Remote Data Protection and Disaster Recovery ⟨4

information STORAGE+ SECURITY journal

www.ISSJournal.com

In This Issue:

VOLUME: 2 ISSUE: 3 2005

IDENTITY MANAGEMENT

AS A REGULATORY COMPLIANCE ENABLER ⟨18

It's the cornerstone of most major regs

ISSJ Announcement Meet the New EiCs! ⟨3

RECLAIM YOUR EMAIL

Visit us at NECC
Booth #358
June 27-30 Pittsburgh, PA

Spam and virus protection at an affordable price.

- No per user license fees
- Prices starting at $1399
- Powerful, enterprise-class solution

Barracuda Spam Firewall

BARRACUDA NETWORKS

Order a free evaluation unit at
www.barracudanetworks.com

POWERFUL EASY TO USE AFFORDABLE

Aggressive Reseller Program
Get more info by visiting **www.barracudanetworks.com/NECC**
or by calling **1-888-ANTI-SPAM or 408-342-5400**

information STORAGE+ SECURITY journal

Worldwide Newsstand Distribution
Curtis Circulation Company, New Milford, NJ

For List Rental Information:
Kevin Collopy: 845 731-2684
kevin.collopy@edithroman.com
Frank Cipolla: 845 731-3832
frank.cipolla@epostdirect.com

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

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

From the Co-editors-in-Chief

A Quick Look at the Coming Year in Storage...

BY PATRICK HYNDS AND BRUCE BACKA

WE (PATRICK AND BRUCE) are new to ISSJ. In a future article, you'll learn more about who we are and why we are here. In this issue, we thought we would take a minute to frame some of the discussions that will go on throughout the year.

As you probably all know from first-hand experience, storage is the fastest growing and probably the most expensive resource on your network. With recent legislation concerning privacy, security, and financial reporting – hundreds of laws in all – storage must now be one of the most carefully managed resources on your network. You have two important reasons why you must do a good job: financial and legal.

While disk drives may be cheap, the cost of installing new storage and maintaining what you've got is not. In fact, industry analysts say you will spend three to five times your hardware acquisition costs maintaining that storage over its lifetime. We watch people put lots of effort into the economics of their purchases; we see much less effort put into cost of ownership.

To gain control of your operating costs, you have to gain control of what is going on with your information resources. This is only logical, right? How many of us have control – any amount of control?

If you had control, what would you do? You might classify the data and then provide different levels of service to different classes of information...critical data gets careful (read expensive) handling; less critical data gets less expensive treatment. Brilliant! (Obvious?) Can you do this with the infrastructure you have today? Are you positioned to do it tomorrow?

Industry analysts also tell us that much of the data on our networks is junk. (How much? Thirty to forty percent.) Do you have the infrastructure in place to sweep the junk out of the environment? Probably not, but it's worth a year and half of growth in terms of space utilization based on the fact that these same analysts tell us that the storage needs of most networks are growing at 18–25% a year.

As we go forward, we will tell you about some of the clever things people are doing to meet their legal obligations and reduce their operating costs. We will give you a sense of what is real and what is not. For example, the first thing you need to manage anything is control. If you don't have the policies and technology in place to control your storage use, complicated conversations about compliance aren't worth the time. You won't be able to do anything with the answer. If you are talking compliance and you don't have control – stop talking! Go get the technology that gives you control.

If you have control, you need classification. Not all information is the same. No one can afford the cost of treating all data as though it were mission critical. Classification is an emerging area for storage management technology. But let's take one of the myths off the table right now.

Retrospective classification – classifying the data that is already on your network – can only be done with the metadata attributes that are already there. If an attribute was not attached to this data when it was created, you can't afford the cost of figuring out what is missing. For those of you who are older, this harkens back to the old economics of system conversions. How much does it cost to move data from one system to another? Pretty much whatever it cost to create the data in the first place. Same deal with retrospective classification. How much does it cost to add classification attributes that are not there already? About as much as it cost to create the data in the first place. No one can afford this.

That's it for now. These are the issues we will be wrestling with in the coming months. Get on board! We hope you enjoy the ride as much as we will!

About the Editors

Patrick Hynds is the Microsoft Regional Director for Boston, the CTO of CriticalSites, and has been recognized as a leader in the technology field. An expert on Microsoft technology (with, at last count, 55 Microsoft certifications) he is experienced with other technologies as well (WebSphere, Sybase, Perl, Java, Unix, Netware, C++, etc.). A graduate of West Point and a Gulf War veteran, Patrick has experience in addressing business challenges with special emphasis on security issues involving leading-edge database, Web, and hardware systems. phynds@sys-con.com

Bruce Backa is the founder of NTP Software. He has acted as chief architect, technologist, and project manager for assignments involving large-scale technology and implementation strategies. Bruce has held the positions of director of technology and business research for the American Stock Exchange (AMEX) and director of technology for American International Group. He has also been responsible for the architecture, implementation, and management of a worldwide client/server networking infrastructure for a Fortune 10 company. bbacka@sys-con.com

Relieving the Pain

Automating remote data protection and disaster recovery

BY ROBERT FARKALY

AS ANY IT manager will tell you, ensuring the reliable backup and recovery of centralized data is difficult enough. Add the myriad challenges associated with backing up data from multiple remote locations and the complexity increases exponentially. That's why remote data backup and recovery has become the bane of many an IT manager's existence.

The problem of protecting remote site data is significant. Today's remote offices tend to have increasing amounts of critical data residing locally, whether it's customer databases, e-mail, applications, or financial information. In fact, according to a study done in 2004 by Strategic Research Corporation, 60% of the data resides outside the data center and as much as 75% of this data is unprotected and unmanaged by any IT staff. What's worse, Strategic Research Corporation also reports that up to 50% of remote backups fail.

Because there's more unprotected critical data residing outside the data center, companies are searching for solutions for backing up, recovering, consolidating, and managing distributed data. These solutions range from the traditional approach of using backup servers and tape drives at each location to new software-driven solutions that automate the process of backing up and consolidating remote data.

The Traditional Approach

Historically businesses that have attempted to protect their remote site data deployed backup servers, backup software, and tape drives or autoloaders at each location. On a daily or weekly basis, remote personnel changed the tapes and sent them to the data center or disaster recovery site.

Clearly this "good enough" approach has a number of drawbacks – oft times tapes are rotated, changed, and mailed to a central site by poorly trained non-technical staff for whom backup is an afterthought at best. As a result, backup policies may not be followed properly, tapes may not be rotated or handled correctly, and failed backups may not be noticed. In addition, tapes may be lost, misplaced, or damaged during shipment to the data center. Local recoveries are equally difficult, generally requiring the assistance of central IT staff.

The traditional approach also requires IT staff to periodically travel to each remote site for system maintenance, upgrades, and troubleshooting. This quickly becomes an expensive proposition for companies with dozens or hundreds of remote locations – and represents a significant drain on costly IT resources.

When Good Enough Isn't

Until recently many organizations managed to make do with the traditional approach to protecting remote data despite its many drawbacks. However, today's fast-changing business climate now dictates that businesses find a better, more reliable way to back up and recover remote data. In particular, many organizations must now comply with a host of new regulations for protecting mission-critical data. These regu-

ACHIEVE STORAGE CONSOLIDATION ACROSS THE ENTERPRISE. PUT AN END TO SERVER PROLIFERATION. SAVE WITH TACIT NETWORKS' WAFS SOLUTIONS.

From the Fortune 1000 to companies of all sizes, enterprises worldwide are joining the movement to Tacit Networks' Wide Area File Services (WAFS) solutions. They're solving their remote IT challenges…and slashing costs in the process…by:

- Eliminating file servers at remote offices
- Eliminating tape drives at remote offices
- Enabling true global storage consolidation
- Using existing storage resources far more efficiently
- Managing and backing up data for remote users at the data center
- Eliminating latency and duplicate files

STACK UP THE SERVICES. STACK UP THE PERFORMANCE. STACK UP THE SAVINGS.

On top of WAFS-based storage consolidation, Tacit Networks stacks an unparalleled suite of low-cost, datacenter-class, centrally managed IT services for remote offices, including:

- EMAIL SERVICES
- WEB CACHING SERVICES
- REMOTE MANAGEMENT SERVICES
- NETWORK SERVICES
- PRINT SERVICES
- FILE SERVICES

Extending IT services to the branch office

THE BOTTOM LINE? YOUR BOTTOM LINE.

Make the move to Tacit Networks and *consolidate* storage across the enterprise. *Drive* stronger information flow throughout the enterprise. *Eliminate* remote office IT infrastructure. And *save* every step of the way with ROI in nine months or less.

Calculate your ROI. Visit our new WAFS ROI calculator at www.tacitnetworks.com/ROI, or call 888-757-TACIT.

TACIT™ NETWORKS

IMPROVE YOUR NET WORK.™

lations, such as Sarbanes-Oxley, govern a number of data protection factors, including how long data must be stored, how quickly it must be recovered, and how many miles apart a central and mirror site must be for effective disaster recovery. These regulations also require that businesses have full audit accountability for IT systems and processes for data protection and availability.

As a result, businesses with distributed data are now examining their backup and recovery practices to identify areas that are potentially non-compliant. These internal audits often uncover significant gaps in their data protection strategies.

For example, a large meat packing company in Canada has a backup server and tape drive at each of its 150 locations. Every night an employee at each location leaves the processing line, removes his gloves and smock, changes the tape, packages up the used tape for shipment, and returns to the line to complete the shift. As one might imagine, this manual process proved highly unreliable. Employees often followed improper tape rotation procedures or forgot to change the tapes altogether; tapes were prone to dirt and contamination; and shipments were lost or damaged. As a result of these factors, the company's auditors now require that it implement a more reliable backup and recovery solution for its distributed data.

Besides the new regulations, other issues in remote site backup and recovery include meeting internal service-level agreements, protecting current investment in IT assets, resources and improving ROI, and providing greater control over distributed information. Together these factors are driving change in how organizations protect, consolidate, and manage their remote site data.

Replication Software: Right Technology, Wrong Problem?

As part of their quest to improve remote site backup and recovery, many organizations are evaluating replication technology. Replication enables a second storage unit to duplicate the data on the original storage unit. Originally designed for duplicating storage in the data center, replication has recently gained momentum as a means of duplicating remote site data, given the relative affordability of the high-bandwidth connections that replication solutions require.

Remote site data replication entails two separate steps – duplicating data from the remote site to the central site, then consolidating and backing up the entire data set centrally. While this approach improves the reliability of backing up branch office data, it introduces a number of additional challenges that businesses must consider before implementing a replication-based solution.

Replication solutions require companies to buy, install, administer, and maintain replication servers and software licenses for each remote location. This not only represents a significant hard cost, particularly for companies with a large number of remote sites, but also necessitates considerable IT expertise to deploy and manage the solution. In addition, customers with heterogeneous environments at their remote locations may face major management or standardization issues as they try to implement an organization-wide solution. They may need to rewrite backup scripts, buy new hardware and otherwise disrupt their existing backup infrastructure.

Another important consideration is the actual location of the backup data. Under a replication scenario, backup data exists only at the central site, which means restoring a local file will typically require the intervention of the central IT staff. What's more the ability to restore a file at all is dependent on the network connection between the central location and the remote site. If the network goes down, the remote site is cut off from its backup data until the connection is restored.

In general, these and other challenges related to using replication for remote site backup stem from the fact that replication software wasn't designed for this purpose. As such, it has inherent limitations that are spurring a growing number of companies to seek solutions that are purpose-built for remote site backup and recovery.

A Better Solution: Software-Enhanced Disk-to-Disk Appliances

In recent months, Overland Storage has introduced a new solution that relieves the pain of consolidating and backing up remote site data. The foundation is a disk-to-disk appliance called the REO that's optimized for backup and recovery. On top of this appliance specialized software called Multi-SitePAC runs that automates the process of protecting remote site data.

Here's how it works. At each remote location, the backup software that's already in place transfers data to the appliance in virtual tape format as if it were still backing up to a tape drive or autoloader. The remote location appliances are linked via iSCSI to a central site appliance. At a user-determined time, the remote data is automatically mirrored to the central site appliance where it can be easily consolidated, managed, and even archived to tape if desired. The entire process is controlled from the central site.

There are a number of benefits to this approach not the least of which is the vastly improved reliability of remote site backups. By automating the process of protecting remote site data, REO with Multi-SitePAC eliminates the potential for human error at branch offices. It also eliminates media and shipping costs – a significant savings, particularly for businesses with many remote locations.

Unlike replication solutions, customers using REO with Multi-SitePAC don't need to purchase, install, and manage additional servers or software. Instead, the platform- and backup software-agnostic appliance fits seamlessly into current environments. No changes to existing backup scripts, systems, or software are required.

Another important benefit is the ease and flexibility of data recovery. Because a copy of the data resides locally as well as centrally, files can be recovered from either location – no matter what the status of the network. In addition, data can often be restored locally without assistance from the central IT staff.

The Bottom Line

Today's distributed businesses are struggling to comply with new regulations, fulfill service-level agreements, ensure business continuity, and reduce storage management costs across their operations. By leveraging new technologies like the REO with Multi-SitePAC, these organizations can deliver on these business imperatives while relieving the pain of protecting remote site data.

About the Author

Robert Farkaly is director of disk-based products at Overland Storage. He has more than 25 years of information technology sales, marketing, and business leadership experience at both start-ups and Fortune 100 companies. Bob is a founding member of SNIA, creator of the SAN Appliance, D2D2T, and backup acceleration appliance market categories.


Is your network
TENABLE?
What happens between the last time a network vulnerability scan is completed and the next? New hosts, new intruders, new ports and new vulnerabilities arrive continuously. Your efforts to defeat them must be continuous as well.
Detect and verify intrusion attempts and vulnerabilities without active scanning. NeVO from Tenable keeps 24/7 watch through a passive monitoring system that helps to ensure comprehensive security with zero impact to your network.
Available for Windows or UNIX. With NeVO, install once and receive continuous vulnerability monitoring.
TENABLE Network Security
www.tenablesecurity.com
(877) 448-0489
TENABLE
Network Security
NeVO

Regulatory Compliance in Complex Heterogeneous Environments

THE ANSWER RESTS IN EXTENDING MICROSOFT TECHNOLOGY

BY MATT PETERSON

IN RECENT YEARS the regulatory pressure on organizations to secure, document, and protect their data and systems has become increasingly difficult to ignore. There appears to be no lack of government regulations — both in the U.S. and abroad — that impose new laws requiring corporate accountability for controls placed on information and technology. Typically organizations implement controls by adding or replacing technology, processes, and staff. While the scope of the regulations reaches beyond Information Technology (IT) controls — covering many aspects of an organization's operations — IT departments seem to bear the brunt of the responsibility.

So what is causing this sudden onslaught of regulations? Often they are the result of public complaints of unacceptable business practices. While a review of the actual language of the regulations may be intimidating, the documentation essentially amounts to requiring long overdue "best practices" such as protecting confidential patient data as mandated by the Healthcare Information Portability and Accountability Act (HIPAA) or ensuring the financial integrity of earnings reported by public companies as in the Sarbanes-Oxley Act (SOX).

While some corporate officials may loathe them, the internal controls mandated by recent industry regulations aren't considered to be wholly unnecessary by all executives. On the contrary, protecting the individual privacy of consumers and preserving data integrity are at the forefront of most companies' IT strategies. For the vast majority of companies, the adoption of these newly mandated policies and practices are just part of a security update plan that makes good business sense. They focus on these issues not only to comply with new regulations but because they see that it helps them serve customers better,

generate revenue, and hopefully turn a profit. At other organizations keeping business practices up-to-date with technology can, even with the best intentions, fall short of creating the level of accountability and security required — it's these companies that the regulations are specifically aimed at. However all companies, those with best practices and those without, are still equally accountable under these regulations.

Regardless of the motivation behind a given set of regulations, they generally require organizations to secure data, ensure the integrity of information, protect the privacy of individuals (employees, customers, clients, and partners), and preserve the availability of information for appropriate parties. From an IT perspective all regulations can be boiled down to three main strategies:

- Ensure that data is protected from unauthorized access (either from within or without an organization)
- Ensure that information is accurate (has integrity) and is available to those who are authorized to access it
- Ensure that systems and processes are in place to satisfy the first two

Alphabet Soup

HIPAA, SOX, GLB, and other sets of regulatory governmental enactments can be difficult to digest and even more difficult to satisfy. In order for organizations to successfully comply with the myriad regulations they face, an understanding of the general requirements, penalties, intentions, and motivation for each regulation is useful.

Gramm-Leach Bliley Act (GLB)

Title V of the Gramm-Leach-Bliley (GLB) Act requires financial institutions to ensure the security, confidentiality, integrity, and protection of customer information. Boiled down to its very core, GLB Title V means that a financial institution must protect the customer information it holds from unauthorized access by those outside of the institution and must inform customers how personal information is used by the institution.

Technology solutions to aid in GLB compliance center on access control, identity and authentication management, and data security.

Health Care Information Portability and Accountability Act (HIPAA)

HIPAA is very similar to the privacy provisions of GLB except that it's focused on the healthcare industry. Under HIPAA, organizations that generate, maintain, or distribute a patient's personal healthcare information must ensure that that information is secure and private.

As with GLB, IT departments in organizations covered by HIPAA center their efforts on data security, access control, and identity and authentication.

Sarbanes-Oxley Act (SOX)

Sarbanes-Oxley was passed by the U.S. Congress in 2002 in direct response to the

corporate financial scandals of the time. SOX affects companies (both domestic and international) that have to file with SEC. The act contains a number of requirements centered on financial reporting and controls aimed at protecting investors by improving the accuracy and reliability of corporate disclosures to the SEC.

AMR Research estimates that U.S. companies spent more than $1 billion on technology in 2004 specifically to address SOX. But Gartner predicts that 80% of those technology solutions will be replaced by 2005 as companies improve their compliance and move from tactical to strategic initiatives.

Technology initiatives to address Sarbanes-Oxley should include authentication and password management to "establish and maintain an adequate internal control structure." Generally these efforts aim to raise the security surrounding data access for all systems in an enterprise.

Title 21 Code of Federal Regulations (21 CFR Part 11 FDA)

21 CFR Part 11 is legislation introduced by the Food and Drug Administration (FDA) that allows the use of electronic signatures, electronic records, and handwritten signatures on electronic records in lieu of handwritten signatures on paper in certain circumstances in the pharmaceutical industry. It includes directions on limiting system access to authorized individuals, the use of authority checks to ensure that only authorized individuals can access a system, and the adequacy of the documentation of system operations and maintenance.

Generally, organizations affected by 21 CFR Part 11 will need to include as part of their compliance remediation efforts a focus on technology initiatives tools to help manage the authentication and identity management of users and systems.

Compliance

Each of these regulations requires the protection of data and systems. Several emphasize the privacy aspect of personal information. Most demand security surrounding the transmission of data between organizational units in a company, between a company and its partners, and/or between a company and its customers. Governments have taken best practices beyond simple recommendations to the level of mandated requirements with specific and significant penalties assessed for violations.

A vast majority of companies fall under one or more of these regulations. While many companies undertook best efforts to ensure privacy and security prior to government regulation, almost without fail, recent regulations have demanded that they take another look at these issues. Responses may range from simply documenting current practices to a comprehensive overhaul of all systems and management operations. Most, probably fall somewhere in between.

For issues of access control, SOX establishes management's responsibility of "establishing and maintaining an adequate internal control structure." For heterogeneous enterprises legacy Unix systems may rely on NIS-based authentication and password synchronization. Typically, when these types of legacy implementations are discovered by SOX audits, they're likely to fall short of the requirements. Both "significant deficiencies" and "material weaknesses" in a system's access control must be reported to the SEC.

The same kind of shortcomings face virtually any multi-platform enterprise striving to comply with GLB, HIPAA, SOX, 21 CFR Part 11, or other regulations. Often these organizations have implemented Microsoft tools and technologies — such as Active Directory or SMS — which include features that address compliance in a Windows environment. Unfortunately creating the same level of compliance for Unix, Linux, Java, or Mac systems is where the greatest challenge lies. These Windows tools don't extend to non-Windows systems so organizations must turn elsewhere to address compliance.

Regulation	Industry	Security	Privacy	Transmission
GLB	Financial	✓	✓	✓
GLB	Healthcare	✓	✓	
SOX	Publicly Traded Companies	✓		
21 CFR Part 11	Pharmaceutical	✓		✓

Summary of Regulations

When compliance initiatives are extended to the entire enterprise, something as simple as password maintenance can require significant investment in additional technology, infrastructure, staff, and processes. Each of these business areas provide another potential "reportable condition" to be discovered in compliance audits. Multi-platform password synchronization solutions or meta-directory solutions are notoriously difficult to manage, require additional infrastructure solely for the purpose of making the solution work, and ultimately drive up IT lifecycle costs while potentially falling short of regulatory compliance.

Take a real-world situation at a large U.S.-based company. At this company the process of de-provisioning an employee (eliminating system access and terminating user rights) for a dominant Windows environment and distributed Unix/Linux environment simply couldn't be done realistically with its existing infrastructure. The Windows de-provisioning process was very simple, straightforward, and compliant due to the effective use of Microsoft Active Directory. Unfortunately, de-provisioning the same employee on the Unix systems required a number of manual "visits" to the Unix servers, which pulled important and highly compensated Unix support staff away from their core responsibilities to focus on tasks that were handled more efficiently by the Windows help desk.

As a consequence, Unix de-provisioning rarely happened in a timely manner (and sometimes not at all) producing significant reportable violations of the regulations imposed on this company. Moreover, the remediation options that were researched looked impractical. All efforts to procure an appropriate cross-platform identity and access management solution proved too cumbersome, too error-prone, or too expensive. The company concluded that developing a solution internally was too expensive. Fortunately, this company did light on an elegant and simple solution (keep reading).

In a nutshell the more complex the solution the more likely it is to be non-compliant. This doesn't even take into account the almost guaranteed increase in overhead expenses as more time, infrastructure, staff, and processes must be implemented to make the solution work.

The same thinking can be applied to other areas of compliance such as data secu-

rity, systems security, Web-based access control, and even security patch management. In all cases, Microsoft tools and technologies provide a highly compliant self-contained solution. However, as soon as non-Windows systems are introduced to the mix, the complexity, cost, and potential for regulatory violation increases exponentially.

Extending Windows Tools to the Entire Enterprise

Concerning access control, 21 CFR Part 11 gives some specific guidance surrounding the procedures and controls for limiting access to only authorized individuals. These guidelines include:

1. Maintaining the uniqueness of each combined identification code (user ID) and password so that no two individuals have the same combination of ID and password
2. Ensuring that user ID and password issuances are periodically checked, recalled, or revised (for example to cover such events as password aging)
3. Use of transaction safeguards to prevent unauthorized use of passwords and/or identification codes

For each of these guidelines, Microsoft Active Directory (AD) provides an ideal platform for ensuring compliance for Windows resources. AD's ability at each is listed below:

1. Active Directory supports standard character-type requirements as well as third-party "passflt" add-ons that enforce complex passwords for rigorous password policies
2. Active Directory includes password reuse restrictions, minimum and maximum password aging, and password lockout features (disable for x time after x failed attempts)
3. Active Directory is a powerful implementation of the Kerberos and LDAP encryption standards creating transaction and session layer safeguards

While 21 CFR Part 11 is simply used as an example here, all other regulations demand roughly the same level of security. Active Directory is an ideal tool for creating authentication and access security in Windows environments.

At the real-world company mentioned above, a Windows consolidation project on Active Directory convinced IT management of the value of AD in creating security and compliance. In addition, the consolidation saved the company millions of dollars annually in password maintenance and support. A desire to extend those benefits to the company's Unix and Linux systems initiated a comprehensive search for password synchronization solutions — none of which delivered the desired results. This was followed by an internal development project that attempted to create a Kerberos and LDAP extension of Active Directory to the Unix and Linux environments. The company quickly realized that internal development would be extremely expensive and not provide the support, update, or stability of a similar commercial solution.

The company discovered a compliance solution that literally extends the scope of Active Directory to include Unix and Linux resources. The particular solution they chose is called Vintela Authentication Services (VAS).

This solution proved uniquely valuable in creating compliance because it extends the compliance-enabling technologies offered by Microsoft Active Directory to non-Windows platforms. A solution that was proven to help deliver compliance in the Windows world can now be leveraged for non-Windows systems. The advantage is a single point of administration, no additional infrastructure, consolidation of management tasks, and reduced operational costs.

Each of the 21 CFR Part 11 requirements listed above, which are addressed by Windows systems through Active Directory, can now be seamlessly and natively extended to the rest of the enterprise. In a SOX audit, consolidating all systems into AD eliminates the potential deficiencies of NIS-based Unix authentication and password synchronization scripts. The same can be said for HIPAA, GLB, EUDPD, and others.

Available Solutions

Many companies have adopted Vintela solutions to help establish regulatory compliance. These solutions functionally extend Microsoft management and infrastructure technologies and products to the non-Windows world. It's all made possible through the native and specialized implementation of standards on each non-Windows system. This strategy extends Microsoft management and infrastructure technologies beyond Windows. With VAS, Kerberos and LDAP are uniquely and specifically applied at the Unix/Linux OS-level. This integration with the native NSS and PAM standards creates a seamless experience allowing Active Directory to act as the single sign-on environment for all systems while still maintaining the individual personality of the Unix or Linux systems. All systems become full members in the AD domain.

Conclusion

The recent proliferation of government regulations has put an added burden on IT departments as they struggle to satisfy regulations while controlling costs and management overhead. These regulations generally demand accountability, security, privacy, and documentation for data and system access as well as data and system integrity. With significant penalties for reg violation, the desire for companies to comply can come from the very top. Unfortunately in all but the rare case, companies must evaluate current practices and implement new technologies or strategies to become compliant. These shortcomings become even more apparent when compliance must apply to heterogeneous enterprises. Most available cross-platform solutions are:

- Complex
- Expensive
- Difficult to manage
- Require an additional layer of infrastructure, which must also be maintained and managed
- And rarely deliver the level of compliance required by the regulation

Microsoft offers a comprehensive set of tools and technologies proven to satisfy compliance issues for Windows networks. Unfortunately when non-Windows systems must also become compliant, organizations traditionally explore additional tools that in turn introduce additional complexity, cost, and potential violations.

Available solutions — such as those from Vintela — extend the scope and capabilities of Microsoft technologies to non-Windows environments. Because of these products, a company's compliant Active Directory implementation can be extended to Unix, Linux, and Java systems immediately bringing the entire enterprise into line with specific regulations. Similarly a compliant Windows management strategy using SMS can be extended to Unix, Linux, and Mac resources.

About the Author

Matt Peterson is chief technology officer of Vintela.

mpeterson@vintela.com

INFITECH

X5 NAS

empower your data network

High Performance
Rack Mount Servers and Storage Solutions

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

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

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

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

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

Powered by NetEngine

Visit Us www.infi-tech.com

or Call 1-800-560-6550

to Find Out More

Weathering the Storm of IT Security Compliance

IT'S 90% PROCESS AND 10% TECHNOLOGY

BY DREW WILLIAMS

IN BUSINESSES THROUGHOUT Europe and the United States, the segregation of IT security and system operations has become entrenched. Further confounding the rift is the pursuit of all things "compliance" (e.g., BS7799, ISO 17799, BASEL II, etc.). Industry analysts and vendors alike anticipate an extension of the compliance movement that focuses on the actual IT audit, which may further confound efforts to reunite IT operations under a common banner. As anxiety heightens over when the next "Big Problem" will hit the Internet, there are some things that systems administrator and C-level executives can do to fortify their IT business processes against that unseen storm that's looming just over the horizon.

Facing the reality that all Internet-connected systems are doorways of risk is not easy for IT administrators. But since more than 90% of all security risks exploit known system vulnerabilities according to Gartner, the controversy of "where to react" transforms into one of "failure to plan." Add to this the fact that organizations can no longer hide behind the "we didn't know what was happening" defense, and matters concerning "security risk management" become issues of "business contingency planning and accountability."

Umbrellas of Compliance

In recent years, many organizations have felt the heavy hand of standards and compliance knocking on their door — especially government agencies and the banking community. For American-based companies, much of the compliance push comes from the vague and elusive Sarbanes-Oxley (SOX) rules for security risk management and accounting. During 2005, while SOX continues to stand at the center of the compliance controversy — with its reach extending into European markets as a new potential benchmark — other frameworks and methodologies, such as ITIL and COSO/CobIT — along with ISO-based standards — are beginning to thunder through the world's business communities.

But what of the hype that surrounds all of these issues of compliance? The seasoned IT manager has heard this rumbling before — in the recent winds of the Y2K storm that passed by a half-decade ago.

Compliance standards are reaction-based initiatives. These new and often ambiguous standards further the confusion IT administrators and their bosses are forced to face as fears of penalties and possible prison time threaten to strike at will. And unfortunately, the IT security vendors are all too well-aware that buzzwords like compliance mean good business on which hundreds of IT security vendors build their marketing models.

Preparing for Foul Weather

Focusing on continued efforts to defend their expensive mission-critical infrastructures from the frequent storms of attacks and exploits, IT administrators are also frequently forced to decide which vendor's story about security makes the most sense (or cause the least amount of confusion). Determining which tools make the right sense to address security risks, while trying to maintain current operational standards of performance puts even more pressure on administrators. "Which anti-virus will best defend my system?" "Will these policy and assessment applications scale to my enterprise?" "Do these free spyware tools really work?" And "What do 'intrusion prevention' tools really prevent?" are all common questions for the bewildered sys admin.

So, which tools make the most sense? How much "security technology" do you really need? And where and when does the "prevention" actually begin?

IT administrators have raised time and again the fact that their concerns aren't necessarily about the rules themselves — rather, they are concerned with what further risks they might be facing by overlooking something while rapidly moving to meet compliance deadlines, or while reacting to specific incidents or reports of attacks.

That said, the following are three basic principles that systems administrators might find helpful when trying to break through the clouds:

1. **Compliance is 90% process and 10% technology.**

   Part of "process" is gaining a full understanding of what's happening "behind the scenes" before beginning to define any sort of policy, or react to any type of mandate.

   While there's a lot written about "intrusion prevention" (IPS) technology, in most cases an incident actually has to occur, or a violation of the defined policy must be recorded before tools claiming to be IPS become active. Realistically, even the "IPS" methodology is more reaction-oriented than preventive.

2. **Defining an operational policy without first assessing the environment to which it is assigned is too late.**

   More than 800 vendors are vying for one's IT security business. Most of them begin their security lifecycle models at the policy and move forward with varying

degrees of success to defend some portion of that policy (assessment, event logging, perimeter defense, etc.). However, since these security policies are often segregated from the rest of the operational controls (i.e., a separate policy for everything else), most times the general market still looks at IT security tools as a way to react to a fraction of a bigger problem (such as a virus outbreak, the threat of denial of service, etc.).

Administrators may find it easier to manage and enforce a policy after first learning as much as they can about their environment, its settings, and what is necessary to optimize that environment. In this case, knowledge before taking action is key in determining which decisions will have the best results. Administrators will find that gaining a better understanding of their environments will greatly simplify the need to react to a mandate or some other external control.

3. **More than 90% of all the exploited vulnerabilities are based on known problems and poorly configured environments.**

In Las Vegas, those odds would make millionaires out of the homeless. When navigating through rough waters and high seas seafarers know that survival depends on maintaining a true course while ensuring watertight integrity throughout their infrastructure. Knowing that there's a nine-to-one ratio of where a problem is going to occur (and often with a three- to five-month lead time) plus the capability of gathering thousands of data points about an infrastructure's most intimate configuration settings moves the concept of "risk prevention" to the level of "security empowerment."

Following a more administrative approach to addressing potential risks, systems administrators should consider a configuration management database or CMDB-driven data repository as the starting point. Administrators could actually prevent most of the risks to their IT infrastructures by gaining a complete understanding of details associated with system settings and configuration controls at all points throughout the enterprise. Defining the policy on which an organization builds a "gold standard" of operation without this critical step results in an ineffective reactionary-based trend in enterprise IT security.

Over the Rainbow

Once administrators have collected that mission-critical data, they can begin to shape an appropriate policy for what should be considered the "gold standard" of operational expectation. Blending the strong integrity of a CMDB-based approach to policy management further capitalizes on the administrator's ability to address the need for pre-emptive control rather than post-event recovery. In a sense, you can't fix what you don't know is broken, but you CAN plan for risks when you know what you have and how it's working before those risks are exploited.

The old axiom that "knowing is half the battle" certainly rings true where your organization's risk management plans are concerned. Organizations can no longer afford to claim "The hole is on your side of the boat."

About the Author

Drew Williams, a long-time information management and security strategist, brings deep industry experience in corporate and product marketing and business development to Configuresoft. He pioneered the vendor security research team model with the industry's first such group, AXENT Technologies' Information Security SWAT Team. Williams was also a founding member of the President's Partnership for Critical Infrastructure Security, a member of the Internet Engineering Task Force on Internet Security, and an initial member of the independently supported CVE development team. He has served as a security policy advisor to major financial institutions, healthcare manufacturers and state governments.

drew.williams@configuresoft.com

Assuring Compliance with Content Security

LESSONS FROM THE TRENCHES

BY KIMBER SPRADLIN AND SKIP DOSTINE

REGULATIONS AND AUDITS have become a way of life for many security officers, especially those in the financial and healthcare industries. For example, the Gramm-Leach-Bliley Act (GLBA) requires banks and financial institutions to establish comprehensive security policies to safeguard customer data. Likewise, the Sarbanes-Oxley Act of 2002 requires all publicly held companies to establish and maintain internal controls over their financial reporting systems and ensure their effectiveness.

At the time these regulations were drafted, however, their far-reaching consequences weren't understood. The cost of non-compliance can be high ranging from a loss of company reputation, to prohibitive fines, to imprisonment. For employees, expectations of privacy have been forever altered. And yet, most compliance regulations don't provide the specifics needed to translate broad security mandates into day-to-day guidelines and procedures. If there was ever a formula for a headache, this was it. As a result, companies are only now coming to grips with what compliance means to their organizations. As IT makes its way forward in this wilderness, lessons have been learned. Here are some of them.

Follow the Framework

What's become clear is that when it comes to compliance, legislators are better at spelling out the end point they want to arrive at, rather than the road to get there. In particular, none of the information security or privacy regulations provide more than the broadest of guidelines when it comes to the nitty-gritty of drafting effective acceptable use policies (AUPs) in the organization. Indeed, while Congress may have set down the broad requirements, it will be the courts that, in

the end, determine how the law applies to real companies under real circumstances. What does that mean to you? In essence, you must connect the dots yourself, tailoring the intent of the regulations to your specific business or industry, as well as special constraints and considerations and other facts of operational life. In doing so, you may find that you are describing policies that are specific not just to your industry, but your company. Two companies in the same industry with similar organizations may, in fact, have quite different policies depending on slight variations in how they do business. The best advice we can give you is to use externally validated frameworks as your guideposts, including ISO 17799, COBIT, which is most commonly used in relation to Sarbanes-Oxley, and the guidelines coming out of the National Institute for Standards and Technologies.

The E-Mail Retention Balancing Act

Some of the most eye-grabbing tales about compliance have to do with e-mail — how long to keep it, whether to filter it, what to do if it's subpoenaed. In some cases, government regulations mandate what needs to be retained and for how long — particularly in the financial and healthcare sectors. Otherwise, the choice is up to you, and, for most companies, the emerging conventional wisdom is to retain as long as necessary and not a day longer. This isn't about burying potential evidence; there are valid legal reasons that once a piece of correspondence isn't absolutely required, you should get rid of it. Not doing so has caused many companies, including Microsoft, major financial burdens.

And then there's the punishing cost of pulling the relevant documents up from a massive offline archive of correspondence. The number of e-mail messages generated in a week by even a mid-size company can number in the millions, and the cost of retrieval rises exponentially with the number of years retained. If you don't have the tools in place ahead of time, the cost of doing so in time to meet court-mandated deadlines can be extraordinarily expensive.

As a result, some companies are starting to filter correspondence up-front if possible — determining message content that in turn determines longevity at the time an e-mail is sent and received. For example, correspondence related to patient care may be retained for the life of the patient — if it has to do with diagnosis. Or it might be retained for a much shorter period — if it's related to billing. We may eventually see in the U.S. what some European countries already permit: an "opt-out" policy in which employees can mark a given piece of correspondence as personal, not business-related. The e-mail goes out and — not being business-related — is classified as such. European

countries also tend to present more real-time policy reminders to employees when an activity is performed. In some environments, for example, each time a staffer sends an e-mail, a prompt message comes up as a reminder that the system is meant for business use only. Those reminders are also another way of demonstrating that a company is doing the right thing — proving that an individual knowingly violated corporate policy.

Whatever policy you set, remember that investigating agencies make a distinction between your written AUP and the de facto policy you actually follow. If, for example, you say that you retain e-mail for a year but your archive extends to three, the de facto retention policy is the one that may apply.

Webmail, IM'ing, and Webcams

During the early stages of regulatory compliance, companies turned their full attention to e-mail as the communications link between their internal staff and the outside world. In practice, however, the picture is a bit more complicated. When we first install our e-mail security tool at a customer site, employees who want to communicate privately switch over to a Web mail account such as Hotmail, Yahoo! Mail, or Google's Gmail. The shift is both immediate and predictable. And when those venues are covered, resourceful employees shift again — to instant messaging.

As a result, companies are now planning from the get-go to monitor traffic on their corporate e-mail accounts, over the full gamut of Web mail services, and on messaging services as well. The scope of scrutiny is wider, but the method is the same. Good compliance tools will look for key words and phrases that could signal trouble, sometimes using standardized templates that attempt to recognize a type of activity an individual is undertaking. Sometimes information in combination can raise a red flag. A medical group exchanging medical terms might not raise suspicion, but that combined with a customer ID number or a social security number may raise a red flag.

A good tool will also let management put restrictions on the kinds of files that can be uploaded and downloaded and consider the content of those as well. In the long run, the answer will be one of employee expectations. The new compliance rules all but mandate that there's no such thing as truly private correspondence on the corporate network. If you have something to say and don't want others to know what you are saying, say it somewhere else. All electronic communications leaving the company network should be viewed the same way. The acid test is would you say what you're communicating via e-mail, the Web, or IM if it was typed in a letter on company letterhead? The courts will certainly view it as having the same weight and remember that it's just as permanent – just because you delete something from an archive don't assume that the communications trail has disappeared. If it was sent to one or many others what was their retention policy?

Employees aren't only clever at figuring out new methods for private communication, but also at figuring out more creative ways to do it. Many companies, for example, have figured out that cell phone cameras are a potential security risk and have banned them from the premises. And yet, the combination of instant messaging and low-cost Web cameras is every bit as lethal, but many AUPs have yet to catch up.

Web Browsing – Get Real

Companies have discovered that just as too lenient an AUP can lead to trouble so can a policy that's too rigid — because it can't be enforced. In most situations, for example, it's simply not realistic to ban any correspondence that is not 100% business-related. Human nature being what it is, even crusty security guys can spend a few minutes browsing ESPN.com. Here's the trap: once you have known violations, if you don't prosecute, then the policy becomes null and void, giving you no legal standing to enforce it. The better course is to create a policy that seeks a realistic balance. For example, you might specify that personnel are allowed to use the Internet for personal use (within the bounds of a company's anti-harassment policy) six hours a month, or only during a lunch break, or for 10% of their time. That gives employees the ability to check their bank accounts and eBay bids, and gives your AUP the flexibility it needs to pass muster.

You may find it also makes sense to have different AUPs, depending on the circumstance. To cite one extreme example, one customer, an energy company, had an AUP for the crew on an oilrig that essentially said: "Anything goes." The isolated environment and lengthy stays justify what, in a different setting, would be an irresponsible AUP. But note that when drilling crews return to the mainland, that policy stays back on the rig. While this approach may make logical sense, without automated tools to assist you in enforciong these different AUPs, putting this approach into practice is nearly impossible.

An Ongoing Process

It might sound like a cliché, but compliance turns out to be a process, not a goal. One of the biggest issues facing IT is convincing management to fund additional compliance projects, as well as maintaining the existing ones. When the regulations were first introduced, their visibility in the press, particularly with Sarbanes-Oxley, alerted executive teams to the need for funding. What IT departments are now discovering is that sustainable funding for ongoing compliance is much more difficult to secure. In some companies, the finance department expected that the budget would return to prior levels when in fact, compliance is an ongoing, never ending process. The people in the trenches know that, and the challenge is in communicating that message above.

Compliance regulations are here to stay. They will be tested in the field, refined by the courts, and, no doubt, augmented by further legislation down the line. The biggest lesson learned is one any Boy Scout can relate to: be prepared.

About the Authors

Kimber Spradlin is a senior compliance architect at NetIQ corporation with eight years of experience in the information security field. She is a security subject matter expert currently focusing on understanding the needs of, and communicating with, the regulatory and policy compliance market.
kimber.spradlin@netiq.com

Skip Dostine is the product marketing manager for NetIQ's Marshal Content Security Solutions. With more than 25 years of international technology experience, Skip's background includes sales, product planning, project management and engineering, as well marketing and operations.
skip.dostine@netiq.com

Do Not Pass Go!

JUST PROCEED DIRECTLY TO JAIL

BY WINN SCHWARTAU

I'M GOING TO make two predictions.

One: Every single American will have his identity stolen in the next five years.

Two: Some of the management folks who read ISSJ will go to jail in the next five years for poor security practices.

OK, time to explain. In the last year or so, unless you are dead, you've seen the headlines about countless private databases that have been compromised by criminal hackers, insiders, lost or stolen computers, misplaced tapes, and other abuses of private data. The numbers are staggering.

In 2004, there were 9.3 million cases of identity theft – and those are just the ones that were reported! Heaven knows how many other clueless Americans are wandering the streets and malls with debt loads that are ballooning because of organized crime. Guess: 50 million or more? Whatever. Too many.

This year is shaping up to be a banner year for stolen IDs. Lexis Nexis: 49 known and reported hacks. BofA. Credit unions. Medical databases. HIV lists. Voter registrations. Manufacturers. Retailers. It is truly a sad state of affairs that so much data is being negligently released to the bad guys when we, the security folks, have offered the solution to the database folks for decades.

Yet your management refuses to implement the solution, which accounts for prediction number two: Someone is going to jail for malfeasance, and that will be the managers at companies who have made a conscious decision NOT to build the appropriate security controls into their databases and mass storage media. (Are you looking over your shoulder?)

The solutions are trivial and based on the most basic principles of information security developed more than three decades ago: confidentiality, integrity, and availability. Organizations do a fine job investing in infrastructure, redundant infrastructure, real-time mirroring, fault tolerance, contingency planning, and business continuity to make sure that the third leg of the security triad, availability, is as close to perfect as can be. Why?

Business is business, and the loss of availability means a loss of revenue and a loss of revenue is the worst thing that can happen to most companies, so they spend a lot of money making sure the data is available and the systems are working. Good job guys. But, you missed the point and now meet your cellmate, Bubba.

Where are the confidentiality (keeping secrets a secret) and integrity components to protect the data itself and the gazillions of people whose records you are trusted to safeguard? It ain't there, Jack.

Jeeez, and it's so simple conceptually, and, yes, a bit harder to make work in far-flung enterprises. But, as any security expert witness in a court of law will swear, "It's been well within the capabilities of the database and storage industry to enforce both confidentiality and integrity for almost 30 years."

The simple answer is cryptography. Cryptography is the security tool that will solve both the confidentiality and integrity problems every storage system faces.

- If the data on the chosen storage media are encrypted, the bad guys would have to steal both the machines, the tapes, or the data AND have the appropriate key to decode the data.
- By encoding the data with an additional hash, the reliability and accuracy of the data (we call this integrity) can be maintained. Banks have been doing this since the appropriate cryptographic tools were made a National Standard in 1976. Yes, Virginia, there is a Standards Clause!

Admittedly there are no plug'n'play solutions for the enterprise, but many

Hashing Algorithms

The key in public-key encryption is based on a hash value. This is a value that is computed from a base input number using a hashing algorithm. Essentially, the hash value is a summary of the original value. The important thing about a hash value is that it is nearly impossible to derive the original input number without knowing the data used to create the hash value. Here's a simple example:

Input number: 10,667
Hashing algorithm: Input # x 143
Hash value: 1,525,381

You can see how hard it would be to determine that the value 1,525,381 came from the multiplication of 10,667 and 143. But if you knew that the multiplier was 143, then it would be very easy to calculate the value 10,667. Public-key encryption is actually much more complex than this example, but that is the basic idea.

Public keys generally use complex algorithms and very large hash values for encrypting, including 40-bit or even 128-bit numbers. A 128-bit number has a possible 2^{128} or 3,402,823,669,209,384,634,633,746,074,300,000,000,000,000,000,000,000,000,000,000 different combinations! This would be like trying to find one particular grain of sand in the Sahara Desert!

enterprise management tools offer crypto as an option. I bet you didn't buy it, eh? Consider how you can dovetail these critical security services into your existing architectures:

- Employ your existing PKI backbone with crypto services not just for transmission but for storage as well.
- Have X.509 management somewhere? Tie it into your cryptographic services.
- Add a couple of crypto experts to manage the whole thing. Trust me, you need them and it's worth the expense.
- Still using passwords? Well, stop it! They're useless. Especially so when insiders go rogue on you. You absolutely have to have a minimum two-factor authentication system to identify the user of sensitive data properly.
- Security criteria since 1981 have clearly spelled out that we must accurately audit all access to, transmission of, or modification of sensitive data. This means the bad guys can't spoof (at least not too easily) the identity of your staff, and you can trace their activities.

The technology exists...and has existed for a long, long time. Today there's no excuse, other than greed, not to invest in technology to protect the people who trust you. I assume that's why several huge class action suits are working their way through the system based upon the theory that the management of the company was negligent and didn't take care in protecting critical private information. From what I've seen of the cases, the management is going down.

You can identify your company's future guest of the feds by answering the following questions:

1. Does your company use cryptographic services to protect data in storage?
2. Do you still use passwords to access control systems as well as databases?
3. Do you employ audit controls to record who does what?
4. Have you ever suggested to your management that increased security is a good thing for your company and its customers?
5. Have you told your management that there are risks to your customers' (et al) privacy without additional security controls?
6. Has your management told you, "No, we're OK," or "We don't have the budget for it," or "It's an acceptable risk"?

You see where this is going, so get your subpoena shoes on 'cause some of you will be in the witness box describing how poor the security is at your company.

Controlling privacy in your databases is the law.

Most companies do a poor job of security and privacy management.

The lawsuits have begun. (Think Napster, P2P, RIAA, and MPAA).

The economic losses are extreme and provable.

Someone is going down. Today or tomorrow, it's either your company or the next one. Someone's going down. Just make sure it isn't you.

About the Author

Winn Schwartau is CEO of www.TheSecurityAwarenessCompany.Com and Trusted Learning, Inc. www.TrustedLearning.Com. He's a popular author and speaker with thousands of credits to his name.

winn@thesecurityawarenesscompany.com

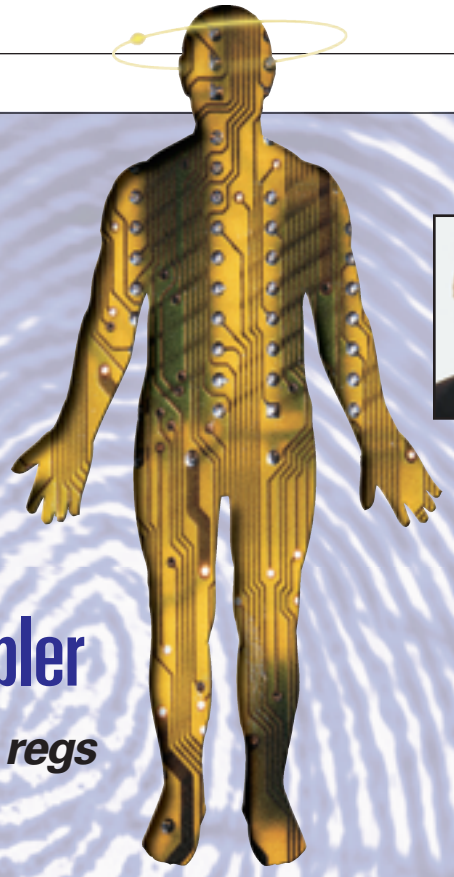

Identity Management as a Regulatory Compliance Enabler

It's the cornerstone of most major regs

BY SUMNER BLOUNT

OVER THE PAST several years, a number of factors have conspired to cause the security of information to become a critical business issue that's core to the operation of most companies. These factors include the recent corporate financial scandals, the rise of terrorism, and the increased concern over the privacy of user information. With security and privacy becoming more important everyday, the failure to maintain security over sensitive information could result in irreparable damage to a company's reputation.

These trends have resulted in new governmental regulations relating to financial reporting, security, and privacy. The importance of regulatory compliance has now become a critical boardroom issue. Companies that don't comply with these regulations risk legal action, as well as stiff fines and restrictions. As a result, regulatory compliance has become one of the top business drivers and the concern of security officers at most large enterprises.

Most new regulations don't prescribe specific technologies that have to be used to achieve compliance. In fact, many regulations can be met only with improved procedures and processes, some of which might not even involve new technology. Still, many corporations are finding that the old "paper and pencil" approach to regulatory compliance might get them through initial compliance relatively unscathed, but it's not a viable long-term solution. They are finding that full compliance is immeasurably easier if a common way of managing all their users and their access to confidential resources is implemented.

Classification of Major Regulations

Governmental regulations cover a wide range of target areas. However, the regulations that impact the IT infrastructure generally fall into one of three major categories:

- ***Governance*** – These regulations deal with issues related to the transparency and accuracy of financial records, the retention of records in the corporation, and requirements of disaster recovery and business continuity. In some cases (notably Sarbanes-Oxley), this type of regulation was heavily driven by corporate scandals and financial fraud. In short, they are intended to ensure that proper controls exist to guarantee that corporate reporting is accurate, timely, and complete.
- ***Privacy*** – These regulations are often specific to a single vertical market and dictate how a customer's personal information must be handled. There are regulations that specify what type of personal information may be kept, how it's handled (including who, if anyone, it may be given to), and what actions are required in the event of a breach of established privacy restrictions.
- ***Security*** – The role of security regulations is to protect a corporation's critical infrastructure, as well as to protect against certain external threats. Although security is a key element of many regulations, there are very few that focus exclusively on security issues, and they tend not to be formal regulations, but simply frameworks and policies that represent "best practices." In general, these regulations specify how users will be identified, how their access to sensitive resources must be controlled, and how that access can be tracked and audited.

Some regulations focus only on one of these areas. However, others include requirements that span areas, sometimes including each one of the above areas.

Figure 1 lists the major governmental regulations that most companies are required to comply with.

Table 1 summarizes the intent and purpose of each of these major regulations.

Common Requirements for Regulatory Compliance

Each of these regulations is targeted at addressing different problems, often for a different category of company. Still, there are a number of common requirements on IT in almost all of them. This commonality is important because it allows a single compliance effort to leverage its efforts across the range of regulations an individual company must comply with.

More specifically, the types of issues addressed by these regulations include:

1. ***User Authentication*** – How are users identified to a system? How secure is the method used? Are there adequate procedures for creating, managing, and changing user passwords? Are there password policies that ensure strong and changing passwords?
2. ***User Authorization*** – How strong and flexible is your method for ensuring that only properly authorized users have access to protected data and applications? Are these controls reviewed regularly to identify role conflicts that would lead to unauthorized access? Are there clearly defined rules for the treatment and processing of private information (health, financial, etc)? Are there controls so that the owners can grant or withhold permission for various people to view their information? Are users removed from the system automatically when the need arises (such as after an inactive period or inappropriate user behavior)?
3. ***User Administration*** – Do you have clear processes and controls in place to create access rights for each user? Are the necessary approvals part of the defined process? Is there an automated workflow mechanism in place to ensure that this approval process is done consistently and

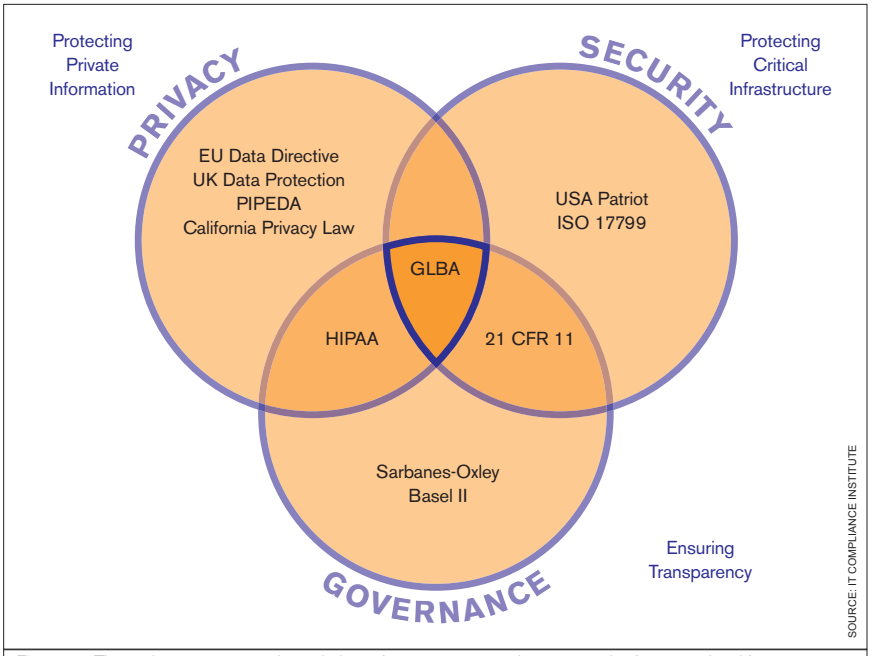

Figure 1: The major governmental regulations that most companies are required to comply with

Vertical	Regulation	Summary of Regulation
Financial Services	Gramm-Leach-Bliley (GLB)	Defines privacy requirements for customer personal financial information. Also, it restricts the use & transfer of information between organizations.
	Basel II	Defines requirements for risk management within a capital framework. The higher the risk, the more capital is required for a bank.
Healthcare & Life Sciences	HIPAA	Governs privacy and security of private user health information. Applies to providers (hospitals), payers (insurance), as well as other health-related entities such as pharmacies.
	21 CFR Part 11 (FDA)	Defines requirements for storage and access to data relating to drug development cycles. Since many FDA-regulated manufacturing processes are now done online, this regulation determines how to handle quality assurance when approvals are done online.
Cross-Sector	EU Data Directive	Provides a framework of data protection and privacy requirements that is to be reflected in national law by all member states to provide a minimum level of protection throughout the European Union nations.
	UK Data Protection Act	Provides additional data protection and privacy requirements over and above what is required by the EU Data Directive. Controls how information is kept, and provides users with the right to inspect and correct any data held about them.
	ISO 17799 (BS 7799)	Standard that sets out requirements for a "best practices" Information Security Management System (ISMS). Corporations want to conform to this to ensure that they are using "best practices" in their security infrastructure.
	Patriot Act	Requirements to combat and report money-laundering activities. Also, contains a number of anti-terrorism statutes.
	Sarbanes-Oxley Act	Requires public companies to document internal controls that relate to financial reporting. Strong internal security means better controls, and therefore easier compliance.
	California Privacy Act	Combats identity theft. Defines requirements for notification of suspected breaches of personal information. Applies to any company doing business in California.
	PIPEDA	Controls use of personal information by corporations. Requires consent for collection and use of private information. A Canadian regulation.
	NORPDA	Establishes a national standard for notification of consumers when a database breach occurs.

Table 1: The intent and purpose of each of the major regulations

formally? Are there controls to ensure that individuals can't expand their access rights inappropriately? When someone leaves the company, are their access rights terminated immediately? Are there regular reviews of all user accounts to ensure that they're correct and appropriate?

4. ***Auditing and Reporting*** – Are there comprehensive capabilities to provide real-time auditing of all important security events as well as user access? Will segregation of duties be enforced consistently so that one person doesn't have (for example) the ability to both initiate and approve a request? Will inappropriate or suspicious access be identified and corrected quickly? Are there controls to recognize attempted breaches? Are breaches identified and resolved quickly? Are there regular procedures to review all system activity to ensure that problems are identified quickly?

How Identity Management Can Aid Regulatory Compliance

The secure management of users and their access to sensitive resources is a cornerstone of almost all the major regulations that companies need to be concerned with. An integrated approach to identity and access management (IAM) across an enterprise can therefore be an important element of any regulatory compliance strategy. In fact, a centralized and automated way of dealing with user identities and their access rights is virtually a requirement for any sustainable and cost-efficient compliance effort.

Identity and access management solutions bring together people, processes, and technologies to enable organizations to manage their relationships with users throughout the user lifecycle, creating access and security policies, enforcing those policies, and automating the process of creating, and modifying and disabling digital identities. Identities can be people – such as employees, customers, suppliers, and partners – or resources – such as software programs, Web Services and machines on a network.

Let's be more specific about what an identity and access management infrastructure actually includes. Although various analysts sometimes include directories or meta-directories in their definitions, there's general agreement that the core capabilities of this type of solution include:

- ***User Administration*** – All users must have electronic identities, and these identities need to be created, managed, and reviewed periodically to ensure compliance with relevant regulations. In addition, the management of these identities needs to be delegated to the appropriate group or business unit so that it can ensure that the user's attributes and access rights are correct and current. Users also have to be able to self-service their own accounts based on a set of access policies that have been defined. Lastly, an integrated workflow capability is important so that appropriate management approvals can be granted for all identity and access requests.
- ***Access Management*** – The core of any robust IAM solution is the access management component. This capability provides a policy-driven infrastructure to securely control all user access to protected applications and information. Without this kind of technology, security is generally implemented in each application, thereby creating "silos" of security. Such an environment provides a number of compliance challenges if only because it's harder to ensure that inappropriate user access doesn't occur when each application is doing the security enforcement.

  Typically, access rights are based on the user's role, so that an integrated role-based model is an essential element of a comprehensive IAM solution. A common element of virtually all regulations is also a robust password management capability, so this should be considered essential in any IAM platform.
- ***User Provisioning*** – A robust Web-based provisioning system provides a common automated foundation with links to legacy systems and workflow procedures to automate the granting, management, and revocation of access to digital resources according to ever-changing business and/or regulatory requirements. As a typical example, one of the most common areas of non-compliance with these regulations revolves around inadequately removing a user's access when that user leaves an organization. These so-called "orphan accounts" are estimated to be roughly 30% of all the accounts at most large corporations. Provisioning solutions should not only automate the assignment of resources to new employees (for greater and faster productivity), but also automatically remove access rights and accounts on termination. The absence of such a capability would almost certainly cause non-compliance with any regulation that required strong controls over the access rights of departed employees.

  These capabilities are critical for any IAM platform. Without the full suite, enterprises face a far more difficult task of compliance, because there's no common model for handling user identities, their access rights, and the allocation and de-allocation of their resources.

In addition, it's critically important that these core components be integrated into a single platform, rather than a set of separate, but relatively non-integrated, technologies. Native and tight integration allows a common role model, for example, that makes compliance easier because it's far more straightforward to manage all user access, as well as identify areas that need management attention (for example, the segregation of duties violations).

Summary

The emergence of a number of important governmental regulations is having a profound impact on most corporations, regardless of their vertical market. Yet, despite the amount of effort being expended to achieve compliance, we've seen that there are some basic requirements on the IT security infrastructure that are common across all these regulations. Corporations should consider the role of an integrated identity and access management platform in helping them meet these regulatory requirements.

About the Author

Sumner Blount is the director of security solutions marketing at Computer Associates. He has been associated with the development and marketing of software products for over 25 years. He has managed the large computer operating system development group at Digital Equipment and Prime Computer, and was director of software for Pathway Designs. He was instrumental in the original conception and development of the DCE technology from the Open Software Foundation and served as the DCE program manager at DEC.

sumner.blount@ca.com

www.sys-con.com/edge2005

web services EDGE conference&expo

FALL SERIES

[Engage and Explore]

the technologies, solutions and applications that are driving today's **Web services** initiatives and strategies...

CALL FOR PAPERS NOW OPEN!*

Coming to a City Near You ►

Web Services Edge
Fall Conference Series

3 Dynamic Conference Programs
Targeting Major Industry Markets

20+ seminars within 5 tracks will address the hottest topics & issues:

- ► Web Services: The Benefits and Challenges
- ► Web Services Security
- ► SOA (Service-Oriented Architecture) and ESB (Enterprise Service Bus) Strategies
- ► Interoperability, Incremental Integration, & Open Source
- ► The Management Process in Developing a Web Services Strategy

Program Features:

- ► Keynotes
- ► Tutorials
- ► Panel Discussions

Why Attend:

- ► Improve the return on your technology investment
- ► Develop & sharpen your strategy and identify key action steps
- ► Find new ways to reach and impress customers with Web services
- ► Maximize the power of your enterprise
- ► Protect your business from security threats
- ► Assess Web services as a viable option

Attention Exhibitors:

- ► An Exhibit-Forum will display leading Web services products, services, and solutions

The Westin Washington, D.C.
Washington, D.C.
September 7-8, 2005

1

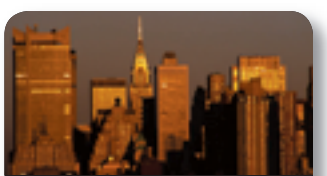

Helmsley Hotel
New York, NY
September 19-20, 2005

2

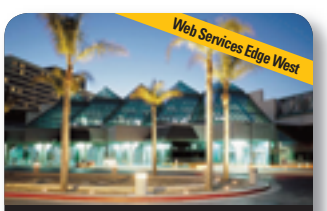

The Westin Santa Clara Convention Center
Santa Clara, CA
October 24-25, 2005

3

Register Today! www.SYS-CON.com/Edge2005

Sponsored by

WebServices JOURNAL · XML JOURNAL · .NET JOURNAL · eclipse developer's journal · WebSphere JOURNAL · information STORAGE+SECURITY journal · wldj · JDJ · LINUXWORLD MAGAZINE · MX developer's journal · asp.netPRO · SD Times · CoDe · Software Test & Performance

*Call for Papers email: grisha@sys-con.com

For Exhibit and Sponsorship Information ► Call 201 802-3066

Produced by SYS-CON EVENTS

Compliance & the Role of Security Patch & Vulnerability Management

DIFFERENTIATING BETWEEN HIGHER-RISK SYSTEMS AND HIGHER RISK

BY SEAN MOSHIR

IT SECURITY PROFESSIONALS who are already managing the bottom-line expectations of their boardrooms while guarding their organizations against myriad security threats have a new "C-level" challenge — that of compliance. These professionals must now meet voluntary and mandatory regulations such as Sarbanes-Oxley, the Gramm-Leach-Bliley Act, the Federal Information Security Management Act (FISMA), and the Health Insurance Portability and Accountability Act (HIPPA).

A critical piece of the compliancy puzzle and an instrumental component already supporting the enterprise IT security posture, patch and vulnerability management is quickly becoming a multi-faceted solution that IT security professionals are employing to address this challenge.

However, determining the exact role of automated patch and vulnerability management and policy and how it supports and meets voluntary and regulatory compliance has become the question. Here we offer an answer to help the security professional understand the role of patch and vulnerability management in supporting auditable compliance — the state of compliancy at which an organization's IT control and security status and reporting ultimately meet the requirements of many of today's corporate policies and regulatory acts. Moreover, it's important to know that as a security professional the ability to achieve a state of auditable compliance is real and attainable.

The Dual Dilemma: Security and Compliancy

IT departments need to secure their infrastructure against threats such as viruses and worms and manage their network's availability. They also must demonstrate to auditors the adequacy of this security through measurement and process. The roadmap to regulatory compliance (see

Figure 1) shows how the three technical competencies of an enterprise IT security program — command and control, threat mitigation, and audit and monitoring — strongly support the business goals of management, security and availability.

According to Yankee Group analyst Phoebe Waterfield, "management" is the Achilles' heel of many organizations and is central to demonstrating regulatory compliance. For example, she says, keeping up with critical patches for desktops, servers, and applications demonstrates effective management of IT systems and controls. Moreover, there's an urgent need for more effective processes to ensure that critical patches and fixes are applied to higher-risk systems in a timely manner. To this end, many organizations are now turning to strategic security solutions like patch and vulnerability management to improve systems security management processes and meet the burden of proof regulations impose.

Systems Management Under Security Scrutiny

The most frequent question organizations ask about regulations is how to interpret them and apply real-world security solutions to meet compliance. Regulations don't stipulate any specific security tools or products. Instead, they ask each organization to demonstrate how they are protecting the information contained in IT systems to a level commensurate with its value. For example:

- Applying adequate security protections for customer data according to stated corporate security and privacy policies
- Following best practices to secure a network's perimeter and access to systems on the network
- Protecting confidential information and limiting access to personally identifiable information

The lack of specificity in regulations creates uncertainty about what compliance means, what auditors are looking for, and what vulnerabilities are considered unacceptable risks. IT departments can judge their general compliance readiness using the roadmap depicted in Figure 1. More than likely, some of these areas lack definition in most organizations.

The most common missing component is security management. This is a serious issue because it prevents IT personnel from demonstrating and reporting the effectiveness of the security components already in place. For example, a business may have a systems management tool such as Microsoft's SMS and use it for keeping endpoints such as workstations, laptops, and servers updated with the latest patches. However, this tool doesn't provide the necessary reports that give an accurate and timely picture of which systems are missing patches. It's no longer enough for organizations to patch; they must now prove that systems are patched to a reasonable level and measure how effective patching processes are at reducing vulnerabilities on their network. Figure 2 shows how manual or less than best-of-breed patch management processes leave systems exposed for long periods of time.

Malicious code threats propagate in days, not weeks. For example, according to

Yankee Group research, the Nimda and Slammer worms spread worldwide in less than an hour, but the signature updates to protect networks weren't available for 24 hours. Then enterprises had to get these updates out on their networks to be protected. The impact and cost of virus and worm infections are largely attributed to unpatched and misconfigured systems running a vulnerable service. According to Gartner, over 90% of security exploits are carried out through vulnerabilities for which there are known patches.

Regulations Demand IT Control and Command

Regulations require organizations to demonstrate compliance through risk management and sufficient IT controls. Regulations don't always cite specific controls, so provided below are descriptions of how auditors are interpreting the law and some emerging standards that are used for industry benchmarking.

Sarbanes-Oxley Act

The Sarbanes-Oxley Act (SOX) is the single most important piece of legislation affecting public corporations since the U.S. securities laws of the early 1930s. The act charges managers of public companies with the task of certifying that they have an operational system of internal controls over financial reporting. SOX followed in the wake of several major accounting scandals and raised heads because it holds corporate executives directly responsible for the accuracy of the financial statements their companies make to the Securities and Exchange Commission (SEC). Section 404 of Sarbanes-Oxley deals with IT systems. It requires an annual evaluation of the internal controls protecting the systems used to prepare a company's financial statements. Companies must vouch for the veracity of their financial data and, significantly, the effectiveness of these internal security controls. Sarbanes-Oxley was supposed to come into effect in November 2004. However, the SEC recently pushed this deadline back a year, a relief to many publicly traded U.S. companies that now must make significant changes in their infrastructures to meet SOX IT control requirements.

The Federal Information Security Management Act (FISMA)

The Federal Information Security Management Act, enacted in December 2002, requires that federal agencies ensure the effectiveness of the information security controls protecting federal operations and assets. FISMA specifies what agencies must do to strengthen security. It's an agency-wide information security policy specifying how to classify data as confidential and how to protect data and systems according to their criticality. It mandates a set of internal audit activities, including periodic testing of IT security controls and submitting an annual report on FISMA compliance to the National Institute of Standards and Technology (NIST). Federal agencies are specifically required to report how they ensure all systems are patched in a timely manner.

The Health Insurance Portability and Accountability Act (HIPPA)

The HIPAA Privacy Rule of 2002 was first implemented on August 26, 1996 by the U.S. Department of Health and Human Services (HHS). HIPAA requires the secretary of HHS to publicize standards for the electronic exchange, privacy, and security of health information. The Privacy Rule of HIPAA, finalized in 2002, attempts to ensure the security — in particular, the confidentiality — of health information. The policies and procedures set down by HIPAA apply to health plans, healthcare clearinghouses, and any

healthcare provider that transmits health information electronically. Under the Privacy Rule, HIPAA mandates adequate technical safeguards to prevent intentional or unintentional access to confidential health information.

Homeland Security Presidential Directives

In the past few years, the White House has released a number of Homeland Security Presidential Directives (HSPDs) to address computer-related threats against the U.S. HSPD 7 requires all federal departments and agencies to protect critical U.S. infrastructure from possible terrorist attacks. This includes working with state and local governments as well as the private sector to prioritize, remediate, and protect their electronic resources concurrent with FISMA.

Emerging Industry Standards

In June 2004, the Basel Committee on Banking Supervision released its new standard — the Basel II, International Convergence of Capital Measurement and Capital Standards: A Revised Framework. The Basel Committee is independent of any government organization and prepared this standard for the banking industry's voluntary participation. Basel II Section 745 describes in detail what constitutes proper internal control review and stresses the need for the accuracy and security of data.

The Committee of Sponsoring Organizations (COSO) released an updated framework for enterprise risk management (ERM) in September 2004. ERM is used by companies for various reasons such as compliance with applicable laws and regulations. It outlines eight components necessary for risk management: 1) internal environment, 2) objective setting, 3) event identification, 4) risk assessment, 5) risk response, 6) control activities, 7) information and communication, and 8) monitoring. Many companies use this framework to comply with regulations such as Sarbanes-Oxley.

Standards in healthcare are set by the Joint Commission on Accreditation of Healthcare Organizations (JCAHO) and the National Committee for Quality Assurance (NCQA). These non-profit organizations are committed to raising the standard of healthcare in the U.S. Their rigorous published standards for healthcare organizations form the basis for accreditation for HIPAA compliance and to inspire confidence in their services among clients. In 1997, JCAHO and NCQA began collaborating on the issue of patient confidentiality and published a report called "Protecting Personal Health Information: A Framework for Meeting the Challenges in a Managed Care Environment." This framework advises healthcare organizations on the use of technology to secure sensitive information.

How Much Patching and Configuration Updating Do Regulations Require?

Regulations have led to the public scrutiny of security practices, adding legal burdens and forcing organizations to work harder to ensure the safety of their IT systems. Although only FISMA specifically cites patching as a requirement, regulations have added urgency to the task of keeping up with critical patches and configurations for desktops, servers, and applications. The burning question for many is: Given that vendors such as Microsoft are releasing 40-plus critical patches a year, and even more for Microsoft Office and Internet Explorer, do regulators need to see systems patched or fixed to the highest level at all times?

The answer is no. Regulations call for a demonstrated ability to manage and patch systems according to risk. For patch and vulnerability management, regulations require a business process that takes into account both business and security risks. This means a business process supporting the following risk-based principles:

- Critical patches are applied more quickly than less critical patches.
- Patches and reconfigurations are applied to overexposed or higher-risk systems before they are applied to low-risk systems.
- Patches and fixes are only applied to systems when the benefit of the patch or configuration outweighs the associated business disruption.

A process that includes these three principles meets the needs of both HIPAA and Sarbanes-Oxley because it treats each patch or configuration fix according to risk. Organizations only need to demonstrate that an effective process is in place to patch or fix systems.

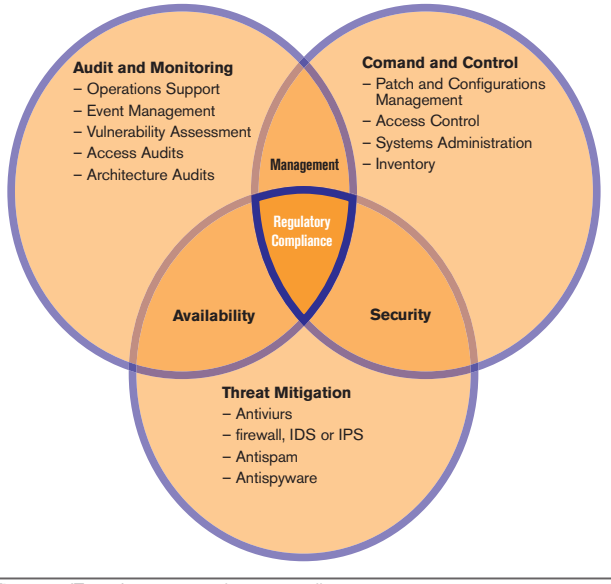

Figure 1: IT road map to regulatory compliance

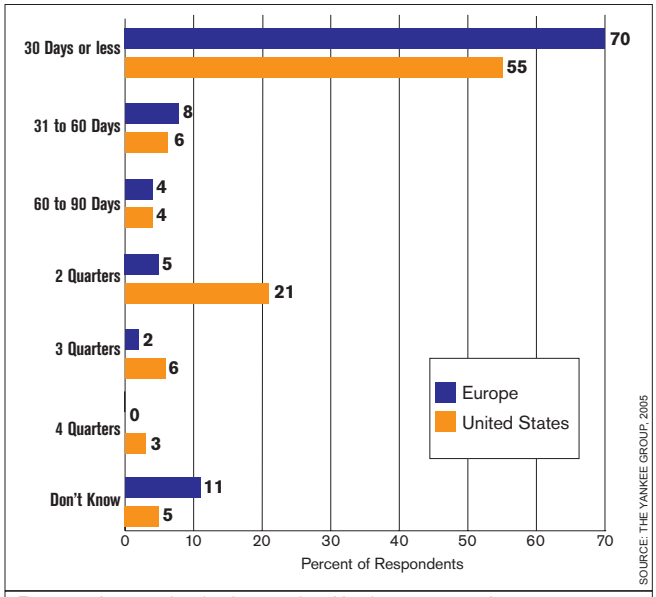

Figure 2: Average time it takes regulated businesses to patch systems

Patch and Vulnerability Management Supports Compliance

Automated patch and vulnerability management provides a high-performance, enterprise-scalable solution to secure systems management in regulated environments. Moreover, agent-based architectures that power many of the leading patch and vulnerability management solutions on the market today, automate the discovery and distribution of patches and agents on all remote endpoints, including Unix, Linux, Windows, and Macintosh.

Most important to compliance efforts are the management, reporting, and alerting features of these solutions. Patch and vulnerability management technology should provide the ability to create and manage arbitrary groups of computers. Also it should provide the ability to automatically enforce mandatory patches and fix baseline policies on the members of each group so as to provide a patch-and-configuration compliance assurance mechanism. Additionally, custom reports produced by the patch and vulnerability management solution should then be able to identify all the computers that are out of compliance with the corporate policy.

A patch and vulnerability management solution alert feature that assists compliance efforts should automatically notify administrators of new patches and computers that don't meet baseline patch and configuration policies. This removes the need for administrators to track multiple vendors' patch releases and perform manual audits. Furthermore, a value-add provider of patch and vulnerability management technology and services researches, tests, and approves patches from an array of vendors including Adobe, Citrix, Microsoft, IBM, Oracle, Symantec, McAfee, Sophos, RealAudio, MP3, Compaq, Novell and many other applications. This directly supports end-user testing efforts and speeds up system compliance and in parallel supports an organization's overall security posture.

Achieving Compliancy While Instilling Trust and Managing Risk

Effective patch and vulnerability management instills partners and auditors with a feeling of trust that security processes are meeting regulatory and corporate compliance. This trust permeates the organization as a whole, easing many of the difficulties that arise between business and IT when urgent patching is needed. For IT professionals to balance regulatory security responsibility with pressing business needs, they need a system that differentiates between higher-risk systems and higher-risk vulnerabilities. In the end, the real challenge is making the right decisions about when and where to patch or update based on business risk and eliminating the burden of proof that regulations and management impose.

References:

- *Balance Regulatory Compliance Needs with Secure Systems Management,* Yankee Group, January 2005
- *Sarbanes-Oxley Compliance: Management Technology Controls*, WatchIT.com, 2004

About the Author

Sean Moshir is chairman and chief executive officer of PatchLink Corporation, a company he founded in 1991. During the course of his career, Sean has focused exclusively on strategic security software and services that supports business continuity with an emphasis on patch, vulnerability, and compliance (PVC) management functionality and support.

seanm@patchlink.com

ISSJ | Advertiser Index

What to Look for in an Endpoint Intrusion Prevention Solution

A SAFETY CHECKLIST

BY SAMAN AMARASINGHE

EVEN IN THE best of times, security products that aim to thwart worms are playing catch-up. Anti-virus and anti-malware products are populated with signatures created for attack specific signatures, which are created and distributed only after an attack is underway. Internet worms propagate too quickly for such reactive solutions to be effective. This is a major problem in maintaining information security as well as providing business continuity for many organizations.

In the "good old days" worm creation was literally a mischievous playground for young hackers to show off their skills. Unfortunately, worms today are part of the criminal syndicate. Worms are used in a high-stake, high-tech version of the neighborhood shakedown. Enterprises dependent on Internet commerce are extorted by threats of distributed denial of service (DDOS) attacks. To pull off DDOS attacks criminals need large remotely controlled "BOT" networks. Typically worms are used to take over the unprotected machines of unsuspecting users to create these BOT networks. According to Symantec, the number of bots jumped more than 15 fold in the first six months of 2004.

Driven by changes in the attack landscape, pandemic worm breakouts (e.g., the Slammer, Blaster, and Sasser worms), as well as Sarbanes-Oxley and other compliance requirements, enterprises are increasingly taking on high-profile anti-malware or worm mitigation projects.

These projects typically involve an Endpoint Intrusion Prevention Solution. There are a multitude of different host-based solutions and technologies in this space. Selecting one that addresses your organization's needs is no easy task. However, any solution you choose should address the following seven criteria:

1. Accuracy

Accuracy, or the ability to correctly identify an intrusion, is required in a good solution. Even more importantly, the solution shouldn't tag a normal operation as an intrusion – a false positive. Unlike an Intrusion Detection Systems (IDS), where false positives are just a nuisance, each false positive in an intrusion prevention system will disrupt a normal business operation. So a solution that doesn't treat a false positive as a software bug but asks you to live with it isn't a viable solution. As billions of normal events occur between attack events, even the smallest rate of false positives will make the solution negatively impact business operations at a rate higher than the worms themselves. Accuracy is the most important criteria because it can have a continuous and immediate impact on your normal business operations well before you encounter an actual attack.

2. Maintainability

Maintainability is an obvious and important criterion for a product that will be deployed enterprise-wide. A solution that requires individual attention in each installation or tuning every time a system is updated, upgraded, or used for a different purpose, will become a management nightmare. Unlike signature-based systems, where the vendor does day-to-day signature creation, policy-based and learning-based systems offload most of the work onto you. In a policy-based system, you may need to fine-tune policies on every machine and work to eliminate the false positives that show up. You also must convince yourself and your management that the policy you created is sufficiently stringent to catch the next attack. In a learning-based system, you need to teach the system by stressing it with all the possible execution scenarios and hope that you don't miss any critical ones. This requires that the system be in detection mode in production long enough for the system to learn all the common scenarios or that you set a sufficiently tight policy before a system is put in protection mode. During this time it's vulnerable to attack. A solution where the cost of operation is higher than the cost of cleaning up a few worm infections a year provides little ROI.

3. Scalability

Scalability of a system to an enterprise-wide deployment is important because all mass worms to date have attacked the entire infrastructures – including servers, desktops, laptops, even embedded controllers. So, rolling out a solution that protects everything in the enterprise is critical. This applies especially to machines running Windows, since most of the recent worms have targeted the core infrastructure programs in those machines. Any solution that requires constant individual attention at each endpoint doesn't scale. Critical centralized components also hinder scaling.

4. Coverage

Coverage measures the range of attacks a solution can protect against. Contrary to marketing claims no solution

can systematically and comprehensively handle all intrusions — both known and unknown. It's essential that you understand what class of attacks a solution covers. See that the coverage complements the existing layers of security in your organization. Although more than 500 mass intrusions were detected during the last year, you should be protected from most of them by commonly deployed products such as anti virus systems. Since it's impossible to predict future attacks, look at the past and see what attacks got through the existing layers of protection. Focus on specific intrusions that your current security systems didn't handle satisfactorily. What would have happened if you had this product deployed during that intrusion?

Each intrusion exploits an application's vulnerability. Thus the published vulnerabilities in your critical applications and operating systems are an important resource for understanding how comprehensively a solution covers the kind of attacks that can take advantage of these vulnerabilities. A comprehensive list of vulnerabilities can be found in the CVE list or Microsoft security bulletins. Since most attacks exploit an existing vulnerability, a solution that can't cover the most important vulnerabilities, it isn't the solution for you.

Since no enterprise can implement barriers against every conceivable intrusion it's critical to prioritize. Some kinds of intrusions have devastating consequences; others are important; many are rare; some are still theoretical or imaginary and only exist in the minds of researchers. Trying to protect an enterprise against all these possibilities creates an enormous cost in implementation and management and makes the infrastructure less stable. Worst of all, the noise drowns out absolutely necessary and critical protections. Therefore, you need a good understanding of the attack landscape and must precisely define the remaining holes in your current protection shield. This way you can prioritize and protect against the critical attacks at a reasonable cost while ignoring the noise.

5. Proactivity

Proactivity, the ability to stop an attack with the least amount of attack-specific information, is extremely important against zero-day attacks. A solution that requires a new signature update to stop an attack is no use against rapidly propagating worms such as Slammer, which only took 10 minutes to go from 0% to 90% infection worldwide. It is not sufficiently proactive to protect against modern worms. Solutions that require some knowledge of the vulnerabilities provide some level of proactivity. To date, attack writers have relied on the vulnerabilities revealed by the ISVs when they release a patch. In these cases vulnerability information was available before an attack. However, during the last two years, the time between the revelation of the vulnerability and the release of an attack targeting it has decreased. So the best solutions are the ones that stop attacks without any special knowledge of the attack or the vulnerability. A solution that's technically capable of stopping an attack but not sufficiently proactive will gear up to stop a day-zero attack only after that attack has created havoc in your enterprise.

6. Uncircumventability

Uncircumventability of a solution is essential for a viable defense against attackers who are knowledgeable, resolute, and resourceful. Shockingly, many solutions out on the market are easily circumvented. To test a solution fully against zero-day attacks you would need to create a new attack, which isn't practical. However, there are other alternatives. Attack tools are available on the market that use a collection of existing vulnerabilities and attacks to probe a system. Make sure you test the solution on an unpatched system vulnerable to those attacks. Another approach is to ask competing vendors to break each other's products. The Internet has become a powerful educational tool for attackers and you can also benefit from it too. You may be surprised to find information on simple ways to break most products on the Internet. Don't be an emperor without any clothes by using an easily circumvented solution and convincing yourself and your organization that you're well protected.

7. Containment

Containment indicates how successful an attack is before the solution can detect and stop it. A solution that stops an attack before it's loaded into the system or before a single instruction from the malicious payload gets executed provides the best containment. If the attack partially executes, you may be required to do a detailed forensic analysis. For example, California privacy law SB 1386 requires that companies disclose any possibility of a security breach. Furthermore, executing even a small number of instructions provides an attacker with an opportunity to circumvent the solution.

What Is an Ideal Endpoint Intrusion Prevention Solution?

Accuracy:	An ideal solution produces zero false positives.
Maintainability:	An ideal solution is easy to maintain and administer.
Scalability:	An ideal solution scales across the entire enterprise.
Coverage:	An ideal solution completely covers the class of attacks you're trying to protect against.
Proactivity:	An ideal solution proactively blocks attacks without any special knowledge of the attack or the vulnerability.
Uncircumventability:	An ideal solution is impossible for hackers to circumvent.
Containment:	An ideal solution stops the attack before it causes damage and spreads.

Author Bio

Dr. Saman Amarasinghe is cofounder and CTO of Determina, Inc., an associate professor of the Department of Electrical Engineering and Computer Science at MIT, and a member of the Computer Science and Artificial Intelligence Laboratory (CSAIL).

saman@determina.com

The Dark Corner of Regulatory Compliance

EXTENDING IT REGULATORY COMPLIANCE BEST PRACTICES TO THE BRANCH OFFICE

BY NOAH BRESLOW

FOR A DECADE or more, internal and external regulations have impacted corporate IT and shaped the way companies are required to do business. Today you would scarcely find a corporate CIO, IS director, or IT architect who would build or maintain a storage infrastructure without sweating the details of how to meet the regulations and policies that govern the data their companies produce. The ongoing health and well being, if not the future of their companies, depend on this diligence.

Each government or agency regulation carries with it very specific IT-related requirements that impact corporate IT decision making. The Sarbanes Oxley Act of 2002 dictates that companies must secure, store, and archive all documents, records, and business data or their corporate executives may face stiff fines or jail time. The Health Insurance Protection and Portability Act (HIPPA) requires that healthcare providers, health plans, and public health authorities guarantee that data is standardized and scrupulously protected. The Occupational Safety and Health Administration (OSHA) says that employers must maintain and archive worker injury and illness records and that those records must be easily accessible to auditors. The list of regulations goes on and on, and the conscientious corporate IT professional must design an information management compliance (IMC) strategy that will protect his company without breaking the bank.

But this same IT professional often has a serious issue that he must face when designing his IMC or IT compliance strategy. Worst of all, it's an issue they often don't have a solution for. These professionals may have been meticulous in protecting, managing, and backing up all the data that resides in their data centers and other key locations, but more often than not they have been unable to achieve that same level of compliance for data that lives at the "dark corners" of their network in branch and remote offices.

Needless to say, the consequences of an overlooked branch office ICM or IT compliance strategy can be huge. In one recent case it was reported that a single lost e-mail caused a ripple effect that cost a financial services company over $400 million in regulatory fines. In another, the FDA refused to allow a major pharmaceutical company to produce a certain drug because the servers storing the company's files didn't meet the agency's requirements. When you add up the cost of legal actions and penalties, and then total in revenue shortfalls due to lost productivity and opportunity, it's easy to see how regulatory compliance IT oversights are costing companies millions of dollars every year. And with 87% of corporate employees working in locations that are remote from their headquarters' location, the branch office is one of the prime areas where these oversights are occurring.

The Possible Dream: Protecting Branch Office Data

It's not that IT professionals and the companies they support are willingly negligent. It's just that certain hard truths about managing enterprise branch office data make it impossible. Those charged with managing the data generated by these offices may attempt to do so by hiring or assigning branch office personnel to backup and maintain it, or they may try to manage this data "long distance" over a Wide Area Network (WAN). In either case, these efforts are often doomed from the start.

In the first case, a lack of appropriate corporate resources often considerably weakens the plan (assigning workers, who have other responsibilities, to oversee data backup never seems to work and hiring dedicated skilled IT workers at each branch office can just be too costly). And in the second case, WAN latency and unreliability create a nearly intractable roadblock that simply can't be overcome. Studies show that nearly 75% of data resides outside the corporate data center, and that most of that data is not reliably backed up on a consistent basis. These unreliable efforts to manage branch office IT data are often the reason why these backups never occur or don't complete successfully.

WAN latency and lack of proper corporate resources are forcing data center IT managers to physically ship data between locations, make numerous site visits, or spend hours on the phone. All this in an attempt to make sure that their companies are in compliance with government and internal regulations. And in the end, they are often left in the same place they began, with mounds of branch office data left at risk.

For corporations to ever completely and effectively comply with regulations across the organization, it's clear that they must centralize their corporate data. Consolidating and centralizing all data in a single location, lets companies keep tabs on all their data, enabling them to take full advantage of the regulatory compliance technology and strategies they've installed at their data centers. A single repository for all corporate data means that there is only one point of "entry or exit" for data; and that means that IT

“Learn what Microsoft® doesn’t want you to know!”

It is almost certain that there are people outside your organization who know an account and password that has complete access to your data and your network.

We have a White Paper that tells you how to determine if you are at risk, and what you can do about it.

While others have paid thousands of dollars for this information, you can get it for free—but only if you act immediately!

Register and download your FREE copy today:
www.e7software.com/risk or call 1-800-824-4717

We are looking to collect some additional information on the size and scope of the corporate data that may be at risk. After you register, if you complete our brief survey, you will be entered to win one of 5 iPod® shuffles that we will be giving away. But remember, the real value is the White Paper. Don’t leave your data at risk. To fill out our survey go to **www.e7software.com/risk** or call 1-800-824-4717

Register now before it is too late! The drawing is on June 30th.

e7software™

managers and architects have only one place to look to completely protect and manage that corporate data.

So why haven't corporations centralized their data to better protect and monitor it? Again, the culprit is most often the wide area network. Given the architecture and logistics of the WAN, if corporations were to centrally locate all their corporate data – effectively moving it out of the dark corners of the enterprise - workers in branches and remote offices would have a difficult time, at best, in accessing it. WAN latency and inconsistencies make accessing centrally located data painfully slow and sometimes impossible for remote workers. For corporations and workers who have come to rely on LAN speed and reliability this is an unacceptable compromise. Choosing between centralizing data but limiting access or leaving data unprotected in branch offices is an agonizing dilemma for corporations, and most often they simply take a chance on one or the other and then hope for the best.

At the same time, corporations and the IT professionals who support them find themselves dreaming of what might be possible if the WAN weren't a factor. Once you address WAN latency, the issues start to disappear. Once you figure out a way to ensure that data transfers will survive WAN disruptions, you stand a chance of centralizing all of your data. And that's where recent technology advances have started to play a major role in assuring regulatory compliance for all data throughout the organization.

Breaking the Wide Area Barrier

A technology called Wide Area File Services (WAFS) has been developed that utilizes a combination of techniques to nearly eliminate the WAN latency and unreliability that makes protecting branch office data so difficult. Best-in-breed WAFS software solutions deploy compression, bandwidth optimization, and data aggregation techniques that dramatically reduce the number of remote procedure calls (RPCs) or round trips that data has to take across the network, essentially optimizing the transfer of file applications such as Microsoft Office and Adobe Photoshop that were never designed for the distances and latencies of the WAN. These same WAFS solutions then deploy read/write file caching, internal locking, persistent logging, and server failover to ensure that any data that's being accessed or saved across the WAN is safe from packet loss or network disruptions. The combination of these technologies creates a branch office IT solution that essentially breaks the wide area barrier to branch office regulatory compliance.

A WAFS solution for the enterprise typically consist of two appliances: a server appliance that's installed in the corporate data center and a branch office appliance that's installed at any branch office that needs to access, consolidate, back up, or protect file data. For remote users, the WAFS solution makes file data appear almost as if the data were stored at their branch office, even when that data resides on a file server thousands of miles away. But for IT managers, it's important to note that neither the branch office nor data center appliances are storage repositories – and so they are not subject to the same regulatory requirements as data stored on a file server. The authoritative copy of all data in the WAFS system resides on file servers at the corporate data center, allowing branch offices to leverage centralized compliance infrastructure.

Best-in-breed WAFS products also offer a number of other benefits that meet more than just a company's branch office regulatory compliance and data consolidation requirements. To begin with, WAFS allows users to share and collaborate on consistent and coherent data in real-time between any locations anywhere in the world, increasing productivity and business continuity. Secondly, the most advanced WAFS solutions will also enable organizations to deploy a "suite" of centrally managed, low-cost branch office services such as print, e-mail, management, and Web caching, without having to deploy servers for these functions at each branch office. Finally, WAFS is extremely cost-effective, driving branch office IT costs lower and insuring a return on investment often in under a year, making it a solution that meets the needs of both the IT and business/management teams.

No More Dark Corners

With WAFS, corporations can now leverage their core regulatory compliance IT infrastructure and strategy to better ensure that storage, backup, and oversight of data from even the darkest corners of their organization will cost-effectively and seamlessly meet internal and external regulations. WAFS has been purpose-built to enable this kind of storage and server consolidation, and that makes WAFS a natural technology for companies to deploy as a key element of their regulatory compliance strategy. Whether it's compliance with Sarbanes Oxley, HIPPA, OSHA, or any number of the hundreds of other federal or industry regulations, WAFS lets corporate IT professionals move all of their data out of the dark corners of the network and into the light of the IT regulatory compliance strategies they've worked so hard to deploy.

"A technology called Wide Area File Services (WAFS) has been developed that utilizes a combination of techniques to **nearly eliminate the WAN latency and unreliability that makes protecting branch office data so difficult**"

About the Author

Noah Breslow is the vice president of marketing at Tacit Networks. He has significant experience in marketing, operations, and engineering in emerging technology companies. He holds an MBA with Distinction from Harvard Business School, a BS in computer science and engineering from MIT, and a U.S. patent in network protocol optimization.

breslow@tacitnetworks.com

Proven Strategies for Protecting Storage Data at Rest, in Flight, and Offsite

PREVENTING STORAGE BREACHES

BY DORE ROSENBLUM

BASED ON RECENT incidents, C-level executives are quickly realizing that in today's increasingly regulated and distributed environments, it's no longer sufficient to rely on status quo barriers of protection for critical corporate information. Instead, security executives are now faced with developing a comprehensive, ground-up strategy to protect critical information at all times from attack. This includes security for data-at-rest and data-in-flight. It also extends to data managed at offsite locations by outside service providers (e.g., disaster recovery services). Regardless of where the data resides, companies are expected to assure their customers that their sensitive information is being handled with the best security practices and procedures.

Why is demand now stronger for storage security solutions? Regulatory compliance, security de-perimeterization, and storage consolidation are combining to increase the urgency of implementing information privacy solutions.

1. ***Regulatory compliance is driving companies to implement stronger security to protect consumer privacy in the event of a breach.*** California Senate Bill 1386 has recently been in the news because it requires companies to disclose the loss of personal information from a California resident unless the information is encrypted. Without this regulation, many of the recent identity theft stories would never have become public. Lawmakers in Washington are working on legislation that would mandate similar identity theft disclosures on a nationwide scale. Other regulations that protect sensitive information are usually industry specific, including HIPAA for the heath care segment and GLBA for the finance sector.

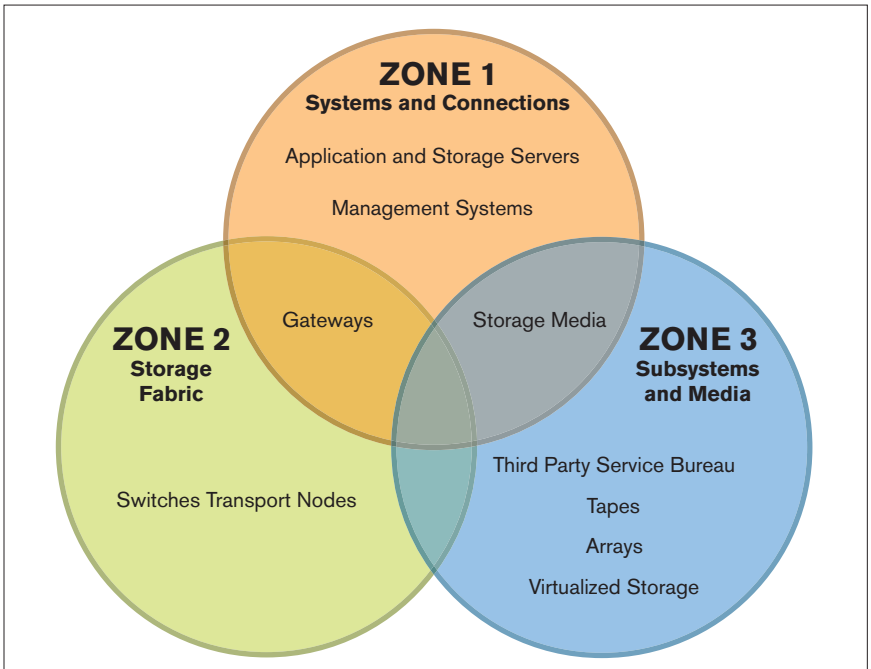

Figure 1: Threat zones

2. ***Security de-perimeterization works to address insider threats, which are often greater than outside threats, with respect to information privacy.*** Instead of relying solely on external perimeter security at the corporate firewall, companies are now building security layers to protect internal servers and storage resources. Conceptually, all users are treated as outsiders and must be authorized to access resources. This new security model requires stronger protection of stored information.
3. ***Storage consolidation minimizes storage costs by leveraging shared resources (e.g., tape silos, disk arrays) and outside providers (e.g., vaulting, disaster recovery services).*** While saving costs, consolidation opens new threats by allowing data to be shared by many different servers. In the past, a hacker would have to compromise each server to access information. With networked storage, a hacker can potentially gain access to all stored information without having to compromise any servers.

With the threat of public exposure it's easy to see why companies are focused on building security layers to protect their networked storage. These resources represent prime targets, because regulated data is usually stored on these networks. As with IP networks, storage networks are susceptible to published security threats such as system breach, spoofing, denial of service, unauthorized access, internal attack, and data theft. Many such threats are being explored and are in varying stages of being addressed by a variety of

industry consortiums and standards bodies including the Storage Network Industry Association (SNIA), Internet Engineering Task Force (IETF), and the International Committee for Information Technology Standards (INCITS) Technical Committee T11 for device-level interfaces.

Network Storage Threats

There are three threat zones that affect networked storage regardless of the network protocol employed (see Figure 1). These threat zones are systems/connections; storage fabric and management services; and subsystems/media.

System/Connection

The system/connection threat zone includes computer systems such as application and management servers, and gateway devices that connect to storage infrastructures. The storage network may become vulnerable to unauthorized data access, denial of service attack, and/or service loss if the administrative or application access to the system or device is compromised. Unauthorized systems access is often obtained through poorly managed configurations, unused services, or default settings. Once overcome, these systems can attempt to compromise media servers or issue abusive requests to storage subsystems for the purposes of data theft, corruption, or service denial.

Storage Fabric

The second threat zone is storage fabric and transport. In the case of Fibre Channel networks, this includes the directors/switches along with SAN extension solutions across MAN/WAN networks. Threats at this layer include:

- ***Data access from an unauthorized server:*** Storage administrators can direct specific storage traffic through segregated switch ports – essentially configuring which storage sources and destinations can communicate. Zoning and LUN masking are used to create the logical isolation, although determined hackers can bypass these security measures by spoofing or attacking the fabric management. This attack could result in material compromise of the storage network and pose a serious threat to data integrity.
- ***Eavesdropping of data-in-flight:*** As storage networks are extended across public MAN/WAN networks, data should be encrypted to ensure privacy. Encryption solutions are typically used to securely tunnel storage data across lower-speed IP WAN networks using IPSec. However, due to the stringent performance and latency demands of real time applications (e.g., synchronous mirroring), companies have often utilized WDM or SONET networks without encryption.

Subsystem/Media

The third threat zone encompasses storage devices, subsystems, and media. This threat to media is often viewed as a more serious risk than access to data in transit, because the potential exposure is much larger considering the amount of information stored on disk or tape. By securing the media, security professionals can protect against two threats:

- ***Lost/stolen media:*** Most stored data on tapes is left in-the-clear, unencrypted, on removable media, which can be lost, stolen, or compromised. Unauthorized users can readily read

tape data, analyze confidential information, and in some cases rebuild entire systems. Given that the data is removable, the perpetrators have more time and resources for tape inspection. As with offsite tape, failed or old disks with recoverable data are often sent to outside repair facilities where sensitive data may easily be retrieved using simple off-the-shelf data recovery utilities.

- ***Unauthorized access to information:*** In addition to fabric-based controls, storage arrays can restrict access to authorized servers by using access control lists. This added protection can be easily spoofed but it does provide extra security against unauthorized data access.

A storage security solution should address all three layers of network storage threats while complementing existing security solutions.

Storage Security Solutions

An effective storage security solution that is purpose-built adds an extra layer of defense to secure storage applications and enables customers to meet their regulatory obligations. By exercising best security practices, an ideal solution would protect against the following threats:

1. Eavesdropping of data-in-flight across public MAN/WAN
2. Lost/stolen media (e.g., lost tapes, stolen disks)
3. Data access from an unauthorized user/application/server

Data encryption is the most effective storage security solution to protect against eavesdropping and lost media on disk or tape storage. Any data that is transmitted or sent offsite should be encrypted, especially if it traverses a public MAN/WAN network or is stored at a third-party location. A combination of granular access control with data-at-rest encryption is required to secure data access. Depending on the environment, access control and encryption can be provided by the application, storage application (e.g., backup application), operating system, or a storage security appliance.

Key factors to consider when deciding which encryption solution to select include the security requirements, operational requirements, and performance requirements.

Security Requirements

- ***Strong encryption:*** 3DES or AES encryption is common for securing traffic on IP networks. Look for a storage security solution that offers strong encryption with 128-bit keys or longer. AES-256 encryption offers the strongest commercially available encryption.
- ***Secure key management appliance:*** Data encryption is effective only if the encryption keys are secured. An appliance that is FIPS certified ensures system-level testing. Review the FIPS certification to determine what was actually tested – the system or a component. A system test covers elements of the overall solution, not just a component.
- ***Secure key management tools:*** The system key used for encryption must be entered securely and should always be available for recovery. To ensure a secure system key, look for a solution that has FIPS random number generation with common pattern checking, M of N backup/recovery of keys, and automatic export/storage of encrypted keys.

Operational Requirements

- ***Operational deployment:*** Deployment of storage security solutions differs dramatically depending on the implementation. Some solutions require agents to be installed on all servers, while other solutions require only appliances to be deployed on the storage network. Look for a solution that delivers the security desired without requiring significant operational changes for a simpler and lower cost deployment.
- ***High availability:*** Most organizations build high availability into the network to avoid single points of failure. A solution that supports redundant designs is ideal.
- ***Secure management:*** Look for SSL Web-based management and SSH console access to ensure secure management access. SNMP MIBs should provide link status and other operational information to management systems.

Performance Requirements

- ***Wire-speed performance:*** Most vendors will claim wire-speed performance, although few can actually deliver full duplex wire-speed performance when encrypting data. For disk applications, look for performance measurements when encrypting all frame sizes, supporting full-duplex traffic, and enforcing multiple encryption policies simultaneously. For tape applications, look for performance measurements that run traffic to multiple drives concurrently to ensure the system can scale to support your daily backups.
- ***Very low latency:*** Disk-based storage solutions require minimal latency delays measured in microseconds, to ensure judicious response time of applications. For example, a read operation requires two round-trip transfers to send data across the SAN. A storage security solution with latency much less than 100 microseconds to encrypt primary storage is desirable; otherwise, applications with complex transactions will see performance delay. Tape solutions have less latency issues because data is generally flowing in one direction. However, it's important to measure the ability of the storage security solution to meet your backup window.

Summary

Secure storage networks have suddenly become a key initiative for organizations seeking to meet and maintain regulatory data protection compliance requirements. Despite conventional attitudes that storage networks are naturally secure, recent tape and data breaches have demonstrated the need for improved storage security practices. Not only do storage breaches subject sensitive information to potential exposure, the breaches can also violate regulatory requirements intended to protect sensitive data. Security is an ongoing process. Risks, benefits, and costs must be balanced and managed as new threats emerge. Although storage security planning and training may initially appear time-consuming or burdensome, an effective storage security strategy will yield measurable, long-term benefit to any organization that must protect sensitive data. Over the last few years, storage security technology and best practices have matured to a point where storage security compliance is now easily achievable at low expense and without impacting network application performance.

About the Author

Dore Rosenblum is vice president of marketing at NeoScale Systems.

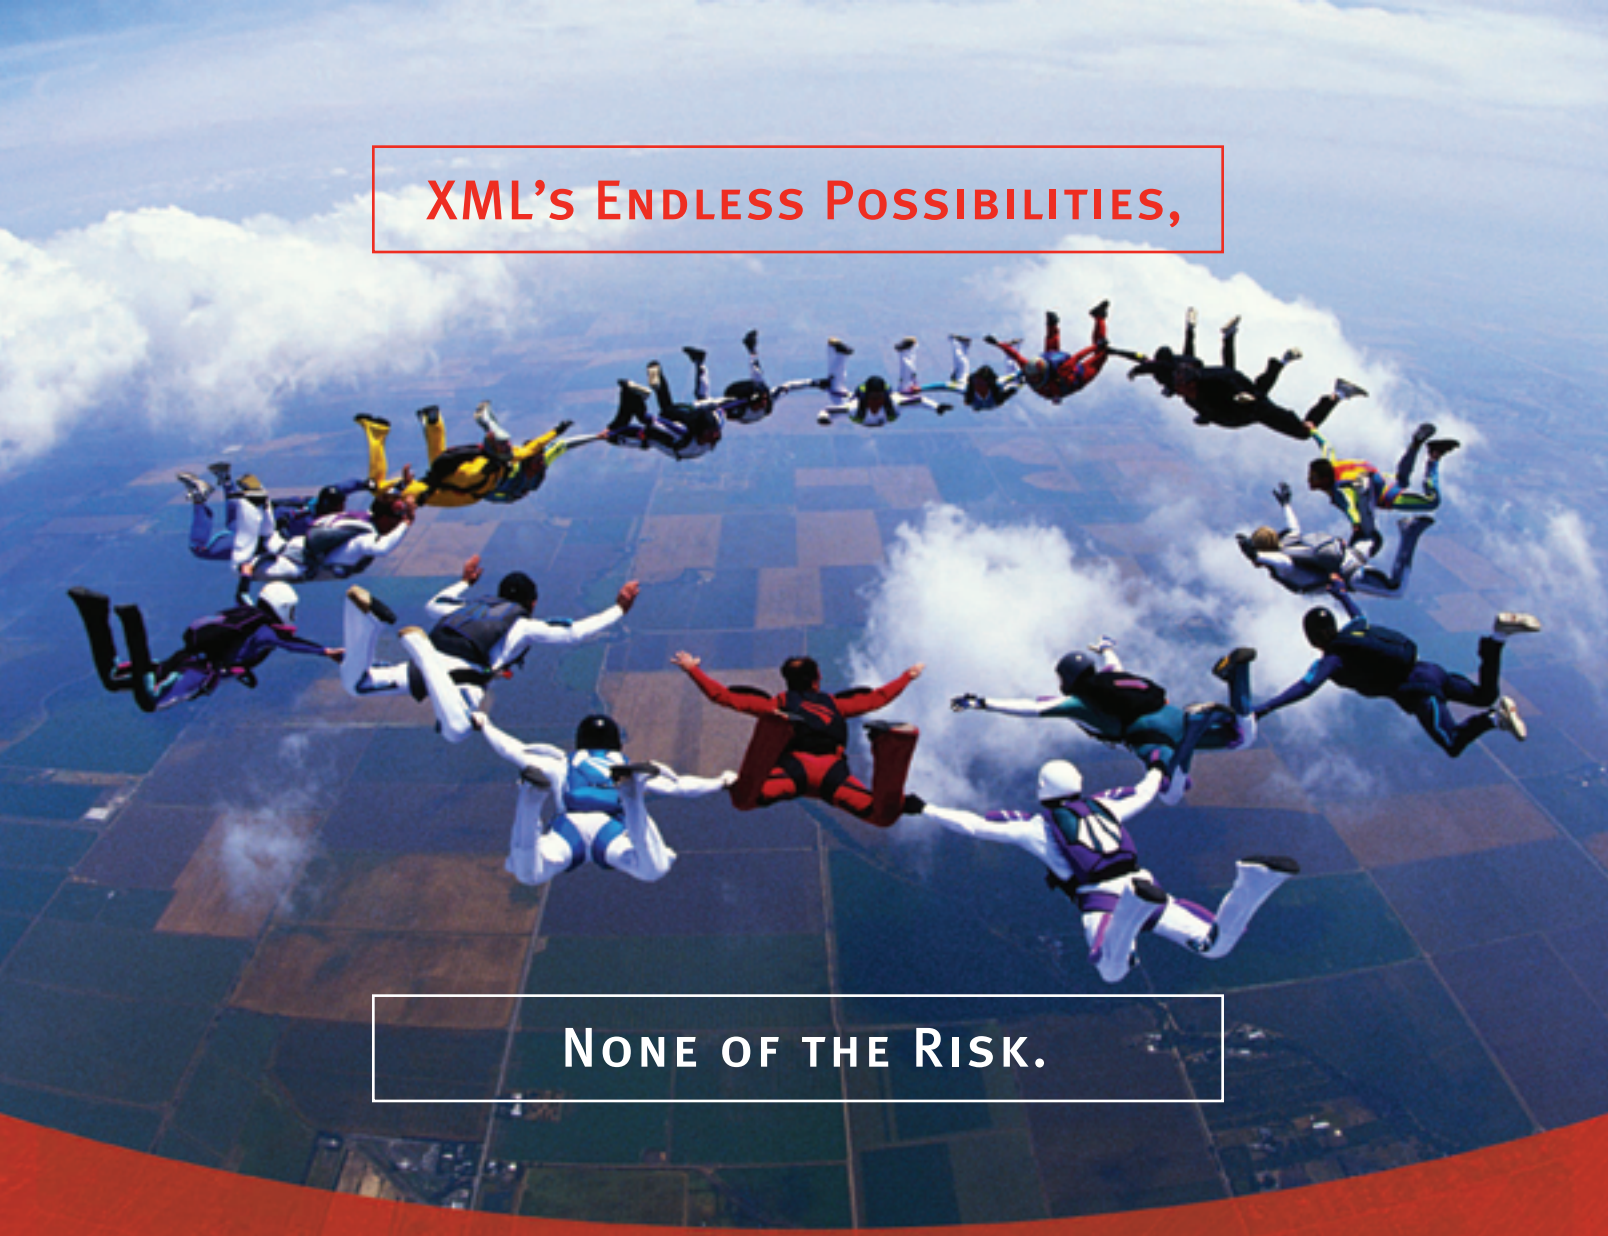

Forum XWall™ Web Services Firewall - Reinventing Security

Security should never be an inhibitor to new opportunity: Forum XWall™ Web Services Firewall has been enabling Fortune 1000 companies to move forward with XML Web services confidently. Forum XWall regulates the flow of XML data, prevents unwanted intrusions and controls access to critical Web services.

Visit us at www.forumsys.com to learn more about how you can take your next leap forward without increasing the risks to your business.

FORUM SYSTEMS™ – THE LEADER IN WEB SERVICES SECURITY

SPEED
ESSENTIAL™ SECURITY
from SafeNet

Now you can have both speed and security.

SafeNet's SONET encryption.
The protection you want, with a lot more speed than you're used to.

When speed is essential, SafeNet is a necessity. We offer the only family of SONET encryption products with a throughput of up to 10Gbps – plus security at the physical, data link and network layers. We give you the highly secure AES algorithm with a 256-bit key length. And SafeNet solutions can secure OC48 and OC192 networks – but will also blend transparently into OC3/OC12, or OC3/OC12/OC48 systems. So if you need protection that runs fast and deep, call SafeNet today and ask about Speed Essential Security. It's where high speed meets high security.

For a free copy of the **Frost & Sullivan white paper**, "WAN Services and Encryption: Protecting Data Across Public and Private Networks," visit *www.safenet-inc.com/hse/14*

Call 1-800-697-1316 to be SafeNet sure.
www.safenet-inc.com/hse/15

SafeNet®
The Foundation of Information Security

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