. Keep enough, and the company will have financials ready for audit and escape going to jail for violation of regulations such as SOX and CSB 1386. Keep too much, and the cost of storage and resources needed to archive and manage all of the old information could affect the corporation's profitability. Worse still, if the volumes of data aren't managed properly, when it does come time for an audit, finding the correct information could mean weeks of hunting through terabytes of information and, potentially, never finding it at all.

So what can we do? How does the data storage affect the overall enterprise's security posture? And what can we do to get the data at hand and on demand?

Determining How Much is Enough

One of the first steps is to identify the types of information that will be critical in the future. There are a few basic rules that a company can employ to decide which items need to be saved and which don't. Take, for example, old versions of a document such as a press release. The draft is sent around to a number of people, marked up and re-distributed, and then finalized and put out on the wire. Do all of the people who were associated with the release need to keep all of the versioned copies? Probably not. But if the users have saved these versions in their Inboxes, then it's a good bet the company is paying to back up and store all of them.

The cost of data storage varies based on the ways in which it will be accessed later. Offloading files to a series of DATs that sit in a box on a shelf somewhere is going to cost far less than keeping files in physically secure areas, in encrypted format on always available repositories on a SAN (storage area network). So, old copies of log files from testing and prototype machines may lend themselves well to less expensive storage methods than the log files from the corporation's production mail server.

While the final determination of data valuation depends on each company's own business requirements, the following considerations will help with the calculation:

- How current is the data?
- How frequently is it used?
- How much did it cost to accumulate/ generate?
- What impact does it have on the business?
- How much does the company profit from the data?
- What would the company lose if the data wasn't available?
- If lost, how much would the company have to spend to get it back?

Safety and Access Control

Once the data is valued, the threats and safety requirements for the data must be determined. To do this, first understand the types of threats that can put the data at risk, the ease with which they can be executed, and the cost of the damage. Then, use the data valuation metrics discussed above to form a basis for establishing a balanced approach to risk mitigation (see Table 1).

Another facet is the analysis that defines types of threats, and the impact, ease, frequency, and probability of exploitation. Current threat analysis models are far different from those generated years ago because today most corporate data is accessible to

Threats	Prevent	Lessen Impact	Recover
Theft	Security Guard	Locked Cabinets	Police/Lawyers
Network Intrusion	Best Practices	Separate Network	Policies
Availability	Clustering	Mirroring	Tape Libraries
Hardware Failure	High Availability	RAID	Hot Swap
Sabotage	Police and Monitor	Firemen	Insurance
Power Failure	Facility Location	UPS	Electrical Generator

Table 1: Examples of Threats and Potential Techniques to Mitigate Risks

Nature of Threat	Power Outage	"Hacker"
Impact	Data loss, work stoppage	Defaces web site, lost customers
Ease	Low	High
Frequency	Every 100 days	Every 100 seconds
Probability	Y%	YY%

Table 2: Example Threat Attributes

more users than ever before. This broader access has introduced layers of complexity in the user population. Years ago, a bank only had to worry about protecting their assets in relation to the few employees with hard-wired terminal connections back to the mainframe on their desks. Today, federal institutions, end users, and financial partners and networks all have some form of access or other. Suddenly the 200 ACF-2 accounts on the mainframe somehow need to extend role-based responsibility to millions of incoming users. With more users near the data, without the right access controls in place, exploiting a vulnerability can be very easy to accomplish and to repeat at a high rate of frequency.

When you look at threat attributes, don't just concentrate on the logical. Data storage is just that, storage, so many of the threats that need to be mitigated include physical safety (see Table 2).

With these metrics, companies can step through the risks, both physical and logical, to data that is stored on the network and begin to build procedures to protect that data at an acceptable business level. Some additional questions to ask are:

- Can the data be corrupted either in transit or in storage?
- Can it be stolen for personal gain?
- Who can access the data?
- Is the access logged and archived?
- Is the stored data tamper proof or tamper evident?
- Are there copies of the data?
  – And are they secured to the same level as the 'originals'?
- How is the physical security: electricity backups, fire protection, air conditioning?
- Will any of the data be stored off-site with a third party? (All of the above apply again)

On-Demand for Efficiency

Just knowing how and what data needs to be stored, and putting in the proper controls to protect it, won't guarantee that

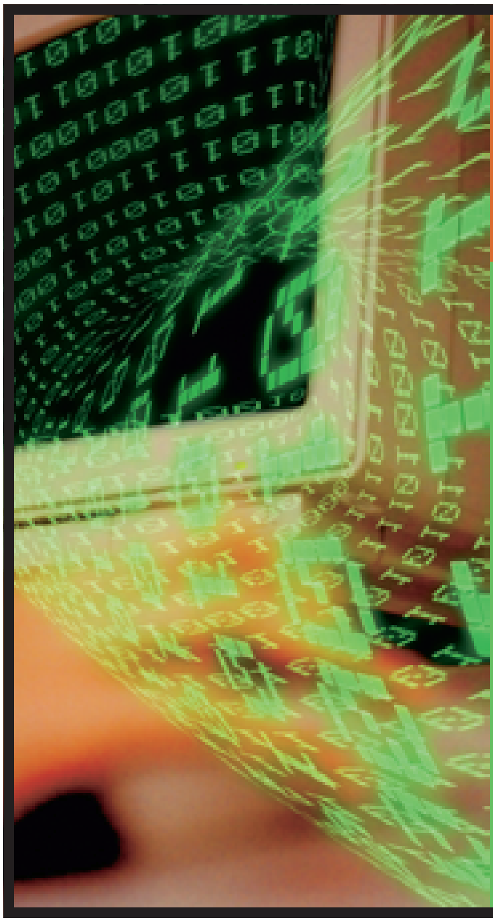

the data will be available when and where it's needed, nor that it will be stored in the most reliable manner. If the data can't be accessed when it's needed, it's not much use. Archived data that has been taken off-site and may take days or weeks to retrieve from storage could be in potential violation of audit policies.

To ensure that data is where it's needed, when it's needed, companies need to look at their own on-demand infrastructure. One of the top priorities is meeting existing and future SLAs (service level agreements) for availability. Another critical point is management of the SAN itself. If more storage space is needed can it be discovered, provisioned, and made available automatically? If not, are the consequences when data is lost or someone gets paged at 3:00 a.m. on a Sunday morning to go into the data center and provision additional storage? Finally, are there metrics in place to predict and plan for storage needs and alert if anomalous storage usage is occurring? Anomalous storage use can be a sign that an attacker is flooding a system and setting off high levels of logging which can quickly fill a server hard drive.

"The default plan for being prepared when audit and forensics investigators come knocking **is to have everything logged and backed up – somewhere, somehow"**

Some additional questions regarding on-demand are:

- Is there sufficient capacity to accommodate growth?
- Is the infrastructure reliable and resilient to attacks such as DoS?
- Do the devices provide high availability and failover?
- Do any mechanisms need to be synchronized for archival purposes?
- Are the devices protected and maintained?
- Are the connections fast enough?
- Are there redundant paths?

Triggers, Reporting, and the Law

While the previous three points address the basics of security and storage management, there are a few additional issues to consider. To maximize storage capacity, some companies may choose to employ triggers that set off higher levels of logging detail when certain events occur. For example, let's say there's a company that doesn't, as a rule, log event information on an end user's machines. However, what if one employee sends an e-mail, flagged by the secure e-mail content monitor, that contains information regarding an upcoming acquisition? In this case, the company might start archiving a full log file trace of this end user's machine to gather data before an SEC investigation occurs. Trigger logging can be quite useful for companies that need to preserve their storage space while tracking legal or audit related data.

And what about reporting? Can the company generate usage and access reports from the stored logs and information on the SAN itself? If best practices are being followed, do the logs reflect this and can the reports prove it out? If someone is accessing backed up data, that shouldn't be; will there be a reported record of when and where and how this access occurred? And if the attempt is thwarted due to strong host-based access control or other measures, will that information show up in the reports?

Finally, a company must ask if any of the log data is impacted by legal requirements. Questions such as how long must the data be retained, and how many backups or copies are necessary, have to be answered. Most of the recent legislation revolves around proving that best practices and controls are in place. It is most often the log files and archived data which show the historical, forensic reality of whether the controls were there or not. In certain cases, critical deleted data on an end user's machine that has been stored, logged, and archived could mean the difference between the user going to jail or the board of directors.

Summary

Management of the business in a continuous and efficient way requires management of storage, securely. Storing data without taking into consideration the security requirements and potential threats is not sufficient in today's enterprise. Legal requirements, audit needs, and shareholder interest all demand that corporations not only protect live data, but log, archive, and store critical, historical data in a safe and retrievable manner. Storage and security are intimately linked, and nowhere is this more apparent than in the realm of archived logged data. No company can afford to be a pack rat with mountains of unsearchable information: keep log data safe and secure by assessing what needs to be stored, mitigating the threats, and keeping the appropriate information available as needed and on demand.

About the Author

Diana Kelley is a security strategist for CA's eTrust brand of security management solutions. She is responsible for evangelizing the eTrust brand portfolio and helps guide CA's security business. Diana has extensive experience creating secure network architectures and business solutions for large corporations and delivering strategic, competitive knowledge to security software vendors. Prior to CA, Diana founded Security Curve, an independent provider of strategy, consulting, and education to the security industry. She also held senior positions with Symantec Corp., Baroudi Bloor, The Hurwitz Group, KPMG and other leading firms and consultancies.

diana.kelley@ca.com

The Dynamic Maginot Line

AN AUTOMATED APPROACH TO SECURING YOUR ENVIRONMENT

BY ERIC VISHRIA

FOLLOWING WORLD WAR I, France's Minister of War and Veteran Affairs, Andre Maginot, convinced the French parliament to build a perimeter line of defense from Switzerland to the Mediterranean to prevent Germany from invading France through its previously exploited Eastern frontier.

Though the Eastern frontier remained partially protected, the strategy ultimately failed when at the dawn of World War II the Germans largely circumvented the Maginot Line by invading through Belgium. Once past the line whose building consumed so many resources for so long, the French army was unable and unprepared to defend against the overwhelming German force and quickly fell.

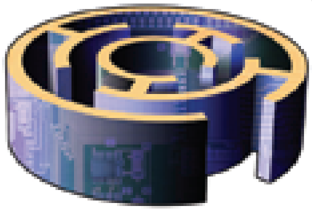

IT organizations have embraced a similar perimeter-based approach to securing their environment, relying on a patchwork of point solutions and ad hoc security schemes that, like the Maginot line, serve to protect the perimeter but leave the foundation of a systems environment insecure. While firewalls and intrusion detection systems provide a good first line of defense, they don't address many of the core vulnerabilities in IT environments.

Enter a new breed of technology known as IT automation software. Automation software complements existing security solutions and can significantly improve the overall security of IT environments by working from the inside out. First, they create a dynamic repository of environmental information that enables quick and accurate vulnerability analysis. Second, they provide a means to execute changes in a systematic and consistent manner based on the information repository.

This two-pronged approach begins by identifying in detail what is present in a given environment. After all, with security it is what you don't know that can hurt the most. Gaining visibility into the environment can be the single most important thing an IT organization can do to secure it. Knowing where servers are located, which applications are deployed, which ones require updating, and even the order in which patches should be applied, all have a considerable impact in securing an environment.

Part of understanding what is present in an environment is understanding how it changes. For example, by keeping track of deviations from an established baseine, automation systems can quickly identify what users – or unwanted intruders such as worms and viruses – have done to specific servers, how configurations have changed, and which machines need to be locked down, and then serve to ensure any backdoors are sealed firmly shut.

Equally as important as identification is the ability to execute changes quickly, efficiently, and using a best practices approach. Many times, seemingly well-guarded environments become vulnerable to attack due to outdated software and understaffed IT departments. Some tools, such as patch management systems, are capable of distributing patches, but operate in an ad-hoc manner completely disjointed from the rest of the server management lifecycle. Consequently, they lack the depth and breadth of coverage to substantially improve the end-to-end security of systems throughout their life.

It is the combination of detailed environmental information and the ability to systematize change that makes automation systems uniquely able to improve the security robustness from the inside out. For example, last year when the Windows RPC vulnerability was announced, organizations with data center automation systems could quickly identify which servers required the RPC port to be open, which servers were unpatched and then actually perform the patching and port shutdown. A point solution only has the inventory of patches and will likely incorrectly patch all systems, shutting down critical applications and resulting in an influx of angry calls to the support desk.

Automation systems on the other hand capture this detailed knowledge in their repositories and can then apply changes precisely where they are needed. This is of particular importance in large heterogeneous environments where different versions of operating systems and endless combinations of hardware and software are running on hundreds or even thousands of servers. Keeping track of everything is only half the battle; ensuring that everything is properly updated in a best practices manner is the other, unwieldy half.

Comprehensive automation systems that take into consideration the people, processes, and technology can turn even the most complex environments into truly impenetrable targets. Automation systems complement your perimeter line of defense systems by reducing the chance for human error, keeping systems up-to-date and ensuring that patches are applied in a timely and uniform manner. Together, these two layers help insulate your environment both against external attacks as well as the unintended consequences of improperly applied patches.

In modern times, the Maginot Line has become a metaphor for something that is relied upon with great confidence but is often ineffectual. Today's IT environment grows in scale and complexity with each passing day. Securing this environment is no small task, but with the right combination of process and automation, the herculean task of proactive security is finally within reach. ■

About the Author

Eric Vishria is the director of product management for Opsware, Inc., a provider of IT automation and utility computing software.

evishria@opsware.com

Identity Theft: More Than A Stolen Wallet

MAINTAINING TRUST IS THE FIRST THING TO REMEMBER

BY JOHN WORRALL

AS THE ROLE of IT administrators continues to expand, it is imperative that companies not lose sight of their core responsibilities: managing and protecting corporate data. This responsibility is becoming increasingly important in the enterprise due to the staggering rise in identity theft around the globe.

A recent report from the Federal Trade Commission (FTC) found that identity theft has achieved the dubious honor of being the most common form of fraud, accounting for 43% of all complaints.

And as more and more corporate and personal information becomes accessible online, that number is increasing. In fact, the FTC reports that identity theft incidents increased 73% from 2001 to 2002.

For a long time, privacy and other forms of e-security have taken a back seat in the enterprise to pressing business issues that consume the attention of both senior management and IT staff alike. It has been common practice to put off thinking about security until the "unthinkable" occurs – a breach. Obviously, that's too late. With this passive approach, companies may be jeopardizing their customers' privacy.

Consider these cases, which have been previously reported in the media:

- The largest identity theft case in history was announced last fall, with total losses estimated at $2.7 million. In this case, investigators arrested a help desk employee of a third-party credit agency who was able to access confidential information about the company's corporate clients.
- A break-in at a health insurance management company resulted in the theft of a file server containing health care information, including some credit card data, from thousands of U.S. military personnel.

Identity theft in and of itself is a broad category, with incidents ranging from petty theft of a single person's identity all the way up to the million-dollar scams described above. But the root cause is the same – the theft of personal information that can be used to obtain credit in another person's name, including bank/credit card numbers, driver's license numbers, social security numbers or even personal information as seemingly harmless as a birthday or mother's maiden name.

But who should take responsibility for protecting people against identity theft? The responsibility has to come from both individuals and organizations holding sensitive data. It's not an either/or situation. For both parties it's largely a matter of awareness. Individuals need to recognize just how easy it is for someone to use their personal information to commit fraud; and organizations need to recognize that it is a privilege to have access to the personal information of employees and customers.

Many organizations don't realize how much sensitive information they carry on their servers and storage devices. Virtually every organization has personal information about its employees that could be used for fraud. Organizations that keep personal information about their customers have an added burden to protect that information. These organizations cut across nearly every industry – from health care organizations to financial institutions to government entities to online consumer sites.

It is important for companies to recognize that identity thieves are less likely to be nameless, faceless hackers than they are to be employees or partners of the company owning the database. This calls for extra time spent ensuring that users of the database have appropriate levels of authentication and access control. Any organization managing identities and customer information is vulnerable to identity theft, and needs to be vigilant about securing that information.

How can organizations prevent/limit identity thefts? First, companies need to determine where the sensitive information exists within their organizations. This is easier said than done because the information could reside on myriad servers and storage systems. You can't protect what you don't know about. Second, companies need to get a true understanding of where and how the information is used to conduct business. Who is it sent to? Under what circumstances is it sent? How is it sent? Who is authorized to access the information in the first place? Where does it come from? Only then can they begin to understand the various points of vulnerability and address them.

Once these first two steps are complete, companies must ensure the systems in place are tamper-proof – making sure information "at rest" is encrypted. This means properly authenticating users (who gets in), monitoring access of the users (where they can go once inside the system), and monitoring the "perimeter" for intrusion attempts. If this is not done properly, identity information can be compromised and the trust of all identities in the system is called into question. A well-managed system for protecting against identity theft includes the following:

1. Properly vetting individuals to assure that the personal information they provide is truly theirs
2. Providing credentials to users accessing the information and providing them with authentication methods to ensure that someone can't access the information using false credentials
3. Implementing the appropriate technologies that allow administrators to access the data they need to effectively perform their jobs, while implementing policies and safeguards that prevent those same administrators from misusing the information
4. Establishing a solid credential-maintenance program – i.e., updating credentials and privileges on a regular basis
5. Quickly revoking credentials and privileges of those who should no longer have access

On the technology front, businesses must move beyond the use of basic passwords for signing onto systems. Technologies that exist today, like two-factor authentication and smart cards, and those that are on the horizon, such as mass-market biometrics, are no longer the exception to the rule. They must become the standard.

No one can diminish the importance of ensuring an employee's computer is up and running, or up-to-date with the latest virus patches. But without working to protect the identities of employees, customers and partners, the loss that could be absorbed by an organization could be immeasurable. If proprietary information is compromised, the trust of the entire organization can be lost, not to mention the loss in actual dollars a security breach could cost a company. ■

About the Author

John Worrall is the vice president of worldwide marketing at RSA Security, a leader in e-security with headquarters in Bedford, Mass..

jworrall@rsasecurity.com

Linux Java Web Services .NET XML Wireless Storage Security
IT solutions GUIDE
TECHNOLOGIES YOU NEED NOW!
How to Manage Your Ideas Using Today's i-Technologies
Top IT Solution Providers!
Candle
ClearNova
Computer Associates
Creative Science
GraphOn
InterSystems
ITKO
Merant
O'Reilly
Parasoft
QuadBase
SAP
Wily Technology
Zero G

Toward Ubiquitous Strong Authentication

THE FOUNDATION OF A TRUSTED NETWORK

BY MARK GRIFFITHS

IT'S ALMOST A tautology these days to say that the Internet has become the life blood for business and personal communications. E-commerce and e-mail are two resounding examples of the transformation exerted by the "network of networks" on people around the globe. Unfortunately, the ubiquity and flexibility of the network has also brought its own set of challenges and security concerns, particularly in the area of user and device authentication.

A strong, ubiquitous authenticated computing environment is needed to address the growing security challenges threatening enterprises today. This article presents a vision for propagating strong authentication across all users, devices, applications, and networks, borrowing from ideas encapsulated in the recently launched Open Authentication reference architecture (OATH) initiative from a wide range of industry players, including hardware and software vendors, token manufacturers, and security companies.

The Need for a Strong Digital Identity

Although recent technology, communication, and geopolitical developments point toward the need for stronger network security, three network trends stand out as driving the imperative for strong digital identities: identity theft, the rise of federated identity networks, and the proliferation of IP devices.

Identity Theft Network Effect

The 2002 Federal Trade Commission (FTC) annual study on consumer complaints cited identity theft, for the third year running, as the most frequent reason individuals contacted consumer protection authorities. While services such as banking, health care, and insurance adopt the network, the fundamental security mechanism for protecting personal information online remains fairly unsophisticated. Since personal information, such as credit card accounts and Social Security numbers, are increasingly used and stored online, an experienced hacker can obtain a dozen passwords from you in a matter of seconds, from anywhere – at any time. A need for strong credentials is important to thwart the "network effect" related to identity fraud. If "something you know" can be stolen through the network, only "something you have" can reduce the threat. A security token in the form of a specialized device or a token integrated within personal digital assistants and mobile phones will be the only viable solution for reducing the threat posed by a global public network.

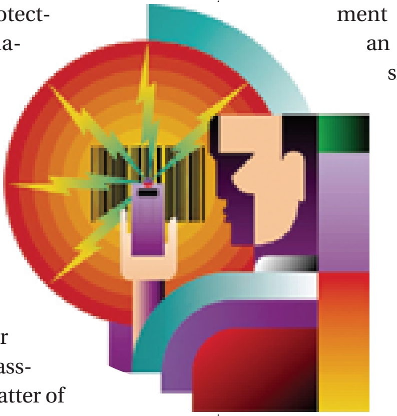

Rise of Federated Identity Networks

While network-based systems are becoming key to the infrastructure that manages corporate content, supply-chain data, and customer services, enterprises are increasingly challenged to provide access to a diverse and dynamic group of end users. The cost and complexity of managing identities across internal and external systems, combined with the necessity of opening up access to data, has created a need for the convergence towards federated identity networks, where identification, credentials, and attributes can be shared among partners. This greatly accelerates the need for stronger identity. If the establishment of technical standards is an important prerequisite for sharing identities, trust is the fundamental business requirement.

To authorize a transaction in a federated identity network, the relying party must be able to trust the credential and identity that was issued and verified by another entity. The strength of this identity must be confirmed and evaluated against the recipient's security policies. When an identity is shared, its strength determines the security that spans the entire access-control chain, creating complex dependencies and liabilities across multiple business and legal parties. The pervasive and interoperable deployment of strong identity technology, security, and operation best practices are therefore key when addressing the crucial issue of trust in federated networks.

Proliferation of IP Devices (Rogue Devices)

Security and trust in any network is a function of all the elements that make up that network. This includes end-point client and server devices that can impersonate users and organizations. As network devices such as mobile phones, PDAs, portable digital music players, set-top boxes, and TPM-based laptops proliferate, the ability to distinguish between trusted and rogue devices is a fundamental security requirement. Since an authenticated device can act as the root of trust, it can also provide the security foundation for a new breed of applications such as identity-based anti-virus solutions and digital information rights management software. From this standpoint, device authentication is a core requirement of any strong identity management strategy.

Realizing the Vision

At the 2004 RSA Conference, a number of industry partners, including chip, smartcard and token manufacturers, operating platform companies, and PKI and VPN vendors announced OATH. These companies realized that for ubiquitous strong authentication to become a reality, corporate employees, Internet users and people accessing everything from health care records to government services, must have the confidence and desire to adopt new technologies such as the tokens described above. To drive this adoption, the technology industry must collaborate to lower the financial barriers and complexity that is associated with strong authentication today. Open technical standards and deployment profiles that promote interoperable solution components are powerful tools for lowering complexity and cost. Therefore, the development of an open and royalty-free specification for strong authentication is the OATH group's initial focus. Open, universal, strong authentication will provide device manufacturers, identity management vendors, security service providers, and application developers with a common framework for the strong authentication of users and devices.

To be effective, a specification must be jointly defined and published by key industry partners that share the vision of universal strong authentication. By laying the groundwork for ubiquity, integration, and interoperability, an open architecture can decrease the risk and complexity of deploying strong authentication products. In turn, the promise of reduced risks and costs will drive adoption across enterprises, service providers, and governments around the world. Ultimately, by making strong authentication part of the network fabric, the entire user community benefits; and by increasing the trust of the network end points, new types of secure interaction will also become possible.

The OATH member companies have laid out a roadmap for the creation of both a strong authentication specification and for the deployment of actual products based on the specification by the end of 2004. If we continue to collaborate, the fastest growing crime – identity theft – could soon become a relic of a bygone era.

For more information on the OATH initiative, please visit www.openauthentication.org

About the Author

Mark Griffiths, VP of Authentication Services, VeriSign Security Services, is a seasoned technology professional with more than 20 years' experience in the computer industry. In addition to his management role in Authentication Services, he also holds the role of VP of marketing, for the division. Prior to joining VeriSign, Griffiths served as the vice president of corporate marketing for VERITAS Software. Reporting directly to the CEO, he led product management, product marketing and corporate marketing during what were arguably some of the most critical years for VERITAS.

Applying Information Lifecycle Management Today

SEPARATING VALUE FROM VISIONS

BY LAURA SHEPARD

WITH VOLUMES OF stored data growing seemingly without limits, organizations are struggling to meet their burgeoning storage demands. While the price of high-performance disk storage continues to drop, it is not dropping fast enough to accommodate the annual doubling of data in more data-intensive environments. The only alternative for many has been manually archiving data from primary disk to tape or other forms of storage – a time-consuming and error-prone process that can inhibit or even prevent access to critical data when it's needed.

Increasingly, Information Lifecycle Management (ILM) is being discussed as the solution to these problems. While much of this concept is based on future developments, a real and significant piece of the functionality proposed by ILM is available today. That piece, referred to as Data Lifecycle Management (DLM), delivers immediate value for data intensive environments.

ILM – the Promise

In theory, under ILM all data is classified and then managed from cradle to grave to ensure that it is automatically stored on cost-appropriate storage devices and given the appropriate level of data protection. In most cases, data goes through a fairly predictable life cycle. It is accessed most heavily in the first few weeks after creation, and then that access frequency drops off significantly as the data ages. Data may eventually be deleted, but an increasing amount of data must be retained indefinitely.

As shown in Figure 1, step 1 of the ILM process is categorization and includes considerations such as criticality of data as well as compliance requirements. In step 2, policies are created to ensure that each category has an appropriate level of access, protection, recoverability, etc. These policies are implemented automatically in step 3. Step 4 is the verification that the system is working and adjustments are made if necessary.

ILM – the Reality

The bad news is that three of the four steps are still manual. The good news is that the DLM solutions available today perform the "automate policies" step, which can save a lot of time and money while helping to manage risk. Policy automation (DLM) solutions keep data available to users and applications while moving it seamlessly among different types of storage without administrative intervention to yield:

- Lower Total Cost of Ownership (TCO) versus buying an all-disk solution to store live data
- Higher productivity versus traditional, off-line
- Lower risk to data availability and integrity versus manual data migration
- Lower data-related liability risk because flexible policies can accommodate a wide range of current and potential compliance requirements

The manual ILM steps and their application are outlined below and illustrated in a customer example that demonstrates the outcome, including a typical representation of the benefits of DLM today.

Categorization establishes the information about the data and can be driven by productivity or non-productivity requirements.

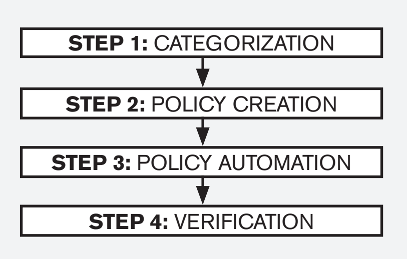

Figure 1: Steps involved in ILM

Productivity requirements dictate that data should remain available as long as it contributes more to the productivity or quality of the work than it costs to keep it available. Accountability elements highlight conditions where a company rule or an outside regulation requires the data to be retained in a certain way, or for a certain length of time.

Categorization Factors

- ***Productivity Elements***
  - Owner
  - Age: When created
  - Size
  - Format
  - Frequency of access and how it changes over time
  - Speed of access and how it changes over time
  - Access Permissions and how they change over time
- ***Accountability Elements***
  - Subject to company policies
  - Subject to compliance or regulatory rules or laws

The goal of policy creation is to ensure that all the factors specified in the categorization process are accommodated in how the data is retained over time, and to take budgetary constraints into consideration. If productivity and accountability elements have been conscientiously determined they should clearly dictate the policy for each data category. Policies need to ensure that frequently used data is on the fastest access media, that no critical data is lost or deleted, and that less frequently accessed data is moved to slower media to save money.

Policy Creation Considerations

- **Persistence:** How long data must be available
- **Location:** On what storage media
- **Access protection:** Degree to which data access is protected
- **Data protection:** Degree to which data is protected from loss
- **TCO/ROI considerations:** Cost of retention vs. the value of the data over time

Policy automation is illustrated in the case study presented below.

Verification is the last step and should be performed at recurring, fixed intervals. Verification consists of checking that the current state of the data fits with the requirements determined in the data categorization and policy creation steps.

Case Study – Widget Co.

Categorization

The design department of Widget Co. designs all of the company's products. A typical design cycle lasts six months and the department needs immediate access to current design-cycle data. To avoid undesirable design elements as well as time-wasting re-invention, they compare against the design data of all products shipped in the past 10 years. These comparisons involve large amounts of data and, because they affect product shipping dates, need to be completed quickly. While there is no current regulatory rule that applies to the retention of this data, the company believes rules are likely to be created in the future. They require that data remains accessible for 25 years to protect against any unforeseen liabilities from either product defect claims, or the introduction of industry-wide compliance rules on data retention.

Policy Creation

Design-related data must be accessible at the fastest possible rates for six months. Design data for the past 10 years needs to be accessed quickly enough to allow for same-day comparison and analysis to occur to protect ship schedules. Because the perceived value to the company of keeping data over 10 years is solely for liability reasons, the only time constraint is that it be retrievable within a reasonable legal discovery period which they determine to be one month.

The requirements for data access and retention indicated by the categorization and policy creation steps confirms that Widget Co. needs an ILM solution. Attempting to address these with an all fibre channel RAID (FC RAID) disk solution would provide the fastest data access and removes the risk, cost, and complexity of manual migration, but would cost almost three times the total current and projected IT budget for storage. Attempting to address cost by using FC RAID only for current design-cycle data while placing older data on a low-cost offline archive fixes the cost problem, but all policy implementation would be manual and the responsibility of the IT department, and historic designs would need to be restored from archive, adding an estimated 20 days to every release.

Sourcing a DLM Solution

Now the IT department has all the core information they need to source a solution for the automate policies step. From the categorization and policy definition steps they know the design department needs the fastest possible access to all current design cycle data and access to data from the last 20 design cycles within a few hours. The legal department needs to be able to access all design data on product releases in the past 25 years in under a month. All design data up to 25.5 years old is considered important, but only current design-cycle data is considered critical.

Taking the policy implementation requirements above, Widget Co. issues an RFI for a solution to meet their requirements. The resulting submissions fall roughly under the same approach: a mix of FC RAID for the current data and less expensive storage media like Serial ATA (SATA) and tape where the policy implementation is automated by DLM software intelligence (see Figure 2). The benefits are that current design data is on the fastest media; and because of the automated DLM implementation design, data under 10.5 years old is stored on less expensive media, but does not have to be restored from archive. Design data over 10.5 years old is automatically identified by the DLM software as ready for archive, human error is removed from the policy implementation, and the system conforms to data access standards. No drawbacks to this approach are identified.

Refining DLM Solution Criteria

Widget Co. requests quotes for DLM solutions. The submissions show the company the importance of several criteria they had not allowed for. They then refine their criteria to verifying that the right DLM solution should:

- Be proven to scale to a capacity that addresses 10-year projected growth
- Be proven to scale to a performance level at which same-day analysis of historic data comparisons can still be performed in one day at data sizes projected for 10 years in the future.
- Have standard interfaces and allow the maximum policy flexibility to accommodate possible future accountability requirements

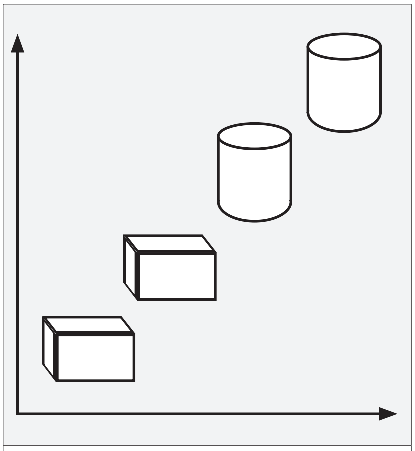

Figure 2: DLM solution

Results

The IT department estimates that over the next 10 years they will save:

- Over 65% on the storage capacity versus all disk since they have planned SATA and Tape as the majority of their planned capacity
- Approximately 50% on storage management costs through the removal of manual monitoring and provisioning of available capacity and data movement

The design department estimates that over the next 10 years they will save almost $1 million in personnel costs by eliminating the 20 days per release accessing data from the offline archive would add to comparisons with historic product designs. An additional benefit to the company of this approach is that it reduces the design cycle by 20%, allowing them to become 20% more productive.

Widget Co. is satisfied that the system will allow them to find any data for which they could reasonably have any liability and to demonstrate its integrity, and because it works with their backup system and fits the IT budget, no additional analysis is considered necessary to approve the implementation.

They decide to perform a verification of the system function every six months, adjusting policies to meet any new requirements or change any that are not achieving the desired results.

Conclusion

As the Widget Co. example illustrates, while ILM solutions do not currently deliver on all promised areas, existing approaches, or DLM solutions, do offer significant value for environments seeking to reduce costs, increase productivity, and meet specific retention requirements.

About the Author

Laura Shepard is marketing product line manager for SGI InfiniteStorage (www.sgi.com/storage). Having worked with data-intensive computing and storage environments for seven years, Laura has seen data solutions grow from gigabytes to petabytes.

shepardl@sgi.com

Before Signing on the Dotted Line...

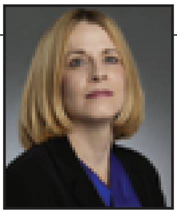

EVALUATE FOR SECURITY

BY MARY ANN DAVIDSON

A RECENT REPORT FROM PricewaterhouseCoopers confirmed that most security breaches occur in stored data. Exponential growth in storage capacity, coupled with emerging regulatory requirements, has led to an even greater increase in storage network vulnerabilities. Today, organizations are forced to recognize the critical importance of securing all types of data – from corporate confidential documents to enterprise instant messages to global personnel records. To meet these challenges, organizations must deploy a smarter, more cost-effective approach to security and veer from the prevalent method of developing and implementing patches only after problems are discovered. This article outlines a step-by-step process for organizations to use as they evaluate technology products.

Internal Perspective

Organizations can make better decisions about security products and reduce the potential back-end costs by researching a few key vendor practices; examining the vendor's corporate culture, specifically the security of their development process; insisting on a response plan for times when vulnerabilities are found; and demanding third-party assessments.

When researching information-technology products, organizations must investigate the vendor's security practices and determine the true cost of the product. A product's true cost is often not just the licensing costs, but also the time and money invested in patching a product once a vulnerability is discovered. Organizations need to make educated purchasing decisions rather than dedicate resources to applying patches after procuring a product, a process that can prove more costly in the long run. An educated purchase can prove less costly down the road. For example, the estimated cost to deploy a patch for a recognized software flaw runs on average $900 per server and $700 per client. If an organization misses a patch and gets hit by a virus, the cost will be magnified.

Vendors must demonstrate that security is a priority at each step of product development and delivery. Some software vendors provide training in secure coding practice and compensation tied to secure coding objectives, thereby strengthening the company's security culture. Organizations with a chief security officer and a team that analyzes product development for weaknesses, or hacks its own products, are clearly dedicated to security. It is better that the vendor notice product weaknesses before the flaw causes problems. The vendor should also run its own enterprise on its products; if a company doesn't trust its own products to secure secrets, why should you?

Patch Management

Before signing on the dotted line, organizations should be convinced of two important elements: the vendor has an aggressive plan to handle problems that may arise; and the vendor has a strictly adhered-to incident-response policy to determine the severity level of a vulnerability. These two elements help to mitigate security vulnerabilities should they arise.

Subsequently, the vendor should finish all relevant patches before announcing a security alert. Information distributed randomly to a handful of customers will exasperate rather than calm the situation. Further, the vendor's security policy should treat all customers equally by providing the same level of notice to all customers, regardless of their size or industry.

Validating Security Claims

Third-party validation represents a critical step in purchasing secure products. Vendors that are serious about security will submit their products for rigorous security evaluations conducted by independent authorities. These evaluations are recognized globally by various governing bodies and provide organizations with a level of assurance about the product's features and security claims. Sometimes, evaluators find product weaknesses and vulnerabilities that are corrected before the evaluation is completed or the product is released.

These evaluations are not without a price. However, reputable vendors know that remedying vulnerabilities found during an evaluation is cheaper than fixing a product already in use. For example, while the cost of an evaluation can reach $1 million, the cost to create and issue a patch for multiple versions of a product that is available on 20 different operating systems can easily cost that much, not including the cost of patch application. Clearly, creating secure products is in the best interest of the vendor and buyer.

Although this due diligence adds a step to the product procurement process, it raises the bar for security across the board. If the industry fails to follow these guidelines, it risks government agencies regulating the process. The U.S. government has already instituted compliance regulations such as Sarbanes-Oxley and the Health Insurance Portability Accountability Act (HIPAA) to govern the way the financial and healthcare industries guard their stored data.

Security as a De-facto Purchasing Criteria

"IT" now stands for "infrastructure technology," and needs to be as robust, secure, and reliable as physical infrastructure. We never worry about bridges failing, nor should we worry about some of our most critical IT systems – such as SANs, NAS, DAS, and backup environments – going down because of design defects.

Adhering to these security guidelines and choosing more robust products are prudent moves that will cut costs and improve business in the short and long term. ■

About the Author

Mary Ann Davidson is the chief security officer at Oracle Corp., and is responsible for security evaluations, assessments, and incident handling. She represents Oracle on the Board of Directors of the Information Technology Information Security Analysis Center (IT-ISAC), and is on the editorial review board of the Secure Business Quarterly. Mary Ann has a B.S.M.E. from the University of Virginia and an M.B.A. from the Wharton School of the University of Pennsylvania. She has also served as a commissioned officer in the U.S. Navy Civil Engineer Corps, where she was awarded the Navy Achievement Medal.

mary.ann.davidson@oracle.com

Someone set off your firewall?

10,000
Complicated regulations to navigate
1
Architecture to simplify the journey