SYS-CON

The Largest i-Technology Event of The Year! 〈30

# information STORAGE+SECURITY journal

www.ISSJournal.com

## In This Issue:

PREMIERE 2005 VOL.1 ISSUE 1

## Compliance Essentials: 〈4 Fullfilling the Requirements

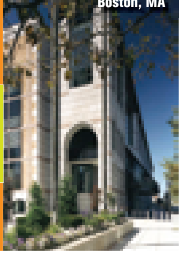

$5.99US $6.99CAN

12

0 71486 01793 6

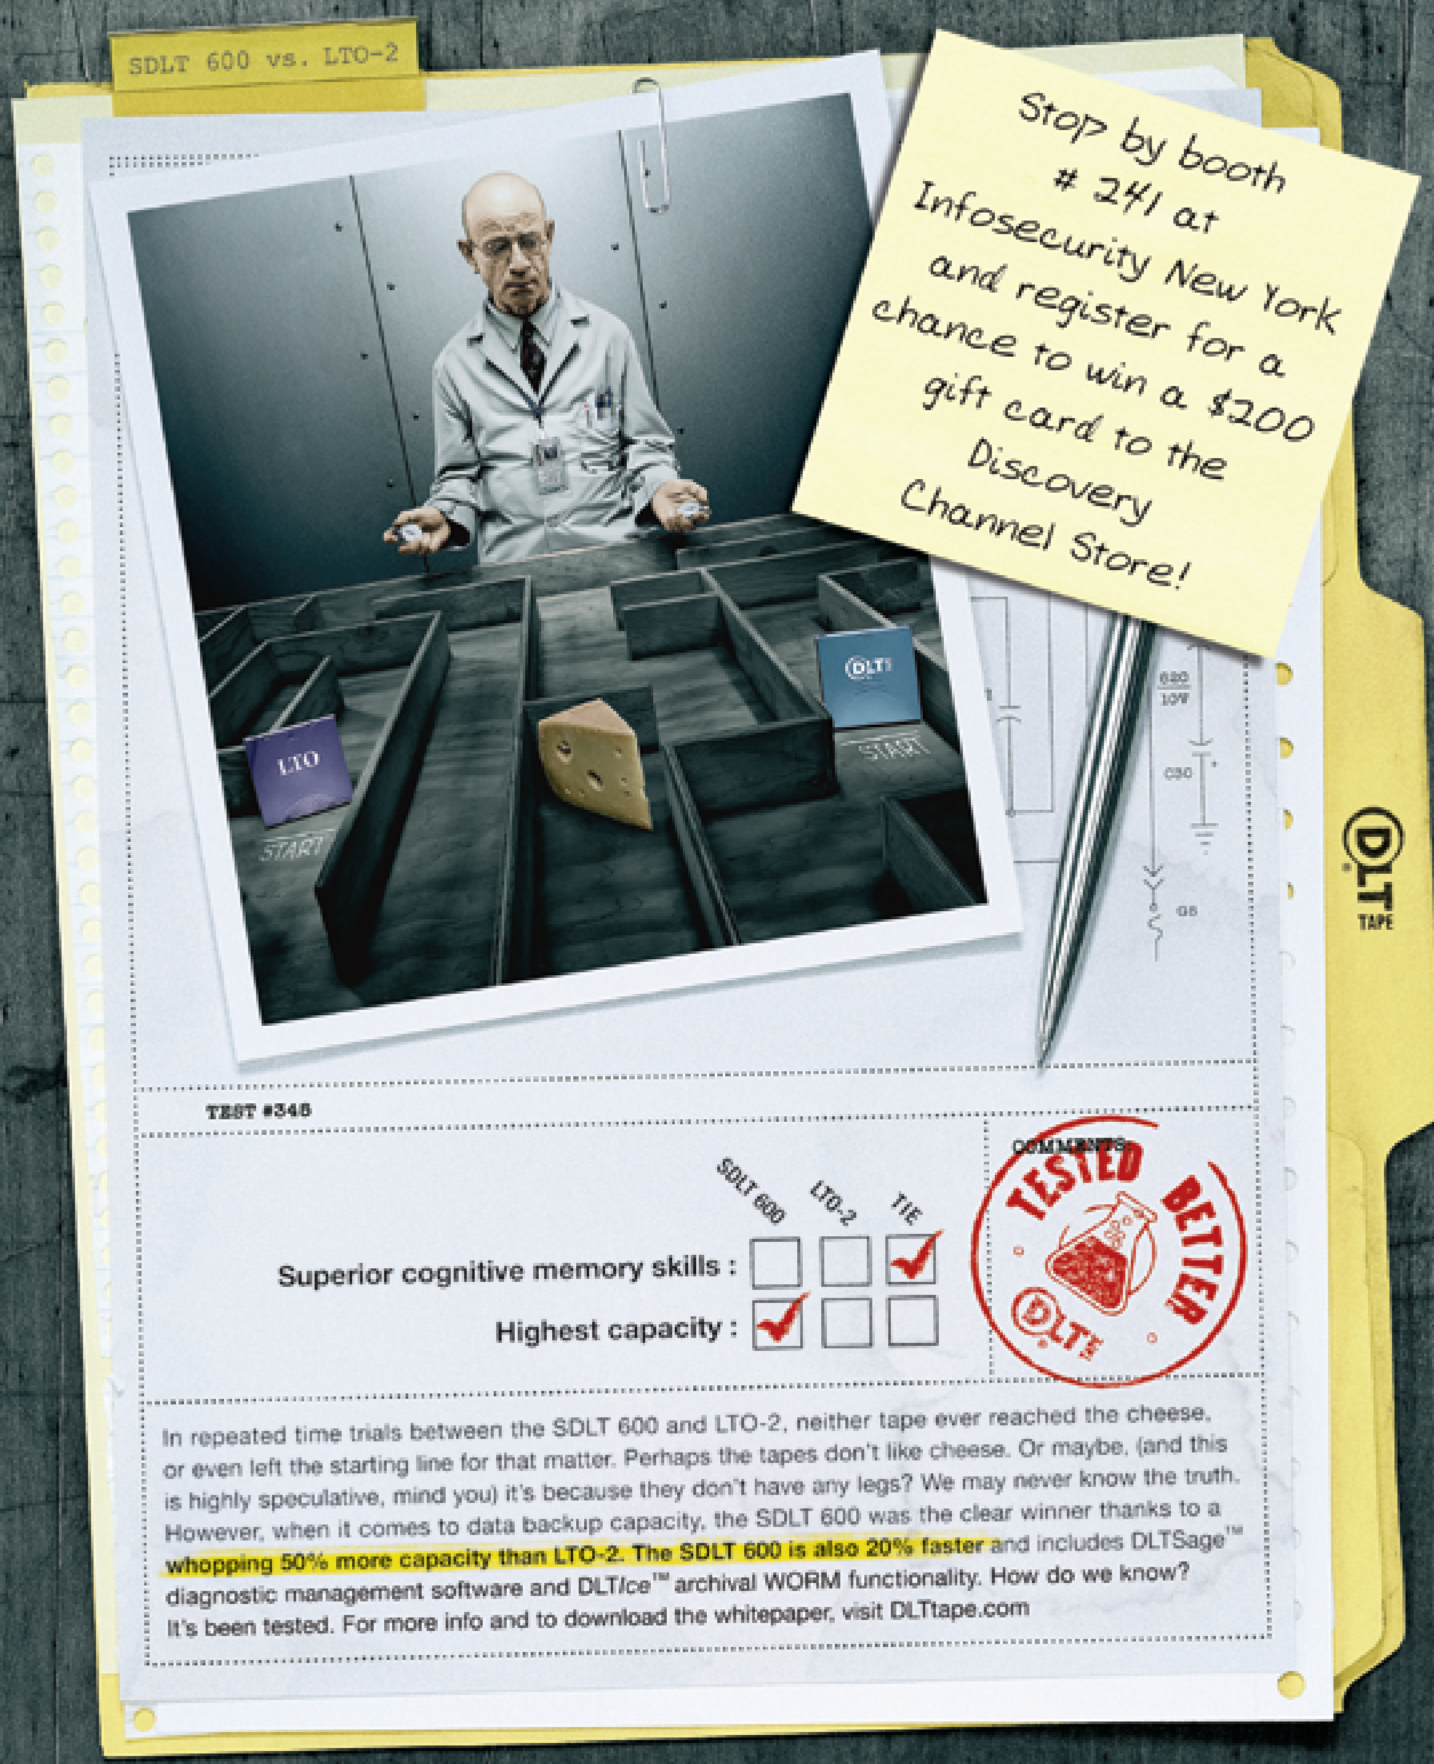
SDLT 600 vs. LTO-2
Stop by booth # 241 at Infosecurity New York and register for a chance to win a $200 gift card to the Discovery Channel Store!
LTO
START
DLT
START
DLT TAPE
TEST #348
SDLT 600
LTO-2
TIE
Superior cognitive memory skills : □ □ ☑
Highest capacity : ☑ □ □
COMMENTS
TESTED BETTER
DLT

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

**Advertising Sales Manager**
Dennis Leavey dennis@sys-con.com

**Associate Sales Managers**
Kristin Kuhnle kristin@sys-con.com
Dorothy Gil dorothy@sys-con.com
Kim Hughes kim@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
Natalie Charters natalie@sys-con.com

**Online Editor**
Martin Wezdecki martin@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
Andrea Boden andrea@sys-con.com

**Web Services**

**Information Systems Consultant**
Robert Diamond robert@sys-con.com

**Web Designers**
Stephen Kilmurray stephen@sys-con.com
Matthew Pollotta matthew@sys-con.com

**Accounting**

**Financial Analyst**
Joan LaRose joan@sys-con.com

**Accounts Receivable**
Stephen Michelin smichelin@sys-con.com

**Accounts Payable**
Betty White betty@sys-con.com

**Customer Relations**

**Circulation Service Coordinators**
Edna Earle Russell edna@sys-con.com
Linda Lipton linda@sys-con.com
Monique Floyd monique@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

All brand and product names used on these pages are trade names, service marks or trademarks of their respective companies.

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

From the Group Publisher

# An A-Z of Security and Storage

## *IT PROFESSIONALS, TECHNOLOGISTS, AND BUSINESS LEADERS REWRITE THE LEXICON*

BY JEREMY GEELAN

SPARE A THOUGHT FOR THE COMPILERS OF DICTIONARIES IN THE DIGITAL AGE. Technology is always moving beyond the confines of the alphabet.

If you were given only 26 choices, for example, what would you list as the chief concerns of IT professionals today? In the storage space alone, there have been more product announcements from suppliers of storage systems in the past six months than in the previous two years. And in the security space, not a week – sometimes not even a day – goes by without a new offering..

So, what should today's i-technology abecedary look like? A for Authentication, B for Backup, C for Clustering, D for Denial-of-Service, E for Encryption…

How about A for AIT (Advanced Intelligent Tape) or D for DAS (Direct Attached Storage)? And what about B for Bots, which are siphoning and transmitting sensitive information from compromised PCs, receiving and spreading malware updates, and being used in distributed, denial-of-service attacks on a wider scale than ever before.

Should F be for Firewall or Fibre Channel, H for Host-Based Security or HIPAA?

By the time you get to S you'd literally have to abandon all hope of narrowing the choices: SAN, Sarbanes-Oxley, SNIA (Storage Networking Industry Association), SNMP, Spam, SSL…Why, with just 26 choices you'd probably never even reach U for USB Drives, V for Virtualization, or W for Worms. Let alone Z for Zero-Day Attacks.

Then would come the colloquies like "Disaster Recovery," "Utility Storage," "IP Spoofing," and the like. Never mind SAN/NAS/RAID, less familiar acronyms are arriving thick and fast, like DHS (Department of Homeland Security), SEP (Security Experts Panel) and even new institutions – like the Internet Storm Center (ISC), an all-volunteer early warning Internet global monitoring organization (http://isc.sans.org/).

Often, amid this slew of technologies and innovations, each new approach seemingly spawns a secondary headache – such as the trend towards networked. IP SANs, which many see as likely to unleash security problems since those who would seek to do harm are so familiar with the IP protocol.

Some say that, in the great scheme of things, neither storage nor security is a front-burner issue – business is. Certainly it is true that, as a recent report noted, IT professionals are often embroiled in operational and tactical considerations, with little time or resources left over for a more strategic approach, and so an understanding of where the storage-security nexus fits in the overall business puzzle is important. But the devil is in the detail, and detail is what we will bring you in each issue.

Here at *ISSJ* we will cover what's new, what's best, and what's next in the ever more important nexus of security and storage. We'll look at key issues, such as whether open-source software means better security or worse.We'll ask where information lifecycle management is going; we'll explore every aspect of storage networking; we'll drill down into NAS management and object-based storage.

What's needed, *ISSJ* articles will show, is a careful, business-based balance between security and storage. Even the most sophisticated SAN isn't much use if it isn't secure – audit regulations require that companies not only log and archive critical data, but also that they do this securely.

As Lenny Heymann, general manager of NetWorld+Interop said, when we unveiled our preview issue at the Networld+Interop Conference & Expo in Las Vegas: "Today's IT buyer is taking a very pragmatic approach to networking purchasing decisions, and really scrutinizing the full range of implications those technologies might have for their company – so discussions about storage should absolutely include related security issues."

The security-storage nexus is here to stay. So is *Information Storage & Security Journal.* ■

### 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*

# Compliance Essentials: Standard Methods of Fulfilling Requirements

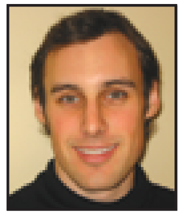

IMPROVE INFORMATION SECURITY AND COMPLY WITH REGULATORY REQUIREMENTS

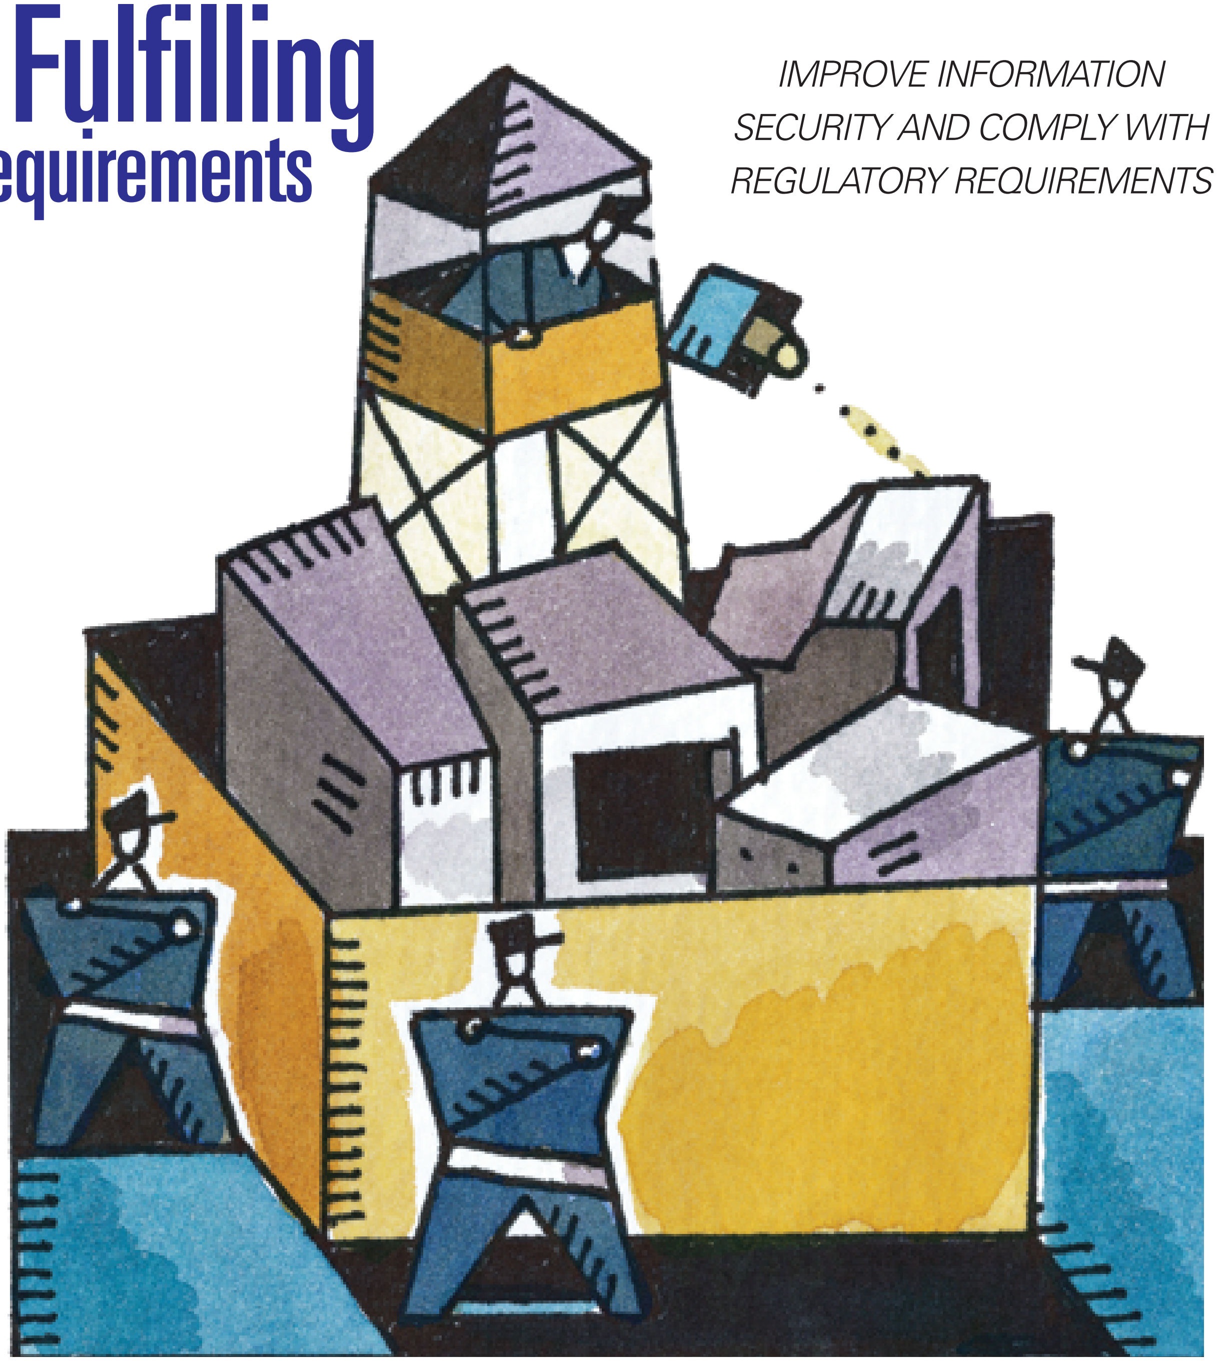

BY RYAN KALEMBER

FROM THE HEALTH CARE INDUSTRY to the financial industry, the influx of network security incidents has impacted any organization that employs the Internet to expedite business processes. As a result, anyone enlisting the services of these companies is susceptible to identity theft or fraud. Responding to this issue, the U.S. government has amplified its legislation dealing with infrastructure security through bills including Sarbanes-Oxley (SOX), the Health Insurance Portability and Accountability Act (HIPAA), the U.S. Government Information Security Reform Act (GISRA), the Gramm-Leach Bliley Act (GLBA), and the Children's Internet Protection Act (CIPA).

These laws require organizations in their respective industries to ensure the safe transfer and storage of personal information. Through strict enforcement of compliance regulations, including tough penalties for violators, the government has dramatically influenced how companies contend with network security issues. In this article, readers will learn the requirements and legal ramifications for each act and gain practical and strategic guidance for achieving compliance.

## Introduction

A reliable indicator of when a particular practice has reached some degree of maturity, or at least adolescence, is the moment when the federal government begins to regulate it. Perhaps an even greater degree of accuracy for discerning that point is when regulations are enforced. An illustrative example is antitrust legislation, which began in 1890 with the Sherman Antitrust Act, but was not enforceable until the passing of the Clayton Act in 1914. Judging by these criteria and allowing for the slightly speedier movement of the U.S. government in the Internet age, information security is on the cusp of its maturity. A variety of pieces of legislation have reached, or will soon be reaching, their compliance deadlines.

## The Legislation

After the headier days of the late 1990s, the federal government took steps to curb irregularities and risks with a series of regulations aimed at particular industries or practices. Public companies with a market capitalization of more than $75 million are perhaps most affected by the SOX Act. This act, among other things, requires checks on the integrity of information involved in the business processes that feed into the enterprise's balance sheet. Certain SOX compliance deadlines have already passed, whereas others are due this year and next.

Two notable regulations are already in full effect. In the health care sector, HIPAA requires a variety of measures designed to safeguard the privacy of patients while facilitating the move to electronically stored (i.e., "portable") medical records.

The GLBA has provisions already in effect that specify how financial institutions can use and share their clients' financial information with other organizations.

The U.S. federal government has not left itself out. The GISRA, which expired in 2002, has had many of its provisions made permanent in the Federal Information Security Management Act (FISMA). Since the Bush administration ordered that funding for IT projects be tied to security compliance, FISMA has become even more critical for both federal agencies and the vendors who sell to them. Important elements of FISMA include the following:

> The National Institute of Standards and Technology (NIST), collaborating with federal agencies, develops mandatory IT security standards and guidelines for nonclassified federal IT systems.

> Agencies develop system configuration requirements and provide ongoing monitoring and maintenance.

> Agencies test security controls at least annually.

> Agency CIOs designate a senior agency information security officer to ensure FISMA compliance.

> Agencies provide an inventory of their IT assets.

Other regulations are tangentially related to the "big four" noted previously. The CIPA is a federal law requiring libraries and schools to take measures to block minors' Internet access to obscene materials, inappropriate e-mail, adult chat rooms, or "hacking."

California has passed a law, known variously as Senate Bill (SB) 1386 or the California Database Protection Act. This requires companies doing business with customers in California to notify them if they suspect that any of their customers' personal information has been accessed by an unauthorized party. Similar legislation has been proposed in the U.S. Congress, although it has not been passed yet.

Finally, the private sector has joined in, with Visa and MasterCard regulating both their merchants and service providers. Visa's initiative is called the Cardholder Information Security Program (CISP) and MasterCard's is called the Site Data Protection (SDP) program. Both programs require that all merchants and service providers are assessed for key information security best practices and, depending on the size of the merchant, evaluate systems involved in the handling or processing of cardholder information for security vulnerabilities.

## Key Trends

Although it may seem that the factors driving the passage of these laws are obvious, it is worth specifying which elements within the broad categories of information security and privacy are tied to each specific piece of legislation. Apart from self-evident issues, the regulations address concerns about the security of personally identifiable information (PII) or accountability for IT systems that process sensitive material. This is in addition to monitoring and maintaining them.

HIPAA, GLBA, and California SB 1386 can be placed in the former category. The prevalence of identity theft has called attention to the security of databases of financial or other personal information maintained by a variety of institutions. This is particularly true when those databases are either accessible from the Internet or, as is more common, are connected to systems (e.g., Web servers) that are.

The second important trend that is driving tighter and more detailed regulations is accountability for IT systems and the processes that rely on them. The past decade saw an IT expansion the likes of which may never be seen again. In addition, the sheer quantities of IT equipment that were purchased provided a serious challenge for organizations seeking to track their assets. Once critical data and processes began to be stored or executed on these assets, the seeds were sown for both information security vulnerabilities and the concomitant legislation.

This has led to specific provisions in several of the regulations described previously. In the case of SOX, public companies' chief financial officers and chief executives become personally responsible for the tabulated results of electronic business processes. This made the integrity and security of the systems that enable those processes critical in ways that they were not before. Human auditors can no longer provide adequate supervision of certain business processes due to the volume of information. This, automated audit mechanisms, highly specific to the related business process, are being developed to provide the oversight required by law.

For FISMA, in addition to the aforementioned required IT asset inventories, the certification and accreditation of systems also feeds into a report card issued by a House subcommittee. FISMA, however, is more directly a response to the 1.4 million documented cyber security incidents involving federal agencies in 2003. This is a statistic from the Federal Computer Incident Response Center.

## Basic Compliance Strategies

In general, the following measures will address the basic compliance requirements for information security regulations:

> Full inventory of IT systems involved in the processing, storage, or transmission of sensitive data
> Information security policy and a corresponding awareness and training program
> Privacy policy
> Computer security incident response plan

### Basic HIPAA Assessment Elements

> Administrative security
> Policies
> Procedures
> Physical security
> Technical security
> Privacy
> Coding practices

### HIPAA Technologies

In addition to process solutions, a wide variety of technologies can aid in a HIPAA compliance effort:

> Firewalls
> VPNs
> Auditing tools
> Password policy enforcement tools
> Intrusion detection tools
> Encryption tools
> PKI
> Digital signatures
> Authentication technologies
> Other access control devices

Beyond these elementary steps, organizations must determine to which regulations they are subject. Although this may seem entirely obvious (i.e., federal agencies are responsible for complying with FISMA, and public companies must adhere to the requirements of SOX), the applicability of some of the regulations discussed in this article is slightly trickier to determine.

For example, any organization that does business transactions with California customers and stores their data on an IT system is subject to SB 1386, even if that organization is not located in California. In addition, companies that have their own health or dental plans and store employee medical information may be subject to certain provisions of HIPAA. Finally, companies that do not consider themselves financial institutions may need to be compliant with GLBA if they collect, store, and share financial information about their customers with their business partners.

Before moving on to specific legislation, it is critical to define the terms "security" and "privacy," as they are employed here. In the information security world, it is often said that it is possible to secure information without making it private. However, it is not possible to keep information private without securing it. Information security is generally defined as the ability to control access to information and protect it from accidental or intentional disclosure to unauthorized persons and from alteration, destruction, or loss. Privacy is controlling who is authorized to access the secured information or the right of individuals to keep information about themselves from being disclosed, depending on the context.

## Sarbanes-Oxley 101

The bulk of current compliance efforts at U.S. corporations are likely directed toward SOX, which became U.S. law in July 2002 and section by section has become effective. A major deadline passed as recently as June 15, 2004, when Section 404 became effective. Section 404 is perhaps the most relevant to information security, as it refers to management assessment of internal controls for financial processes.

In tactical terms, this means that financial reporting systems must have controls that follow internationally recognized auditing frameworks, such as the one provided by the Commission of Sponsoring Organizations of the Treadway Commission Internal Control (COSO). Specific to IT and information security, standards such as Control Objectives for Information and Related Technology (CObIT) and ISO 17799 have been recommended for compliance by the SEC in clarification rulings. It is critical to note that "financial reporting systems" refers to more than simply spreadsheets and databases, and includes informal reporting channels such as e-mail. Reporting systems can potentially include policies, plans, processes, systems, and procedures of all manners at every level of the organization.

Although other types of process development may constitute the majority of the work in a typical SOX compliance effort, information security concerns must pervade any successful effort. Section 404 requires the implementation of controls that protect and monitor the integrity of financial reporting processes. It also requires reporting on the efficacy of those controls. In addition, Sections 409 and 802 have serious integrity-related implications for material changes (to the company's financial conditions) and audit records, respectively.

From an IT perspective, SOX compliance can present a confusing situation at best. Many CIOs have viewed SOX as an audit or financial issue, although this interpretation has proven incorrect. The primary goal of SOX is to ensure the integrity of financial reporting systems. Nearly all of these use IT and therefore must be in the scope of any successful compliance project.

IT compliance efforts have generally taken a five-step approach for each relevant system:

> Determine how the system will be operated and configured once it is in compliance, including processes and controls.
> Assess the current state of the system, performing a gap analysis relative to the compliant state.

PROACTIVE INTELLIGENCE
24/7 MONITORING
TOTAL NETWORK VISIBILITY
REAL-TIME CORRELATION
VeriSign® Managed Security Services
Where visibility and intelligence overpower fear and doubt.

- > Implement any process improvements or new controls, and remedy any identified vulnerabilities.
- > Monitor the system to ensure that it is in line with the compliance requirements (i.e., with vulnerability scanning, intrusion detection, or log monitoring).
- > Report on the compliance status in a format that is intelligible to the audit staff or other management.

## Information Security and HIPAA

Compliance with HIPAA, which most large health care providers should have achieved already, is a complex proposition. For the vast majority of enterprises not in the health care sector, HIPAA will only be relevant to any medical information stored about employees or their spouses on the enterprise's IT systems (see sidebars).

The first step an organization should take is to identify and review all policies relating to physical or electronic access to the relevant data (i.e., medical records) and the protection of that data. The next step in the information-gathering phase is to create questionnaires that address all aspects of data storage, transmission, protection, confidentiality, and privacy for the relevant data.

The second step of a compliance effort is generally a gap analysis, which compares the current state of data security and privacy with "best practices." HIPAA itself has no clearly defined, technology-related or risk-related standards, so a due diligence approach based on best practices is required.

The third step of the plan is generally a "compliance roadmap," which describes how the organization plans to close critical gaps in security and privacy. The actions should be categorized as technology implementations, policy changes, or auditing procedures.

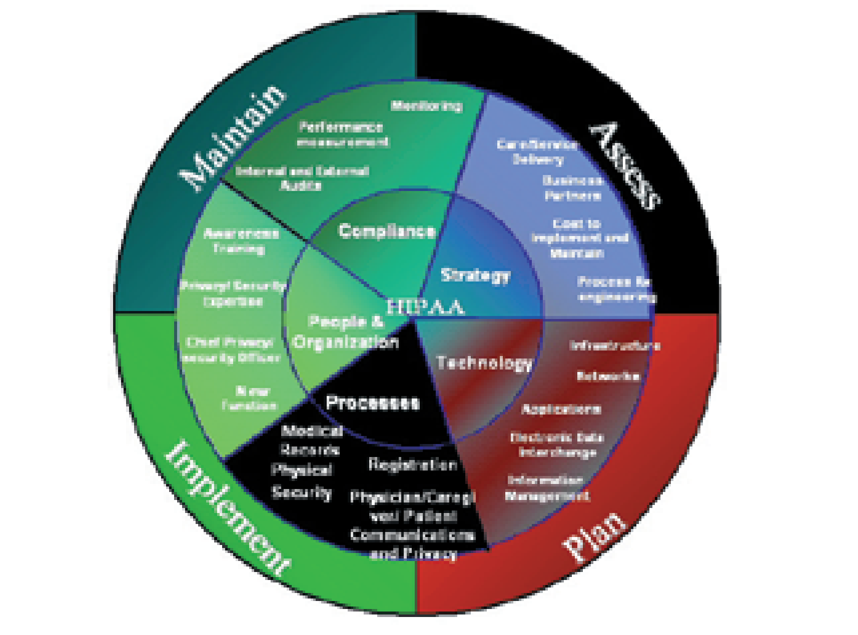

**Figure 1:** HIPAA best practices

| Classify Data | Unique User IDs | Restrict Access | Authenticate Users |
|---|---|---|---|
| Authenticate Customers | Log Access | Password Reset | Provide Logout |
| | | | |
| Store Encrypted | Encrypt Cookies | Use Strong Encryption | No URL Leaking |
| No Caching | Secure Purging | SSL for Transmission | Display Restrictions |
| Use Anaconda | Log All Access | Use Corp Directory | Message Digest (MD) for Password |
| No Secret Display | Account Lockout | Encrypt All in Transit | MD Secret, Encrypt All |
| Opt-in for users | No Caching of Confidential Info | Purge all Confidential | Minimize Display of Sensitive Info |

**Table 1:** GLBA Checklist

This remediation planning should also encompass how the organization will maintain compliance, which could include any or all of the following:

- > Auditing
- > Intrusion detection
- > Enterprise security management
- > Privacy "opt-in/opt-out"
- > Monitoring plan

In any HIPAA assessment, it is critical to note that health care organizations are affected by both HIPAA and state laws, and that privacy regulations such as HIPAA do not preempt state law or other federal law. Any state law or regulation that is contrary or more stringent than the corresponding HIPAA rule retains primacy.

HIPAA has no proscribed implementation measures for either its security or privacy rules, so implementations will vary according to the type and size of the covered organization. Just as with the other regulations mentioned in this article, best practices need to be implemented and followed to achieve compliance (see Figure 1).

## Options for GLBA

Most companies should have been in compliance with GLBA when the deadlines passed in July of 2001 and in May of 2003. However, newer companies or those just starting to electronically store personal information about their clients may still need to take steps to comply. Similar to many of the other regulations, compliance with GLBA can be achieved through information security best practices in general and a few privacy initiatives specifically.

The specific compliance issues brought up by GLBA pertain largely to handling customer information collected via the Web or other sources and the sharing of that information. Basic security measures for Web sites that collect information from customers should be applied, including SSL encryption for transmission, cookie encryption, and account lockouts. In addition, GLBA specifies that customers must be asked to explicitly "opt-in" if the enterprise is to be able to share customer information with other institutions (see Table 1).

## Other Regulations

FISMA compliance efforts have largely centered on key metrics, such as the percentage of IT systems that have been certified and accredited or the percentage of significant new IT investments that integrate security into their lifecycles. Other goals are process-related, requiring each agency, for example, to have a centralized set of procedures to identify, track, and correct security vulnerabilities. To coordinate these processes, many agencies have hired full-time chief information security officers.

CIPA compliance is a significantly easier proposition. Most commercial content filtering software meets the requirements of the legislation, as it has become an essential selling point. Certain configuration changes to PCs in libraries and schools are also helpful, such as disabling administrative access and certain services.

In the private sector, California SB 1386 has simpler requirements. The key step for any enterprise with California clients is to develop and document an incident response plan specifying notification procedures. If such a plan is in place, the organization may follow its own process, rather than the onerous procedures prescribed by the law itself. Protecting systems that store, process, or transmit personal information about organizations' clients is sound business practice in any case, and is the only other general rule to comply with this California statute.

> "Privacy is controlling who is authorized to **access the secured information"**

## Enforcement

Much to the relief of many organizations and executives, the stiff penalties mentioned in much of the legislation have not yet been applied systematically to violators. This relief may be short-lived, as both SOX and HIPAA hold out the threat of prison time for executives who sign off on financial results of questionable provenance. The SEC, which enforces SOX, will likely not pursue vigorous enforcement until it finishes with the Enron and WorldCom cases (and the other corporations), who inspired the passage of SOX in the first place. Once each provision of the law has come into effect and pending clarification decisions are rendered, the SEC should enforce the law pursuant to those decisions.

By contrast, the FTC, which enforces GLBA, has already fined companies for violating the privacy of their customers. The most famous example of this was Eli Lilly, which mistakenly did not obfuscate the e-mail addresses of Prozac patients on a targeted bulk e-mail.

Another federal agency, Health and Human Services (HHS), enforces HIPAA. Security provisions are enforced via the Centers for Medicare and Medicaid Services (CMS), and the HHS Office for Civil Rights enforces the privacy component of the act. CMS is currently assembling an enforcement staff, writing a regulation that outlines the enforcement program, implementing the enforcement system, and beginning to accept complaints. According to CMS, it intends to "provide education and technical assistance to covered entities to help them achieve compliance, rather than seeking out noncompliant entities and imposing fines on them." If a covered organization is identified as noncompliant, CMS plans to work with it to achieve compliance and would only impose civil monetary penalties if these efforts fail.

FISMA is enforced by a combination of government entities. The Office of Management and Budget develops the Federal Computer Security Report Card for each agency using agencies' quarterly-updated plans of action and milestones (POA&M) and IT security performance metrics. Inspectors general and the General Services Administration (GSA) also play a role. Additionally, IT security is a crucial component of a "green" rating on the President's Management Agenda's quarterly E-Government Scorecard.

CIPA is enforced by the FCC, which withholds the discounts offered by the "E-Rate" program to schools and libraries that do not certify their compliance.

The enforcement of other legislation, such as California SB 1386, is more of a question mark. In theory, a corporation subject to a hacking incident would be in violation of the law if it (a) had California customers and (b) could not prove that the database containing the customers' information was not inappropriately accessed.

It remains to be seen how this law will be enforced in practice.

## Conclusion

Although it may be quite easy to become frustrated by the alphabet soup of recent information security regulations, everyone from executives to IT personnel can take solace in the fact that few of the regulations specify any practice that is not already part of the information security canon. "Best practices" is an overused term in the private sector but is nearly ubiquitous in these regulations. By taking a sensible, standards-based approach, organizations can both improve their information security and comply with the vast majority of regulatory requirements. After that, the targeted compliance measures for what the regulations that the organization is covered by become much more manageable. ■

## Additional Resources

- > *Verisign:* www.verisign.com/
- > *American Library Association: Child Internet Protection Act:* www.ala.org/cipa/
- > *United States Computer Emergency Readiness Team:* www.us-cert.gov/federal/

**About the Author**

*With over nine years of industry experience, Ryan Kalember boasts a formidable background in information security. As a senior consultant at VeriSign, he has designed secure e-business architectures, configured and deployed authentication systems, consulted on PKI implementations, conducted enterprise security assessments and secured 802.11b wireless networks. He has also developed guidelines for compliance with ISO17799, the European Union Privacy Directive, the Gramm-Leach-Blilely Act and 21 CFR Part 11, a pharmaceutical and biotechnology industry regulation.*

*ryan.kalember@verisign.com*

# The Make or Break Role of Information Lifecycle Management

## COMPLIANCE, BUSINESS CONTINUITY, AND INCREASING BUSINESS AGILITY

BY DAN SOCCI

The specter of multimillion-dollar fines for regulatory non-compliance is a definite motivator when it comes to data retention. And there are equally drastic consequences, including negative impact on customer service, costs, productivity, and speed to market if data is inaccessible. But while we're all aware of the urgency of setting policies, implementing technology, and instituting processes to manage data effectively, new complications are increasing the challenges and threatening your business.

The problem itself has actually changed shape. The familiar graph showing exponential growth in volume no longer represents the full scope. One issue is the expanding range of types of data that businesses must preserve and access. There are electronic documents such as contracts, invoices, and presentations, as well as CAD/CAM designs and certain types of digitized information such as check images, blueprints, historical documents, medical images, video, instant messages, and photographs. Increasing volumes of e-mail, e-mail attachments, source code, and Web content add to the complexity of the challenge.

Top of mind are e-mail retention efforts because of mandatory compliance requirements, and business best practices so that companies can mitigate risk. Additionally, it makes sense to be able to exploit data assets for increased productivity and to keep a record of communications with customers, partners, suppliers, and employees; and there is future value in providing access to "corporate memory."

More than 10,000 federal and state laws and regulations mandate the maintenance of electronic information – and that's in the United States alone.

In the financial industry, Sarbanes-Oxley and new SEC rules are driving changes; and in health care, HIPAA, PHI, and Part 11 are changing the storage and management of data; EPA and ISO are mandatory for manufacturing firms; as are FDA requirements for CGMP, and 21 CFR Part 11 Life Sciences for pharmaceutical and medical device companies (see Figure 1).

In addition, for every industry, speed to market and effective knowledge management are business-critical issues. And across industries, the increasing consolidation of IT environments, specifically database consolidation, is creating larger data stores where aging algorithms are required to maintain the performance of business applications.

Avoiding legal and regulatory consequences is the negative side of the story. The positive side is that storing and managing information in strategic ways is critical for an agile, efficient enterprise. Many businesses have initiatives in progress to address both the requirements and the opportunities.

One response is storage expansion. According to a study by the Meta Group, the average business is growing storage capacity by about 45% annually, as opposed to 30% at the beginning of 2003.

A word of caution here: expanding storage capacity alone is not an effective solution. Business policies must be shaped to solve the problem, and information lifecycle management (ILM) is essential. We're all hearing the term frequently these days and in many contexts. Our definition here is that ILM is the process of managing data throughout its life cycle, according to the value of that data to the business.

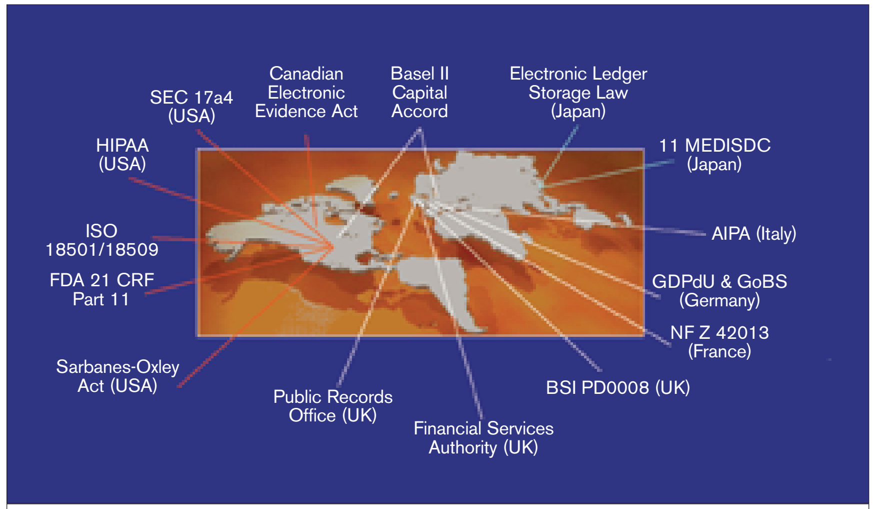

**Figure 1:** The data minefield

It's very important to understand that ILM is not simply a matter of placing data on the most appropriate, cost-effective storage media during its life cycle. In fact, in order to be truly successful with the design and implementation of your ILM solution, it is important to remember that only about 25% of the challenge is about the selection of products. Key to getting it right is implementing the right business processes and supporting them with the right tools.

Data must be stored based on the objectives of the business. How the data is used, its value, and how its usage and value can change over time are all critical to effective ILM. The implementation of these processes may encompass content identification, backup and recovery, replication, archiving, data migration, and data distribution as well as robust indexing and search functions and processes for permanent removal of data (see Figure 2).

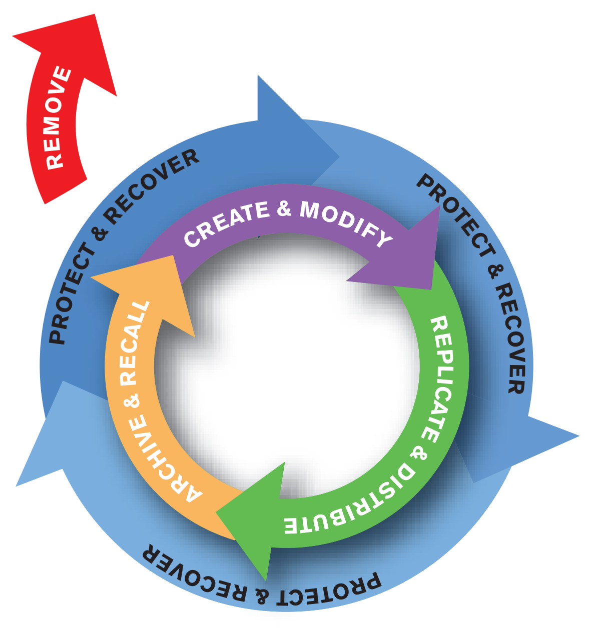

**Figure 2:** Information lifecycle management (ILM)

An effective ILM solution is composed of an interconnected set of business processes, storage components, and data and storage management applications. The fundamental steps to achieve this are defining the business processes and designing and implementing the architecture to support them. A successful ILM solution is a phased approach that can scale and adapt with changing business needs. Only after the design of the business processes is completed should your efforts turn to the careful selection of the right hardware and software to support those processes.

### Essential Steps for Developing an ILM Strategy

ILM is not something you buy off the shelf. A number of steps are critical to developing a successful ILM strategy. For example, an initial assessment and planning stage is necessary to determine the needs of the business, the scope of the project, custom solution design, and requirements for such essentials as administrator training and ongoing security.

The following section briefly describes some of the steps vital to developing your strategy and examines their significance within the overall ILM solution.

- *Application, data, business, and regulatory inventory:* During this extremely important first step your team needs to evaluate your applications, your data, and its value, and build an inventory of the candidates for ILM.
- *Solution design (capacity and performance planning, hardware selection, network integration):* Once target applications and the data stores are identified, the existing business processes must be reviewed and adjusted to ensure the value of the data is reflected. The value of the data is influenced through regulatory requirements, business needs, and cost factors. After the business processes are properly designed to meet all needs, basic metrics such as total capacity required, throughput, availability, and response time are used to design the supporting architecture. The team must also consider factors such as the skill base of the existing system administration staff, the potential for existing hardware to be repurposed, and business continuity requirements as they design a solution customized to fit your business needs.
- *Security audit:* The team must ensure that security matters are addressed in the design of the initial solution. This includes defining processes for ongoing security reviews and patch updates.
- *Document retention policies:* Once the solution architecture is in place, additional business rules must be defined to drive the movement of data between the various tiers of storage and signal the appropriate time for data deletion.
- *Chain of custody:* Data must often be retained and remain accessible even after the application or computer platform that generated it is retired.

"The explosive growth of static data and the need to store and access it effectively **are challenges that are here to stay"**

In these cases, your team will need to transform the data as it moves to different storage layers and validate that each transformation maintains the integrity of the data from an end-user perspective as well as from a legal and regulatory perspective.

- *Auditing:* To ensure that your company is in compliance with all relevant information and retention regulations, your team will need to examine the auditing requirements of each application, which may include logging system administrators' activities, logging the activities of regulatory users, and tracking all operations on data.
- *Billing and charge-back:* Knowing the cost of data storage, and charging it back to the proper business units, is a key component (and opportunity) of ILM. Your team will need to evaluate the requirements and design a solution around the utilities provided by your solution.
- *Backup and disaster recovery:* Ideally, ILM is a fault-tolerant solution that automatically creates and maintains a redundant copy of each document. If, however, further degrees of redundancy are required, you will need to evaluate the business continuity requirements for each segment of your data and develop appropriate backup and replication strategies.
- *Customization:* Designers and developers should be able to customize any element of your enterprise applications – from user interfaces to connectors and additional data processing – to ensure that the entire system works smoothly in an ILM environment.

## The ILM Imperative: Summary and Conclusion

The explosive growth of static data and the need to store and access it effectively are challenges that are here to stay. By implementing an ILM strategy, IT management takes a proactive step toward making information work for the organization in the most effective way possible.

A successful ILM solution extends the business value that storage provides to the company; supports business compliance with government regulations that mandate the retention, access, authenticity, and privacy of business information; improves business continuity; and increases business agility.

In summary, an effective ILM solution must provide:

- A solid set of business processes designed around the value of the data and the business needs for that data.
- Links between business-critical applications and processes and the adaptive storage infrastructure (to ensure that the right data is available anywhere at anytime according to its business relevance)
- An adaptive storage infrastructure that supports different classes of storage based on how the data is used during its life cycle
- Enterprise storage applications that are designed to simplify the management of complex storage infrastructures

To achieve an effective ILM solution that will scale and adapt as business needs grow and change, you must make careful choices when developing your business's ILM strategy, ideally before selecting the appropriate hardware and software. Consideration must be given to an initial needs analysis, detailed project scoping, custom solution design, a security audit, administrator training, and ongoing services to ensure the correct operation and evolution of your solution according to the unique goals and requirements of your business.

**About the Author**

*Dan Socci is vice president of customer support and channels marketing for HP Services.*

*dan.socci@hp.com*

# Will SAN Complexity Keep Storage Networks from Scaling Up?

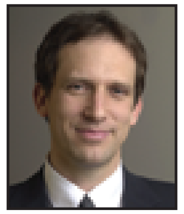

## PAVING THE WAY FOR UTILITY COMPUTING WITH SAN CHANGE MANAGEMENT PROCESSES AND TECHNOLOGY

BY ASSAF LEVY

Storage Area Networks (SANs) have enormous potential to impact much more than storage management. SANs can and should serve as the infrastructure for utility-based processes for the entire IT organization.

Today, this potential is at risk due to the inherent complexity in managing SAN changes, such as adding a server, a switch, or a redundant path between devices. Resolving the SAN change management problem holds great benefits, including:

- Operational efficiencies in managing and growing SANs
- Risk reduction
- Adoption of advanced technologies that facilitate consolidation, virtualization, resource efficiencies and utility-based management

The current state of SAN complexity is an obstacle for scaling the SAN and using it as a shared infrastructure and the basis for future IT improvements. SANs must provide reliable and dynamic service to the IT and business organizations so lines of business can rely on storage to be an always-available utility. However, this is not an easy task, and storage teams at medium and large IT organizations are all facing the same challenges – how to:

- Maintain 100% application availability while applying SAN changes
- Reduce SAN management inefficiencies and support growth without additional resources
- Integrate SAN management into standard IT management procedures through the IT operations team

Storage administration staffs that manage storage networks are beginning to realize that these challenges are difficult to overcome using methods available today or by adding more people to the teams. The core limitation of current methods for changing and growing SANs stems from the sheer number and complexity of SAN access paths and their interdependencies. How will organizations determine the impact of local device changes on end-to-end access paths? SAN complexity increases exponentially as the SAN grows and as new technologies and additional people become involved in the process. This challenge has storage administrators searching for solutions in SAN management software that will allow them to maintain availability while reliably changing and scaling their storage to meet business needs. The solution requires:

- Monitoring and troubleshooting SAN changes and understanding their impact
- Understanding the impact of past changes on access paths
- Conducting root-cause analysis to accelerate problem resolution
- Automating planning and performing simulation to detect errors before they impact the SAN
- Capturing access-path events and change history for auditing and regulatory compliance

Most IT staffers have found themselves poorly equipped to confront the complex maze of access paths in an end-to-end SAN, armed with only archaic spreadsheets and manual tools. Adding staff has not been the answer, since analyzing the vast number of logical and physical interdependencies winding through the SAN gets even more complicated when more people are involved in the change process.

Without a change management framework to validate and automate the change process, the enterprise remains at risk of downtimes, brownouts, security breaches and loss of customer confidence. A recent survey from an IT newsweekly found that managing the complexity of storage networks is one of IT's top challenges. Even a small SAN can have tens of thousands of potential configuration states. A seemingly minor fabric configuration mistake or error in volume masking or cabling can prove devastating, causing data corruption, breaching security and wasting hours of productivity in troubleshooting.

Studies have shown that 25–35% of changes made to a SAN have at least one error, which could be in cabling, port configuration, LUN masking, etc. Many errors – such as lack of redundancy – may remain hidden from view, since data continues to flow until the second path is jeopardized.

Current SAN management tools fail to assist the storage administrator in achieving one of the enterprise's top requirements – end-to-end availability. For a business to realize the full value of a SAN's economy of scale, storage managers must be able to make changes accurately and quickly to keep pace with business requirements. Device monitoring, provisioning, disk utilization, and other software tools provide capabilities that can never be used if the SAN is unstable.

Storage Resource Management (SRM) tools focus on asset management and storage utilization to provide file systems and database utilization levels. However, these capabilities are ineffective if the SAN infrastructure is flawed according to storage analyst Steve Duplessie, founder of Enterprise Strategy Group. SAN change management is a prerequisite "to make all previous investments you have made in storage management, network management, and application management finally return on your investment."

Most storage management problems – and certainly the most complex ones – relate to performing changes. Manual tools are still used to manage changes and are not replaced by SRM tools. SRM software gives IT and storage staffs the impression that they have control over their SAN. In reality, they fail to deal with a SAN's inherent complexity and the difficulty in managing SAN changes.

Urgent and planned changes take days and weeks to complete and storage staffs lack a way to validate changes to ensure that they were made accurately, with accordance to the plan and without any painful downtime. Change cycles for SANs average 10 – 12 days for anticipated changes, and as much as four days for emergency changes. The problem isn't just in validation of the changes and troubleshooting and fixing errors. Storage administrators lack effective control over the different IT groups – such as storage, operations, switches, networking, and cabling – often dispersed across the organization. Control becomes particularly challenging when change directives must be performed in a precise sequence across these disparate groups.

Sadly, the very investment a company made in SAN infrastructure to improve storage efficiencies has become an operational log jam, threatening productivity, business continuity and loss of client confidence.

The solution comes in managing and automating the change process. According to analyst firm Gartner, "improving IT change management processes is generally considered one of the best investments an enterprise can make. Companies that don't properly manage IT changes lose time, money, and efficiency and are subjecting the entire business to undo risk."

Software that detects fatal errors before, during, and after SAN change has recently come onto the market and has been deployed in some of the largest SANs in the world. This software technology continuously maps, simulates, and analyzes the entire storage network in order to troubleshoot errors and find their root causes. Such predictive SAN change management reduces operational complexity, costs, and risk and improves SAN availability, assurance, and customer confidence.

Here is a breakdown of how a predictive change process is implemented through SAN change management software:

- ***SAN change validation:*** An analytical impact model is applied for every SAN change and reports back to the administrator any changes to the SAN along with their analyzed impact on the access path availability, performance, and security.
- ***SAN change troubleshooting:*** Whenever a problem is discovered during the validation phase, its root-cause can be analyzed instantly, and step-by-step roadmaps to resolution are generated. This enables the appropriate fix to take place quickly with very limited impact on service level, potentially before the user is aware of any problems.

- *SAN change audit:* The change management framework enables the capture of the entire change history of all processes and events, in order to generate management summary and trending reports, troubleshoot and validate change implementation, and to facilitate the documentation and audit capabilities of all change history and processes. This is increasingly important as IT comes under scrutiny to provide highly reliable access to information.
- *Planning:* Planning ahead and simulating future changes makes every future change process a predictable and deterministic process. By employing change management software with predictive capabilities, the storage administrator quickly captures and details all required change tasks and actions with their future impact on access paths.
- *Tracking:* The software assists in delegation and coordination of the activities of departments assigned to implement change tasks. The software also logs and tracks every configuration change in the SAN and validates that change tasks have been made correctly, in the proper order and manner, through continuous analyzing of the network.

> "Most storage management problems – and certainly the most complex ones – **relate to performing changes"**

Structured change management's benefits promptly become clear when compared in real life with previous methods, improving accuracy, operational efficiency, and accelerating change times.

Looking to the near future, where utility computing promises to bring new efficiencies to organizations, providing on-demand delivery of applications, computational power, and storage to business units – change management software is essential. Utility computing is based on the ideas of flexibility and efficiency from dynamic allocation of resources to generate competitive advantage and reduce IT costs for enterprises. This of course increases the need for accurate and dynamic changes to storage environments.

Although utility computing is still in its nascent stages, many IT organizations are already taking first steps toward utility computing-based service delivery. These steps usually include changing internal billing to charge for resources used, as well as application consolidation and sharing of infrastructure and applications. Some companies are taking advantage of on-demand pricing from their vendors by purchasing products and services according to actual usage.

To support the utility computing change, IT departments are evaluating software and hardware technologies to assist in on-demand service delivery including storage networks, server clusters, and applications sharing. These technologies promise to provide better resource allocation to meet ever-changing business needs.

## SAN Change Management Enables Utility Computing

Establishing successful utility computing service delivery depends on the control and reach of the storage networking environment. SANs were one of the first utility computing–enabling technologies to become mainstream for many organizations. SANs today have been used for data center and organizational consolidation and, if well-managed, can supply the infrastructure to support dynamic storage changes through centralized storage practices and control.

Utility computing's on-demand delivery of applications cannot take place without an underlying storage networking infrastructure in which:

1. Changes can be made quickly and accurately
2. There is full control over the change process to attain 100% availability of the SAN during any change – large or small.

Change management supports the transition of SANs into an on-demand environment by reducing the risk of business disruptions through better SAN-change planning, predictive assessment, and continuous validation of changes. Additionally, it can increase management efficiencies by freeing the SAN architect to manage the architecture and policies and become better attuned to the needs of internal clients to make SANs an always available utility without adding operational resources.

## Conclusion

As demands for storage capacities rise (the Meta Group has projected a 40–60% annual increase in storage capacities in enterprise data centers), and as utility computing becomes increasingly dynamic, more SAN changes will be required to be performed in a shorter time period, and the technical complexity of networked storage environments will multiply. The confidence of IT professionals in them will likewise diminish without help. Change management technology is required to accelerate the delivery of on-demand SAN service and support on-demand application delivery.

More and more companies are evaluating the need for IT change frameworks to incorporate change management analytical validation models. As the next generation in management software come of age, IT will be able to make the move to the first change management solution for SANs. The most innovative enterprises will take advantage of this opportunity to control and validate their change process helping the SAN reach its promise. ■

**About the Author**

*Assaf Levy is vice president of Product Management and cofounder of Onaro, a Boston-based software company that provides SAN change management software.*

*assaf.levy@onaro.com*

# The Storage Security Problem

## … AND HOW TO PROTECT YOUR NETWORK

BY HIMANSHU DWIVEDI AND ANDY HUBBARD

STORAGE NETWORKS HAVE BECOME critical components of corporate computing environments. Regardless of the type of storage technology, these networks have been designed as if the storage environment and all of the components are already secure because security is provided by other networked systems. Most storage vendors, storage application developers, and storage network designers/engineers operate under the false and dangerous assumption that storage networks are both safe and protected from malicious activity. What's true is that storage networks are just as safe as any other unprotected network. It takes only a single exposed entry point for an attacker to gain access to a storage network and compromise everything the organization is trying to protect.

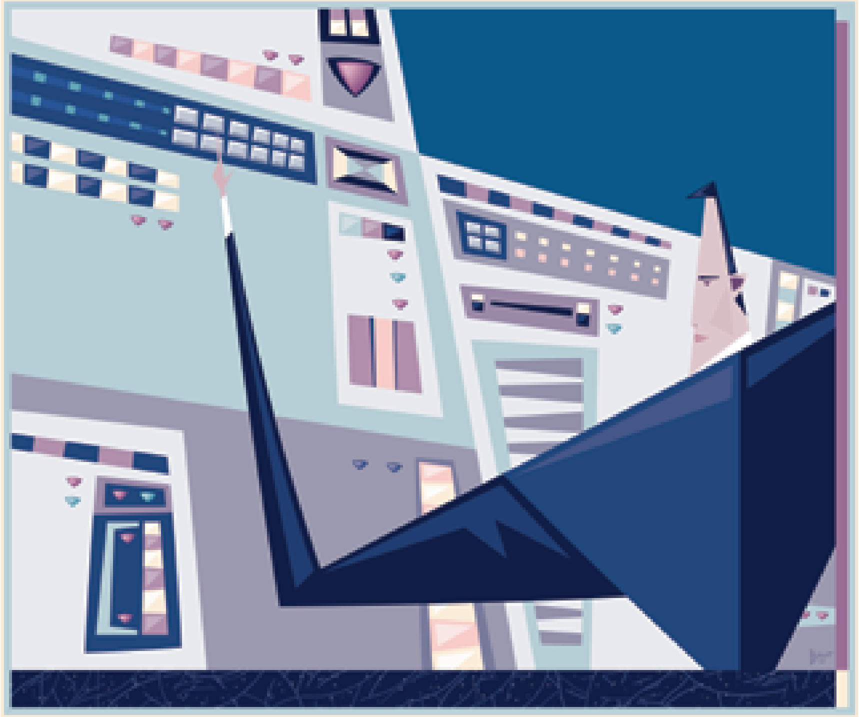

### Elements of Security

There are several basic elements to consider when discussing security. The typical security elements that must be addressed by any secure solution are authentication, authorization, auditing, integrity, encryption, and availability/stability.

Most storage product vendors support these elements to some degree, but not in any uniform, standards-based method. Typically, product vendors focus on only a single component of a storage network, so they only provide for selected elements of security based on a single scenario. This limited focus has a direct impact on the user's environment as a whole.

A complete and secure storage solution must address each element of security. The solution must also address the growth and evolution of the storage environment. In order for products to function together, the newer versions often operate in some form of backwards-compatibility mode. This effectively reduces the security of all of the storage products to that offered by the oldest, and most likely, the weakest version.

The problem doesn't end with backwards compatibility. The storage network environment includes network and host elements that are part of the overall corporate computing environment and may even provide backbone functionality (in the case of switches). These elements are often overlooked as part of the overall security posture.

Overlooked items in terms of security include the storage products themselves as well as any other networking or host equipment that is used to make the environment function. If any one of these elements can be replaced, Trojaned, or subverted, then the entire environment is at risk. While lesser degrees of security may be applied to an environment that is fully contained or localized, the decision to do this and the assumptions made about the design must be understood and recorded. Otherwise, future environmental and functional design changes may fail to take previous design assumptions regarding security into account.

### Security and the SNIA Shared Storage Model

By addressing security in the context of the layering scheme of the SNIA Shared Storage Model, we can easily identify areas where the elements of security can be applied.

If we break the model down into its component parts we can begin to identify where elements of security should be applied to the SNIA Shared Storage Model (see Figure 1). Determining whether or not one or more of the elements of security may be required for the individual layer and how that security is going to be achieved is the important part.

#### Applications

Applications are used to run storage devices, manage storage components, move data, and perform any one of a host of other functions needed for the devices and products in a storage network to function. In effect, every component that makes up a storage network is made up of applications. Therefore, each application must be examined in the context of its ability to be used to attack or defend the storage network. The determination of how security applies to individual elements of the storage network will most effectively be made at the application level.

#### File/Record Layer

Without proper authentication, authorization, auditing, integrity, and availability the components of the file/record layer would easily allow an attacker to bypass security in a number of ways.

Typically, the components of the file/record layer have many of the elements of security built into them. The issue is that the elements of security within these components can be safely ignored if functionality is the only consideration. Databases and file systems are often configured "out of the box" with little in the way of applied security options enabled. This is due primarily to the fact that default installations do not require that either the database or the file system it uses be configured in any way other than simple defaults.

Whether CIFS, NFS, SQL, FTP, or some other proprietary protocol is used, there are risks with the types of communication that are routinely established in the file/record layer of storage networks. These protocols are integral to the file/record layer components and their security components for their ease of deployment and with which disparate systems can be integrated into a shared environment.

#### Block Aggregation

The interoperability and compatibility issues that come from integrating disparate host, network, and device components often introduce new security challenges within the block aggregation layer. Each of these components requires some level of security to function safely and properly. These components must address security at both an individual as well as a unit level. These components may all come from different vendors that have made different design assumptions. The overall storage network design may call for certain component level capabilities that simply do not exist within the component used.

#### Storage Devices

By themselves, storage devices are basically inert objects that await commands from some form of controller (disk, server, storage, etc.). Yet they can understand device drivers, they can understand function calls, and they can establish communication to other devices. Therefore, it is important to understand how these devices function and how they could be compromised. For example, an attacker could use this capability to install rogue applications in virtually any location on a storage network – because that rogue application could interface directly with the storage devices.

#### Authentication

Authentication methods for storage networks like Simple Name Servers, basic end-user authentication, and hard-coded username/password combinations are simplistic and easy to defeat.

Authentication should encompass not only the users of storage systems, but also the devices and applications with which the storage system interacts. In many environments, any component of the storage network can be replaced or added without authentication. And in others, storage applications can be introduced into the environment with no form of authentication other than communicating with the appropriate protocol or utilization of an accepted SDK or API.

Storage networking components can be easily attacked due to weaknesses within their authentication mechanisms. Even environments that have deployed advanced forms of authentication can be attacked if the implementation of these mechanisms is faulty. The strength of any authentication mechanism is based on the quality of the implementation and the strength of the credentials. If the credentials are weak, or if authentication data is exposed due to faulty implementation, the mechanism itself can and will be defeated.

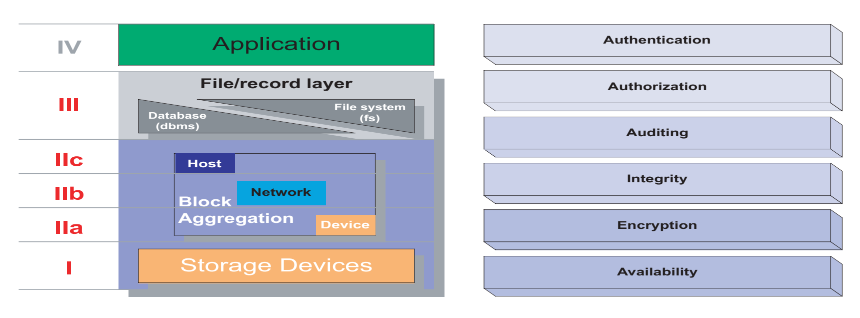

**Figure 1:** Layers of security (right) map closely to the SNIA Shared Storage Model (left)

#### Authorization

In the case of pure networking components, the authorization components are built into the networking gear and may or may not be tied into the advanced authentication/authorization systems that are in common use in larger networks today. In the case of multi-vendor storage networks, there is a wide variety of authorization implementations due to the wide variety of storage hosts, storage devices, and the file system and database components.

User, application, and system authorization are all critical to the security of the overall storage environment. Administrators must ensure that authorization information is not lost during transit from the originating system (the storage client) through some form of intermediary (a storage controller, caching engine, etc.) and eventually to some form of storage device. It is also important to ensure that the credentials that are associated with user access are appropriately understood by all elements of the storage environment and that they can be acted upon (i.e., user rights, disk quotas, or specific file system attributes).

Authorization works best when it reflects discrete roles, which encompass users, devices, and applications within the environment. Controls around authorization must be designed with the overall environment in mind. This makes it difficult for administrators of existing storage networks, especially early adopters of storage technologies, as many of the components that currently exist may have been inherited and therefore may not be fully understood.

Failure to identify how and when objects or resources need to be accessed during design will result in lax or nonexistent access controls or authorizations. For example, access to critical files, especially log, temp, cache, configuration, and database files must be closely guarded and limited to privileged accounts. If these files are not protected with proper access controls, or if the access controls can be bypassed in some way, users can essentially gain access to data that may allow them to elevate their privileges.

> **“Security plays a vital supporting role**
> in enterprise storage networks”

### Auditing

The ability of the systems within the storage environment to capture and retain log information pertaining to access and modification of data is paramount to the security of the overall environment.

All storage network components must be able to capture and maintain log information, either remotely or locally; this includes networking components, hosts, storage devices, and storage applications. While these various components of the storage environment may capture and record log information in different ways, they must have the capability to log pertinent information in context.

Additionally, the ability to log both remotely and locally is important for trend analysis and shared security infrastructure. In order to understand security threats and manage security breaches, the administrator should create a mechanism to allow containment of a remote logging device for the storage network to identify trends, anomalies, and suspicious activity. Most storage products today relegate logging and log reporting to other components of the storage network. While many storage applications and storage products have some capability to capture and display log information, standards and formats are inconsistent, and the amount and quality of detail vary widely.

Many systems are completely proprietary in nature, making the import and export of logging data into a third-party system difficult. As with other networks, many storage network environments support only limited logging capabilities, and administrators tend to accept the default configuration. In other cases logs are not properly protected or may be accessed by users, even those with limited privileges. Malicious attackers know this, and take advantage of both the product's default logging features (which are limited) and the average administrator's reluctance to change them. As a result, attacks sometimes go unnoticed. This dynamic presents opportunity for attack of both storage technology (hardware and software) as well as the networking gear that supports the storage network (routers, switches, and hosts).

Sometimes the simplest solution is the best one. Since the de facto standard for logging of information throughout the computing industry is syslog, it would be ideal for storage network components and applications, in the future, to have some means of delivering log information in this format.

### Integrity

It goes without saying that storage security must not in any way compromise the storage environment or the data it manages. This requires that the system provide some means to confirm that integrity has been maintained over time. While storage solution vendors provide some means for ensuring integrity through their product offerings, the integrity of the system remains open to compromise because there is no accounting for the integrity of the networking or switching components that provide the storage system's foundation.

To the trained security professional (or malicious attacker), these network components are obvious attack points. If the storage vendors don't provide helpful security guidelines for the secure deployment of their components, their customers are at risk.

The integrity of the components of the storage network and the configuration of those components is just as important as maintaining data integrity. If an attacker can Trojan or replace a component of the storage network, then he/she can force nearly any change that is desired into that network, up to and including capture or destruction of data.

### Encryption

Data encryption for storage networks is still in its infancy. Few storage network architectures take advantage of the benefits of encryption, which can be blamed to some degree on design considerations and functionality tradeoffs when encryption is put to use. The process of encrypting data can be very costly and the tradeoffs significantly impact the performance of any network. Encryption brings with it the requirements to both protect encryption keys and escrow them in the case of a catastrophic system failure. While a malicious user may attempt to steal an encryption key and thus be able to steal usable information from a storage network, it is a far greater risk that in the event of a system failure the loss of an encryption key could render all data upon a given disk array completely irretrievable.

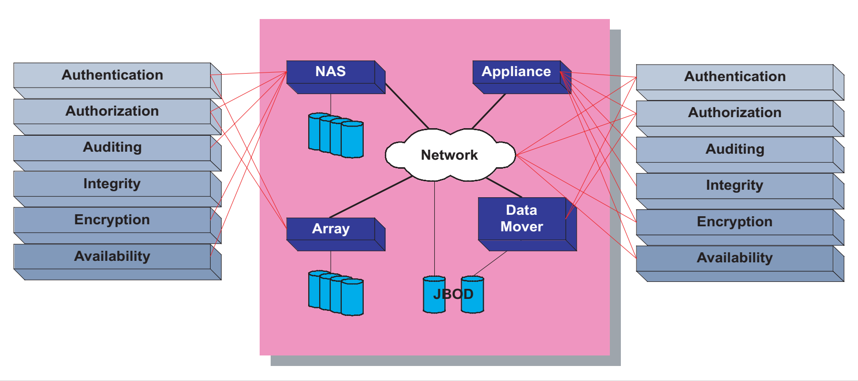

**Figure 2:** Security elements in storage network design

Assuming design considerations and functionality issues are resolved, encryption is not a security panacea. Encryption can protect against data theft, prevent certain forms of hijacking of data, protect network traffic, and even prevent attacking systems from successfully communicating with intended targets. However, encryption cannot protect against the willful destruction of data, which can still be deleted or tampered with in a fashion that will render it useless.

As a security best practice, storage environments must have the ability to encrypt data both in transit and at rest. Since storage environments can be used in many different ways and can have many different customers, steps should be taken to ensure that data is encrypted before it even reaches the storage network. This does not remove the responsibility for providing this capability from the storage vendor, but it is also good practice on the part of the eventual end-user of the environment. This is especially important for users of shared storage environments.

### Availability/Stability

Availability and stability of systems are hallmarks of successful products. Unless alternatives are limited or nonexistent, users will not put their faith in products that are regularly unavailable or are often thrown into an unstable state. Many storage solutions are susceptible to simple denial of service (DoS) or flooding attacks. The likelihood of these attacks occurring is reduced only by the location of the storage network. As storage networks proliferate, they have a tendency to migrate towards the edge of corporate networks, increasing the likelihood that they come under attack. Furthermore, DoS attacks and flooding attacks are common methods used to force systems into an unstable state or force systems to invoke a down-level protocol. This can be part of a larger attack that necessitates the target being weakened in some way. Smart attackers can target relatively unprotected storage networks in order to compromise other corporate information networks or assets.

Overall system security is a requirement for any environment in order to guarantee availability and stability. If the environment cannot resist even simple attacks, then it cannot be maintained in an available state. In the case of some storage network and some storage product designs, availability is addressed by simply supplying more of the same resource to the resource pool. This will not protect the storage environment from automated attacks or malicious mobile code; it will simply result in more of the same type of resource being damaged. The end result will still be a storage network that is unavailable.

## Elements of Storage Design

Storage network design must take security of both the environment and the data into account. Figure 2 describes a simple storage design that spans multiple networks, and presents configurations that enable communication for this type of network along with potential security risks.

### Storage Network Design

As the demand for storage technology increases, it makes economic sense to combine the benefits of storage networks with those of existing network investments. Without proper planning, doing this can actually have negative effects on the security of the existing enterprise network.

Some aspects of a storage network design may look similar to an Out-of-Band (OOB) management network. In these cases the storage network may effectively transit many different security zones, providing attackers with access to a transit network that bypasses security from externally attached networks into the core. Most attackers understand the basics of network management, of which storage solutions may be considered a part, and know how to take advantage of the protocols and applications used to communicate between these systems and environments.

As stated previously, the storage devices and applications may not be the ultimate target of attack, but their vulnerability to attack may make it easy for an attacker to reach resources on the attached enterprise network. In this case, attackers rely on the fact that administrators may cut corners in order to make multi-vendor networking and storage technologies work together.

The converse may also be true. In environments that have grown to depend on storage technology, it is quite possible to introduce connectivity into the storage environment from unanticipated sources. This is a danger in any network, but even more so in storage networks, as many of the components of storage technology within them are critically dependent on the security of the storage environment being maintained.

### Product Functionality/Interoperability

Interoperability and functionality are issues that have plagued network and host systems for years. In the case of storage networks it is again an issue of balancing security needs with system requirements of stability, functionality, and performance. Some storage products require such specific configurations that the introduction of some security technologies has a deleterious effect on system performance. In the case of a localized storage network the risks of allowing some protocols or some types of system configuration are relatively limited as the environment is known and well contained. But, when an environment of this type is expanded or connected to other networks, the previously acceptable risks become security nightmares.

Many storage products actually introduce considerable security risks to a network if all of the functionality of the product is enabled. Some simple examples of this are Web-based management, SNMP-based management, and the use of a large number of ports for communications between product components. Fortunately, each of these issues is easily resolved, but in some cases they require additional layers of protection and design. Many of these issues could easily have been prevented by the vendors through more secure product design.

Additional issues arise when product vendors base their product design on third-party solutions. For example, storage controllers are dependent on the base operating system upon which they run. If that OS is taken down frequently due to patch administration and upgrades, the stability and functionality of the storage solution are reduced.

The problem of product maintenance quickly becomes extremely complex. If the vendor is responsible for support of both the storage component and the supporting infrastructure (the OS) component, then that vendor must devote resources to both understanding the patch cycle of the components and managing each product's maintenance schedule. The vendor must also develop methods of updating the product in a fashion that is easily understood by the eventual end user, who may be a storage operations engineer or a systems engineer.

If the vendor product team is not responsible for the maintenance of the component, then both the component and the storage product are exposed to those attacks to which the component may be vulnerable. Unfortunately, it takes only a few days or weeks for attacks to spread among attackers, often leading to a simple attack vector becoming executed against every buyer of a given product line, while the victim companies await the fix or patch from the vendor.

### Applications

Storage applications cover the implementation of everything from commercial off-the-shelf (COTS) applications to proprietary applications developed in-house, to software development kits (SDKs) and application programming interfaces (APIs) used to enhance storage solutions. These applications represent major components of the inner workings of the storage environment. As a result, they are all the more attractive to attackers and have become the favorite targets.

Application attack techniques have advanced exponentially in the last few years. Unfortunately, quality engineering, testing processes, and security awareness within software development teams has failed to keep pace. Developers of storage solution software and storage applications, both commercial developers and in-house development teams, often fail to consider what would happen if an attacker gained direct network access via the storage application or device.

## Conclusion

Security plays a vital supporting role in enterprise storage networks. As storage networks proliferate and become more integrated within the enterprise network, companies need to put appropriate security plans in place to adequately protect intellectual property. By viewing security as a system of interconnected processes and technologies, companies can still provide appropriate support for requirements such as functionally, throughput, and design simplicity.

This security storage provides a foundation for security professionals who need to understand security issues as they pertain to storage networks. The Security Storage Model puts security in the context of storage and makes it easier for the average storage administrator to include security issues in the design, creation, implementation, and maintenance of any storage network without incurring unnecessary overhead, negatively impacting functionality or compromising the integrity of the data. ■

### About the Authors

*Himanshu Dwivedi is a regional technical director at @stake, Inc., where he leads the Storage Center of Excellence (CoE), which focuses research and training around storage technology, including Network Attached Storage (NAS) and Storage Area Networks (SAN). Himanshu is considered an industry expert in the area of SAN security, specifically fibre channel and iSCSI security. He has given numerous presentations and workshops regarding the security in SANs, and currently has a patent pending on a storage design architecture that he co-developed with other @stake professionals (U.S. Patent Serial No. 10/198,728).*
*hdwivedi@atstake.com*

*Andy Hubbard is a regional technical director at @stake, Inc., and has been a computer security professional for the past seven years. While at @stake he has worked on projects that range from security assessments and secure design of security infrastructures for Fortune 1000 companies to training and curriculum management for @stake Academy. His most recent projects have included product assessments for a variety of storage and enterprise management software products, focusing on functionality, ease of use, and resistance to internal attack.*
*ahubbard@atstake.com*

### ISSJ | Advertiser Index

| Advertiser | URL | Contact | Page |
|---|---|---|---|
| Infitech | www.infi-tech.com | 800-560-6550 | 12 |
| ISSJ | www.issjournal.com | 888-303-5282 | 11 |
| IT Solutions Guide | www.sys-con.com/it | 201-802-3021 | 15 |
| Quantum Corporation | www.dlttape.com | 720-406-5700 | 2 |
| RSA Conference | www.rsaconference.com | 617-848-8756 | 17 |
| SafeNet | www.safenet-inc.com | 800-533-3958 | 36 |
| SYS-CON e-newsletter | www.sys-con.com | 888-303-5282 | 27 |
| SYS-CON Reprint | www.sys-con.com | 201-802-3026 | 29 |
| VeriSign | www.verisign.com/dm/mss | 650-961-7500 | 7 |
| Veritaswww.veritas.com |  | 800-327-2232 | 35 |
| Web Services Edge | www.sys-con.com/edge | 201-802-3066 | 30 & 31 |

THIS INDEX IS PROVIDED AS AN ADDITIONAL SERVICE TO OUR READERS.
THE PUBLISHER DOES NOT ASSUME ANY LIABILITY FOR ERRORS AND OMMISSIONS.

# Bridging the Gaps

## *FINDING THE ROAD TO THE TOTAL ENTERPRISE*

BY DEVIN REDMOND

SECURITY THREATS HAVE dramatically increased for Internet Protocol (IP) networks, applications, and the enterprises that rely on them. These threats come in many forms, from external and internal hackers, to viruses worms; and they threaten enterprises from beyond the perimeter, inside the firewall, and down to individual database files or communications.

With this increase in security threats, a host of solutions has emerged. Each group in an enterprise IT department is increasingly tasked and given budget to solve their security threats with one or more of these solutions. This patchwork of security solutions is where the real challenge for the enterprise begins.

Typically, an enterprise IT department is divided into different departments or areas of responsibility – networking, applications, desktop management, etc. Each group usually maintains its own priorities, agendas, and budgets. Security initiatives are relegated according to the goals of each group (or what they do not want to be responsible for). These three different agendas are the beginning of the breakdown for providing unified security.

For example, the network group will usually focus on protecting network access and access to IP services, using solutions such as firewalls, strong authentication, and remote access via IPSec or SSL VPNs. The application team will focus on protecting their application servers and access to those servers via file encryption, two-factor authentication, and an application extranet with SSL encryption for remote application users. Finally, the desktop team uses some type of application control to prevent hosts from using prohibited applications. To protect the endpoints, the desktop team uses desktop firewalls, IDS, and virus scanning.

In a perfect world – one without time constraints and coinciding schedules and priorities – vendors would have unified solutions for each threat. Without any political boundaries between these functional areas in the enterprise, these groups would implement a unified solution that covers each of their requirements – with a total lower cost of implementation.

Unfortunately, in the real world that's not how enterprise IT departments operate. Rather, most enterprises have overlapping solutions that result in a higher total cost of ownership without solving key threats. As a result, security is not unified in its deployment, leaving a high risk of vulnerability gaps as well as inefficiencies across the enterprise.

A common threat example is a network team that creates a remote access environment with a VPN and strong RADIUS authentication, but they don't have responsibility for the desktop. And the desktop team hasn't deployed a comprehensive desktop security solution. Therefore, users accessing the network remotely can be compromised by hackers and viruses and can compromise the network even though they are encrypted and authenticated.

At the same time, inefficiency emerges as the network team implements RADIUS for user authentication while the application team is using USB tokens for two-factor authentication and file encryption. Not only do network users have to deal with both RADIUS username and password and their token and its related username and pin code, but the enterprise is now paying for two different user authentication solutions.

What can enterprises do to address these challenges? While there is no shortcut, using the following guidelines should ensure that the enterprise goals are addressed along with those of the individual IT teams.

> Take a step back and review each of the security concerns that face the IT teams.
> Match those concerns with corresponding group initiatives to reduce risks.
> Review the various solutions that exist or are being evaluated, identify any overlap between them, and try to consolidate around that overlap.
> Identify the solutions that best meet the variety of needs and reduce the total cost of ownership.

For example, in the earlier scenario, if the network, security, and desktop groups had reviewed their respective requirements they could have prevented new risks, provided a more unified security model, and reduced costs. The network and application teams could have consolidated their authentication model around the two-factor USB solution, and reduced the management and cost of two authentication solutions. Also, those two teams could have also consolidated extranet access and general network remote access initiatives around SSL and IPSec VPNs. Then the two teams could work with the desktop team to protect the desktop and control application access with an endpoint security solution. This process may create some "political" issues but it would also reduce the number of solutions deployed and cut the cost of duplicated solutions, and increase the total budget available to address security issues and provide a unified security approach. ■

### About the Author

*With over nine years of Internet and networking technology experience, Devin Redmond has worked for some of the world's leading technology vendors, including product line management, technical marketing, and business dvelopment roles for Check Point Software Technologies-MetaInfo, Neoteris, Panthesis, RealNetworks, and ViAir. Devin has specifically focused on working with global Internet standards organizations as well as enterprises, key equipment vendors, carriers, and application providers. He is currently a director with MetaInfo (www.metainfo.com) in their efforts to unify IP infrastructure management with key facets of security.*
*devinr@metainfo.com*

# A Holistic Approach to Securing the Enterprise

ASSURING THE SECURITY AND AVAILABILITY OF YOUR IT INFRASTRUCTURE

BY DON KLEINSCHNITZ

THE CONTINUANCE OF MALICIOUS computer attacks has made security a front page topic in almost every boardroom and IT oversight committee. Most IT departments accept that routine updates to software operating environments are a necessary part of managing systems.

It's not hard to convince the IT professional that the protection of data assets forms the foundation of recovering from a disruptive event. But very seldom do we think of security, systems, and storage management as part of a seamless and holistic approach to securing the enterprise. Considering the rate at which vulnerabilities show up in our computing environment and the speed at which they can be exploited, we need to rethink how these three management environments should be leveraged – after all, "The only truly secure infrastructure is a managed infrastructure!"

As the list below suggests, the administrative job of managing and securing the enterprise is complex and convoluted, with loosely integrated software that attempts to automate the normal operations of the enterprise.

> Firewall management
> Virus definition updates
> Data backup
> Applications update
> Software licensing compliance
> Vulnerability assessment
> Disaster recovery
> Storage provisioning
> OS upgrade and provisioning
> Archive policy
> File recovery
> Asset inventory and reporting
> Repurposing
> Common operating environment policy
> Patch installation

However, in today's heavily exploited environment we must ensure that the security, systems, and storage management elements of the infrastructure can not only manage during normal conditions but also manage effectively through the disruption of an exploit. Stated differently, security, systems, and storage management systems must effectively manage during normal state and disruptive state conditions. Clearly, the disruptive case is the more difficult state to manage.

## What Is a Disruptive State?

When an enterprise has entered a disruptive state it is a serious change in status, evidenced by the number of IT executives that suddenly are visible in meetings, phone calls, and triage sessions. The entire enterprise enters a lockdown as the IT departments identify the threat, determine the vulnerabilities, plan corrections and wait for an exploit. The entire enterprise is holding its breath. The IT organization works long hours to secure servers, desktops, laptops, and most recently, handheld mobile devices. Often the more controlled process and management automations succumb to the deployment of individual experts to manually correct known problems and hunt for leaks in the infrastructure. The frequency, duration, and damage that occurs during disruptive states gives rise to new challenges faced by IT management products.

Managing in the disruptive case requires that the management software be capable of managing through three basic transitional phases: understanding the disruption, controlling the transition, and finally, acting in a way that returns the system to the normal state. This *proactive security system* must rely on the underlying infrastructure to take action and remediate the disruption; therein lies the critical connection between security, systems, and storage.

## Understanding Phase

The system must understand and articulate the origin and nature of the disruption. Security sensors provide the knowledge and understanding necessary to warn enterprises of impending disruptive states.

### Control Phase

Once the management state is recognized as "disrupted," action must be taken in a controlled fashion with the goal of returning the system to its normal state. The control phase provides the rules of execution and the instructional intelligence that the infrastructure must follow during the act phase.

### Act Phase

During the act phase the infrastructure must respond to the disruption in a way that restores it to a normal or "safe" state. Act phase activities include many of the same tasks that are undertaken during the normal state but with an increased focus on the speed and reliability with which they occur. As an example, security patches must be deployed quickly without disruption, whereas the normal process of upgrading operating systems and applications is typically done as a normal course of change management. While security patches are being planned and deployed, the enterprise is vulnerable to damage.

Systems and data recovery is another example of similar processes being executed in the normal and disrupted state. Traditional backup systems back up data during normal operations, but they very seldom focus on processes that will allow a recovery within the window required by most disruptive events. Since many normal and disruptive state management tasks are similar, it is logical to conclude that if we architect for the disruptive state we will also realize improvements in the responsiveness of the normal state management tasks.

It is important to recognize the enterprise-wide scope of managing in the normal and disrupted state. During the transition phase the management software must be capable of connecting to and managing the entire computing environment. This environment includes servers, network devices, desktops, laptops, and handheld devices in both wired and wireless environments.

## The Problems

Consider three key pain points often highlighted during CIO discussions.

### Provisioning

The challenge of migrating and building systems at the rate of arrival of new operating systems has become so difficult that some CIOs see it as a career-threatening event. The process involves determining, first of all, what exactly *is* on every machine in the enterprise, setting the standards for a new operating environment, preparing that environment for deployment, and then finally deploying the change. The whole process takes significant manual activity and expertise and can be so difficult that many organizations still have yet to migrate to Windows XP while a new Windows environment is already inevitable with Microsoft's Longhorn. Provisioning is traditionally a normal state management task, but it is a good example of an area that needs significant improvement through automation.

### Patch Remediation

The ability to completely patch and configure machines presents a large problem – primarily because the threat landscape evolves more quickly than the patch process can update the software. Viruses such as Sasser and Blaster are proof that virus writers will continue to exploit vulnerabilities – Sasser was released into the wild less than three weeks after Microsoft announced the vulnerability it exploited. The window of opportunity in which IT can react to vulnerabilities continues to decrease. Patch management is mostly a disruptive state application, but as stated previously, it can be thought of as a highly responsive component of normal state provisioning.

### Protection and Recovery

It goes without saying that generally data should be protected, but organizations should also have a backup and disaster recovery plan that will help them recover in the event of a successful attack. Data recovery has become a heightened concern because the rate of attack is increasing, so the probability of having to recover is higher. Additionally, having an infrastructure where the accuracy of financial reporting, the privacy of personal information, security, and other process certifications is becoming the personal responsibility of executives. This level of infrastructure accountability is driven by regulations such as Sarbanes-Oxley, HIPAA, and FISMA. The scope of recovery solutions must include desktops, PDAs, servers, and laptops; and must have recovery times that are measured in minutes.

## The Strategy

It is becoming increasingly clear that if we are to evolve the task of managing our infrastructure we need to manage both the normal and disrupted states of the enterprise's operation. Ideally, an organization might have a modular suite of applications that participate in the management of the transition from normal to the disruptive state and back again in a controlled and safe manner. The applications strategy is made up of five modular parts:

> ***Installation design:*** A virtual design environment that simplifies the creation of installation and recovery packages. The goal is to improve and reduce the amount of expertise and effort required to create an installation environment.
> ***Software provisioning and delivery:*** A centralized delivery environment that automates the local and remote installation of computer operating environments.
> ***Patch management and help desk operations:*** Local and remote operations that assure the currency of software and automate problem management.
> ***Asset management:*** This is one piece of the life cycle that is often taken for granted, but it is an important foundation. Auto-discovery, inventory, software usage and license monitoring, plus disposal, repurposing, and reporting are elements of the asset management used by most of the applications in this set.
> ***Protection, recovery, and archive:*** A hardware-independent, local and remote, automated backup, recovery, and archive environment. IT needs the ability to get to full working condition in a short period of time.

A holistic strategy will allow IT organizations to become more efficient – personnel will have more time to focus on important projects rather than dealing with urgent security issues. By involving all relevant IT and management groups with a common goal of securing the enterprise, the solution becomes fully integrated, rather than fragmented. The benefits are numerous, including increased security and availability, reduced human intervention, competitive advantage through rapid response to change, and improved governance and compliance.

By implementing a scalable, platform-independent architecture that addresses security, storage, and systems management, IT will find it much easier to stay on top of that checklist.

***"The only truly secure infrastructure is a managed infrastructure!"***

### About the Author

*Don Kleinschnitz is vice president of product delivery at Symantec's enterprise administration business unit. Don joined Symantec as part of the acquisition of PowerQuest Corporation, where he was chief technology officer and senior vice president of Storage Products. His current responsibilities include oversight of the technology within the Symantec Enterprise Administration (SEA) business unit and site management for Symantec's facility in Orem, Utah.*
*don.kleinschnitz@symantec.com*

# Securing Storage

## COMPLETE DATA ERASURE ON STORAGE SYSTEMS

BY LEO COLBORNE

OUT OF SIGHT, OUT OF MIND. When storage systems are upgraded, retired due to proactive maintenance, reach the end of their lease, or are repurposed or resold, companies often delete the data from the disks and forget about it. However, there is a tremendous amount of critical, confidential, and competitive information on those disks that cannot be completely erased by just pressing a delete button.

This exposes competitive intelligence, increases vulnerability to industrial espionage and litigation, and jeopardizes an organization's compliance with corporate governance practices and state, federal, and industry regulations that protect proprietary and confidential corporate, customer, and patient information. For example, regulations such as DOD Pub. 5220-22.M, Sarbanes-Oxley, and HIPAA require proof of secure erasure.

> "Un-erased information **is still accessible**"

Consequently, it is vital that data be completely erased and the erasure recorded to ensure critical and confidential information is secure from accidental or malicious recovery. Done correctly, data removal meets important compliance regulations and guidelines for erasing data, such as sensitive patient records or financial procedures.

### Why Ensure Erasure?

There are several reasons for completely and provably erasing stored data, including:

- Data disposal and erasure has to conform to industry and other regulatory requirements.
- Potential litigation, loss of intellectual property, or financial loss can result from un-secure data disposal.
- Un-erased information is still accessible when storage systems are returned under lease, redeployed, swapped, or repurposed.
- Corporate guidelines require data erasure and removal of proprietary information prior to returning leased systems or repurposing storage systems.
- Some companies or industries require proof of data erasure and overwrite levels.
- Companies have different data disposal standards for different types of information.
- Some companies and industries require a three-pass or greater overwrite process (recommended in DOD 5220.22-M level).
- Companies have strict security requirements, to retain all disks and you need to secure them.

### Delete That Disk

Most companies know how to implement security measures to protect existing data. However, the options for safely and securely removing data from a drive so it cannot be retrieved are not nearly as advanced. These common measures include one-pass overwrites, degaussing, physical destruction, and physically storing old drives.

- *One-pass overwrites:* Replacing data stored on hard disk drives with a variably bit pattern of 1's and 0's that effectively renders the data unrecoverable. A single pass will successfully overwrite some of the data, but not all disk sectors are visible to overwrite applications. This can leave highly critical information perfectly intact. Multiple passes can yield better results, but the overwrite application must be sophisticated enough to locate and overwrite hidden and damaged sectors, as well as produce audit reports for compliance purposes.
- *Degaussing:* Demagnetizing to remove all data. Degaussing can be effective, but it often leaves the disk drive unusable. This is not a good thing when a company intends to repurpose the drives. It is also not cost-effective to degauss large numbers of high capacity disks in storage systems.
- *Destruction:* Physically crush and shred drives. This destruction is extremely effective in erasing data and can be therapeutic for a stressed-out IT professional. However, it is time consuming, costly, and impractical for retiring a large number of drives.
- *Storing old drives:* Physically storing drives. Presumably drives are erased before being stored, but not necessarily. It has been estimated that 85% of business espionage crimes are inside jobs. So, this technique may make it easier for employees to access retired drives to commit these crimes. And physical storage does not meet most compliance regulations for erasure, nor does it protect a firm in the event of litigation.

### Best Practices

The most efficient, cost-effective, and compliant method of erasing data is to completely overwrite the drive to render the data virtually unrecoverable, and to have the capacity to report the procedure. This is harder than it looks, especially with large and complex storage systems. Companies can assign service levels according to the relative importance of the data; with more overwrite passes for critical information. (Common overwrite levels go from three passes for noncritical data up to seven for the most sensitive information.) Once done, the professional service or erasure application should deliver an independent audit and written proof of service completion.

Observing best practices in data erasure has a number of benefits for security-conscious firms. Complete data erasure maximizes compliance measures by managing risk, ensures information in the life cycle disposal phase is really being disposed, enables that utilization and repurposing storage, and lets IT professionals sleep at night knowing they have secured important data on released storage assets.

### Data Erasure Services

A number of hardware and software vendors provide data erasure services for the PC market, but storage systems are relatively ignored. Due to the sheer size and complexity of storage systems, efficient and complete data erasure is beyond the capabilities of the simpler methods. But managing the data life cycle from creation through deletion includes making sure that data has actually been disposed.

Storage system data erasure services can completely erase data on storage assets and prove they've done it. For example, EMC's non-host-based process completely overwrites proprietary and sensitive data, offers flexible overwrite passes and provides audit reports to meet compliance requirements. Any secure data erasure for storage systems should be able to handle the specific requirements of storage assets, be available from highly trusted professional services (for complete security and audit purposes), erase multiple disks and frames concurrently, have a flexible overwrite pattern for differing specifications, be delivered at the customer location to increase security and eliminate delays, and provide an independent audit and documentation of data erasure.

While firewalls and other security measures protect data on the front end of the storage life cycle, it is equally important to protect data at the back end. When it comes to returning, reselling, repurposing, trading, or swapping out storage assets, companies need secure and complete data erasure to meet corporate governance, industry specifications, and governmental mandates. Reliable and proven data erasure services dramatically reduce potential legal litigation due to uncontrolled distribution or viewing, avoid the physical destruction of perfectly good equipment, and address any security concerns. As a result, companies can safely sell or reuse storage equipment and ensure they have the audit trail necessary to meet corporate and industry conformance requirements. Most importantly, this will protect an organization's most valuable asset – its information. ■

**About the Author**

*Leo Colborne is EMC Corporation's senior vice president for Global Customer Service and is responsible for the overall management, operation, training, tools, infrastructure and resources for the company's industry-leading global support organization (www.emc.com/global_services/ini/data_erasure/index.jsp).*

# Plugging The Processor vs Storage Performance Gap

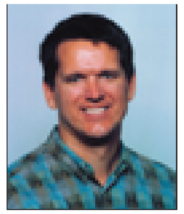

## THE GROWING PRESENCE OF SOLID STATE DISK

BY WOODY HUTSELL

THE EVER-INCREASING SIZE OF applications and databases used to run today's enterprises drives the demand for faster systems. In many cases OLTP (online transaction processing), OLAP (online analytical processing), modeling, and heavy-duty video severing have become so mission critical that system performance directly impacts the bottom line.

While this performance challenge has been met by the processor developers, hard drives (HDD) have not kept pace. Performance improvements of CPU and memory have given rise to a "performance gap" between systems and hard disk drive storage.

Even arrays of 15,000 RPM disk drives are at a disadvantage to processor speed because of the mechanical nature of conventional hard drives versus the electronic nature of processor performance. The spinning platters and mechanical assemblies in HDD systems simply cannot present data quickly enough to today's high performance processors (see Figure 1). This latency leaves many commercial applications running inefficiently and users waiting.

All HDD systems rely on a mechanical moving head and platters to access data. When more and more hard drives are arrayed to increase performance, other problems arise for the data center manager, such as power requirements, heat dissipation, rack space, and an ever-decreasing mean time between failures (MTBF).

How well an application performs is generally measured as I/O operations per second (IOPS); and when performance matters, IT managers have many options. They can add server RAM, build bigger hard drive arrays, or optimize their databases. RAM, monolithic RAID, and database optimization solutions work but they are not the best solution for achieving either the fastest performance or the lowest cost per IOPS. The alternative is a technology that's been around for decades – solid state disk.

This less familiar technology is emerging as the front runner for performance, lower cost per IOPS, and reliability in storage. Solid state disk (SSD) has accelerated applications as high as 25x by eliminating the storage performance bottleneck.

Solid state disk systems use fast-access memory chips as their primary storage medium. SSD does not rely on mechanical parts to input or output data in the way that conventional hard disks do. Rather, SSD uses RAM as the primary storage media. Data is stored directly on RAM chips and accessed from them. This generally results in storage speeds far greater than those that are even theoretically possible with conventional, magnetic storage devices. In order to fully utilize this speed, SSDs typically connect to servers or networks through multiple high speed channels such as Fibre Channel.

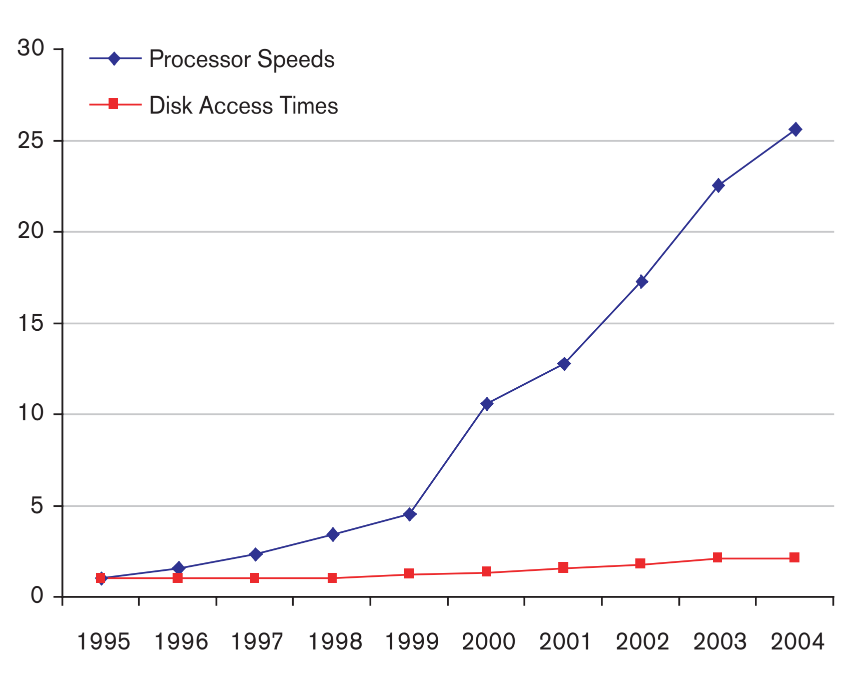

**Figure 1:** Increasing gap between processor speeds and disk access times

SSD delivers low latency and high random IOPS compared to HDD RAID systems. Random I/O performance is a more meaningful metric in assessing the application impact of storage performance than the less practical sequential I/O that is typically published.

Unlike conventional memory, SSD systems are built to be non-volatile. Typically, they include battery power and an internal backup disk. In the case of system shutdown or power loss, the battery powers the unit while the data is mirrored from the RAM to the disk. Internal fans keep the unit cool.

Because there are no mechanical parts in the main data chain to the SSD system, MTBF and reliability are higher and maintenance costs typically lower than with conventional storage.

SSD presents itself in an identical manner to disk or RAID, from a software and system standpoint. Hence, no special management or configuration issues arise. In a SAN environment, SSD can co-exist seamlessly with conventional disk and RAID subsystems. Systems with multiple Fibre Channel ports provide additional throughput and support multiserver connectivity via standard switches.

All, or part, of an application's data may be placed on SSD. For instance, database logs and frequently accessed tables may be placed on SSD, while other components are adequately served by conventional storage. Data that resides on SSD may be shared or migrated in the same way as with standard HDD storage. This is because SSD presents itself to the system and OS in the same way. In many instances, the deployment of SSD has led to significant savings from server consolidation and greater storage capacity utilization.

SSD is not a panacea to all performance problems, however. For this reason, customers usually test SSD solutions before buying, and rely on independent third-party benchmarks to prove the vendor's performance claims. Solid state disks currently hold two different records in the Storage Performance Council's SPC-1 benchmark. They have the fastest recorded SPC-1 IOPS performance; and they have, by a large margin, the smallest price:performance ratio. Having said that, they also present much smaller capacities than HDD solutions, which achieve high IOPS performance by incorporating a large number of disks. Therefore, SSD can be not only the fastest, but the cheapest, performance solution when a fraction of total data is slowing down an entire application.

As applications become increasingly demanding and performance is bound by data access limitations, it is becoming a popular addition for savvy IT departments. Growing SAN adoption, falling SSD prices and an increasing performance gap between application performance and conventional storage are all trends that indicate a growing presence for SSD in the data center. ■

### About the Author

*Woody Hutsell is executive vice president for Texas Memory Systems (www.texmemsys.com), a leading solid state disk manufacturer.*

*woody@texmemsys.com*

web services EDGE conference&expo

# Web Services Edge 2005 East

**International Web Services Conference & Expo**

*New for 2005:*

- Colocating with LinuxWorld Conference & Expo
  badges give access to both shows
- Guaranteed Minimum Attendance 3000 Delegates
- New, Hot Sessions & Seminars

# The Largest *i*-Technology Event of the Year!

## Keynote Speakers & Featured Guests

**Tuesday, February 15**
**11 a.m.**

**Matt Ackley**
Senior Director, eBay Developers Program

**Topic:**
Web Services for e-Commerce

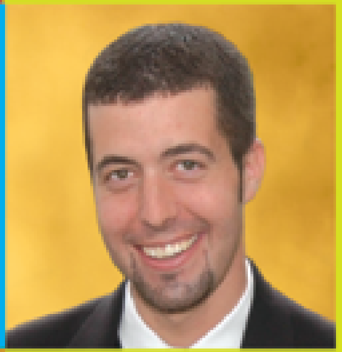

**Wednesday, February 16**
**11 a.m.**

**Ari Bixhorn**
Director, Web Services Strategies, Microsoft Corporation

**Topic:**
Indigo and the Future of Web Services

**Anne Thomas Manes**
Burton Group

**Application Server Shoot-Out Facilitator**

**Hynes Convention Center**
**Boston, MA**
**February 15-17, 2005**

For Exhibit and Sponsorship Information:
Jim Hanchrow, 201-802-3066, jimh@sys-con.com

## Register Now! Hot Sessions & Seminars

- Ensuring Web Services Interoperability
- Web Services Standards: Going Behind the Mask
- The Role of Policy in Web Services Integration – It's More Than Just Security
- Four Abilities SOA will lack without a Registry
- Driving SOA Governance
- BPEL Best Practices from Real-World Projects
- SOA: From Pattern to Production
- Lessons From the Front Line – Building Interoperable Web Services
- Developing E-Commerce Applications with Web Services
- Developing Enterprise Class Web Services
- Orchestrating FORCEnet Engagement Packs with BPEL for Web Services
- CPI: A Global Integrated Problem Tracking and Resolution System using Java Web Services
- Using Service-Oriented Architecture and Web Services to Issue Business Licenses in the District of Columbia
- Developing Web Services with Eclipse
- The Interoperability Challenge of Web Services Security Standards
- Securing Web Services with WS-Security
- Anatomy of a Web Services Attack
- Using a Mobile Phone as an SSO Authentication Device in SOA Solutions
- XML Content Attacks
- XACML and Agnostic Authorities

Register Today: www.sys-con/edge2005east/registernew.cfm

*Come and see what everyone is talking about…*

**The Application Server Shootout!**
*Find out which server best fits your performance and ROI metrics*

## Join us in delivering the latest, freshest, and most proven Web services solutions…

At the *Fifth Annual Web Services Edge 2005 East – International Conference & Expo* as we bring together IT professionals, developers, policy makers, industry leaders and academics to share information and exchange ideas on technology trends and best practices in secure Web services and related topics.

## Exhibit in *The Web Services Pavilion* and see thousands of buyers!

Becoming a Web Services Edge Exhibitor, Sponsor or Partner offers you a unique opportunity to present your organization's message to a targeted audience of Web services professionals. Make your plans now to reach the most qualified software developers, engineers, system architects, analysts, consultants, group leaders, and C-level managers responsible for Web services, initiatives, deployment, development and management at the region's best-known IT business address– The Hynes Convention Center in Boston.

## 3-Day Conference & Education Program features:

- Daily keynotes from companies building successful and secure Web services
- Daily keynote panels from each technology track
- Over 60 sessions and seminars to choose from
- Daily training programs that will cover Web Service Security, J2EE, and ASP.NET
- FREE full-day tutorials on .NET, J2EE, MX, and WebSphere
- Opening night reception

PRODUCED BY SYS-CON EVENTS

**www.sys-con.com/edge**

Sponsored by: Information Storage + Security Journal, wldj, WebSphere Journal, JDJ, XML Journal, WebServices Journal, MX Developer's Journal, SD Times, eaiIndustry Consortium, .NET Developer's Journal, asp.netPRO

All brand and product names mentioned above are trade names, service marks or trademarks of their respective companies.

# Enterprise-wide Intrusion Prevention: Network Security's Next Generation

## STOPPING ZERO-DAY ATTACKS, COMBATING EVOLVING SECURITY THREATS, AND ADDRESSING INTERNAL SECURITY

BY BRENDAN HANNIGAN

NEW SECURITY THREATS are growing in frequency, sophistication, and danger. While perimeter-focused security can mitigate risk from known attacks, real protection comes from identifying and reacting to any new threat the instant it hits your network.

This article looks at enterprise-wide intrusion prevention, a technology recognized by network and security experts as the smart way to combat the many threats facing security managers every day. We'll show how it replaces outward-focused security products with an approach that embeds security throughout the enterprise network.

### What Is Enterprise-wide Intrusion Prevention? Why Do I Need It?

Continued innovation has created many ways to protect against known threats. We evaluate every new attack that hits, spending valuable time analyzing and creating defenses that protect against major worms, viruses, commonly-known hacking vulnerabilities and other threats. Yet a malicious attacker can change only a few lines of code and the same worm, or Trojan will slip right by the reactive signature or patches developed to stop the original. Hackers creatively find new ways to breach traditional signature-based security defenses. Ongoing changes and upgrades in network infrastructures, Web services, and new software continue to create vulnerabilities and opportunities for exploitation.

Perimeter-focused security, which blocks attacks coming from outside, is no longer enough. IT staff really need to understand what constitutes normal network behavior, identify inconsistent behavior, and fix it so business can proceed. Enterprise-wide intrusion prevention profiles network behavior across the extended enterprise, flags anomalies, isolates the source of the issue or attack, and offers a choice of corrective measures to resolve or mitigate the threat. The net gain comes from faster reaction to breaking threats and shortened time to resolution. Business processes suffer little or no impact. That translates into increased uptime and efficiency combined with decreased operational costs and losses.

### How Do I Use Surveillance, Analysis, and Control?

Enterprise-wide intrusion prevention technology models traffic flows, transactions, and network activity and analyzes them to learn what normal behavior, including run-rate activity spikes, looks like. It detects aberrations – changes in traffic levels, communication patterns, or other anomalies that serve as an early warning system for malicious activity – whether from an external attack or internal misuse of the network. Pinpointing suspicious behavior, this technology isolates the source of the anomaly and offers several means of resolution to fix the problem before it causes damage.

Successful enterprise-wide intrusion detection requires a three-tiered approach of surveillance, analysis, and control. Surveillance recognizes malicious activity, catching even the most insidious low-and-slow probes of network defenses without sounding false alarms based on every traffic spike. While firewalls and other appliances provide a limited view from a single point in the network, this technology looks across the entire network.

Behavioral analysis is the key to understanding and applying what is learned from network surveillance. Enterprise-wide intrusion prevention taps both real-time and historical views of network activity to model the behavior of users, applications, servers, and network resources. The latest technology includes a classification engine that profiles network behavior and identifies normal behavior over time. It understands the dynamic complexities of modern networks, recognizing normal and acceptable behavioral changes as safe. It raises an alarm when it perceives potential threats based on deviations from the baseline. Unlike traditional IPS, this technology does not rely on a signature to identify a malicious internal user or an evolving worm. Behavioral analysis recognizes everything from the abnormal behavior caused by a new attack or hacking attempt, to internal threats such as insider scams and stealth attacks. It even finds policy violations among network users who use P2P file sharing and instant messaging, as well as any type of network misuse.

The third element, control, empowers security and network professionals to enforce network behavior. Simply identifying an anomaly is not enough; corrective measures must be taken as soon as possible. New attacks and security threats continue to hit every network with increasing sophistication – and far greater danger. The control element offers a variety of mechanisms for fixing or mitigating the problem. With a choice ranging from automatic remediation to full operator intervention, administrators can address the most critical issues first and focus their valuable time where it's needed most. These systems can address different types of activities in different ways, and are flexible enough to enforce network behavior based on unique customer use. After all, some parts of the network are more critical than others, and different types of threats require different approaches to resolution. Advances in enterprise-wide intrusion prevention technology give IT staff options they have never before enjoyed.

### Top 10 Benefits of Enterprise-wide Intrusion Prevention

1. **Provides an enterprise-wide security system:** Holistic enterprise-wide view of security goes beyond segment-based, perimeter-focused point products.
2. **Stops external threats:** Provides the first (and often only) defense against the proliferation of zero-day, blended, and internal threats, without the time delays or alarm overload of signature-based systems. This means identifying and locating worms, Trojans, denial of service, and blended/hybrid threats quickly and providing automated resolution.
3. **Enforces internal policies:** Exposes and locates internal threats so you can stop them quickly and eliminate future problems, whether from violation of internal policies or intentional misuse. Such misuse wastes resources and exposes enterprises to unnecessary legal and security risk.
4. **Ensures regulatory compliance:** Provides monitoring, detection, alerts, and audit trails to comply with new regulations and compliance issues that demand IT participation.
5. **Avoids legal risks and liabilities:** Provides the processes and information to protect your organization against risks and liabilities such as lawsuits from illegal file sharing of copyrighted material, lawsuits from accidental disclosure of confidential information, and penalties for noncompliance with regulations.
6. **Improve operational efficiency:** Identifies problems quickly, isolating the source of network bandwidth issues or security threats to speed resolution without additional staff.
7. **Secures the "perimeter-free" network:** Protects open, distributed networks from potential threats for the most advanced defense of infrastructures that can't rely on perimeter security solutions.
8. **Eliminates breaches from mis-configured systems:** Identifies network misconfigurations quickly and effectively, isolating the source to close vulnerabilities and conduits for hackers.
9. **Provides live window of network activity:** Gives network and security administrators an instant real-time view into network behavior, along with access to terabytes of data. It identifies issues in real time and archives a complete audit log of activity without costly additional storage requirements.
10. **Combines network and security analysis:** Integrating asset discovery, vulnerability data, and observed network profiling provides context-sensitive detection of known events. Pivoting between security and network data simplifies the process of finding, fixing, and preventing threats.

### Where Does Enterprise-wide Intrusion Prevention Fit In My Security Strategy?

In a crowded security market, every vendor hypes a different technology as the most critical element of a layered security defense. So where does enterprise-wide intrusion prevention fit in your security strategy and network architecture?

This technology incorporates security event feeds and network traffic flows from your existing infrastructure to leverage its data completely. But the most direct value it provides, and the primary reason people choose these systems, is to address the critical flaws of traditional signature-based technologies: addressing internal security concerns and stopping subtle blended threats and zero-day attacks. The bulk of ongoing security expenses, and the biggest nightmare for security and network managers, is identifying, reacting to, and cleaning up damage from the "next big attack." No other technology matches the ability of enterprise-wide intrusion prevention to defend against new attacks that are as unpredictable as they are inevitable. It serves as the first-responder product for identifying, understanding, controlling and fixing any new attack. ■

**About the Author**

*Brendan Hannigan, executive vice president of Marketing & Product Engineering, brings over 16 years of industry experience to Q1 Labs. Before joining Q1 Labs, he was vice president of marketing at Sockeye Networks (a route-optimization firm acquired by Internap), where he led all marketing and product management functions. As director of network research at Forrester Research, he oversaw the firm's most successful practices, covering enterprise networks, security technology, and public network services.*

*info@q1labs.com*

# SANs and NAS

## IMPROVED EFFICIENCY THROUGH VIRTUALIZATION

BY GUY BUNKER

SANS, NAS, ISCSI, virtualization, in-band, out-of-band, the terminology seems never ending when it comes to storage and what's worse, no one will tell you what's best. Unfortunately, it's not that simple. The advent of SANs and the introduction of new technology has increased the number of options available, but there are no clear guidelines as to which one to use and when. There isn't a silver bullet or golden configuration that is good for everyone; the solution has to be tailored to the specific environment. But all is not lost. There has been a lot written about storage and storage architectures, and if all else fails, look at what you are trying to achieve and how much money you have to spend.

"Storage is not just about **the online disk**"

While it is widely thought that SANs are for big enterprises and NAS for smaller ones, this is not true. Most enterprises, whether big or small, now have NAS servers and many are using them for more than just file serving. The cost of SANs has fallen such that they are now a very real prospect for smaller organizations that want to take advantage of improved connectivity and performance to utilize with technologies such as third-party copy and clustered file systems. So it is the applications and the business requirements that should drive the architecture, not the "latest and greatest" technology or the cheapest solution. Storage is not just about the online disk. Backup (which now might be to disk before going to tape), disaster recovery, and legislative compliance all have their part to play. Without a big picture of what needs to be achieved (from the business perspective) the decisions made will be insufficient. Another factor to include is storage growth. If the space required in 12 months is 100% more than you have today, will that influence your architecture decision? What happens if it is 1000% in three years? How long do you plan to remain with the architecture that has been defined? The immediate logical conclusion is to go for the biggest you can buy – now. But we know this is not a pragmatic business decision; the architecture should be designed so that it can be grown – and this might mean starting with NAS and expanding into a SAN just as much as starting with SAN and acquiring a NAS solution later.

Utility computing is a trend we are hearing a great deal about, with many vendors touting it as the next big thing. When it comes to storage, applying utility computing principles and creating a storage utility is a great place to start. By using storage virtualization tools storage can be pooled and then provisioned when required; by having it attached to a SAN it can be allocated to any server that needs it. Additional functionality allows file systems to be grown automatically without the need to take the application using it down. Business reporting tools enable departments (or lines of business) to see how much storage they are using. The IT organization can then choose to apply costs to the storage and could present each business with a bill (a.k.a., chargeback) if they so wished. More often than not it is the insight into costs that is useful, and it can be an invaluable guide as to where best to invest money in IT to get the greatest return for the business. In addition, utility computing is all about improving efficiency through best practice and automation. Again, storage is a great place to begin and putting in some best practices and simple automation – e.g., increasing space on servers when they are running out – can save a business a great deal of money, no matter what its size.

The grid is also seen as the next big thing and again, storage is a key component of a grid architecture. However, most grid applications need a large amount of space to store data centrally; that data is then farmed out and generally processed in memory within the grid so the actual storage requirements are virtually nonexistent on the fringe nodes. For the main central storage, ensuring that the application serving out the data is highly available and that the data is sufficiently protected, i.e., backed up or replicated, is generally adequate. Outside of storage, a general comparison of grid versus utility computing is interesting because, while both have very different applications running on the architecture and so from 30,000 feet look very different, from the ground level there are many similarities: what is being used, how much it is being used, and if it can be used more – to improve either efficiency and/or utilization. ■

**About the Author**

*Guy Bunker is chief scientist at Veritas.*

*guy.bunker@veritas.com*

This is the company that develops the technology,

that encrypts the data,

that's critical to your business,

so the right people get in,

and the wrong people don't.

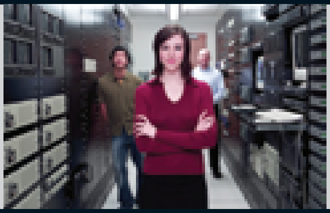

**Single-source security for the life of your information.**

The foundation of information security is encryption. Today, no one has more encryption experience and solutions than SafeNet. We protect and manage highly sensitive financial, medical, and government communications worldwide, even in the Oval Office. You have a choice. You can try to protect your information with a patchwork of hardware and software. Or you can get end-to-end security from a single source—SafeNet. To find out more, call today.

**Call 1-800-533-3958 to be SafeNet sure.**
www.safenet-inc.com

SafeNet®
*The Foundation of Information Security*

APPLICATIONS - AUTHENTICATION - REMOTE ACCESS - ANTI-PIRACY - LICENSE MANAGEMENT - VPN/SSL