Fighting Spyware and Adware in the Enterprise 〈8

information STORAGE+SECURITY journal

www.ISSJournal.com

In This Issue:

3 〉 An Introduction to the EiCs

12 〉 E-mail and the Process of Electronic Discovery

18 〉 Key Operational Issues to Consider for Application Firewalls

20 〉 Demand Vigilance from IT Security

22 〉 Algorithm Agility & OATH

24 〉 Security and Storage Granularity

26 〉 NAS and SAN: The Waiter and the Chef

30 〉 Guided Intelligence Analytical Server

32 〉 Virtualizing NAS?

VOLUME: 2 ISSUE: 4 2005

a "Virtual" Storage Revolution 〈4

SELF-PROTECTING STORAGE SYSTEMS

Plus: Airport Security

16 〉 InteliTrac Delivers an Integrated Airport Security System

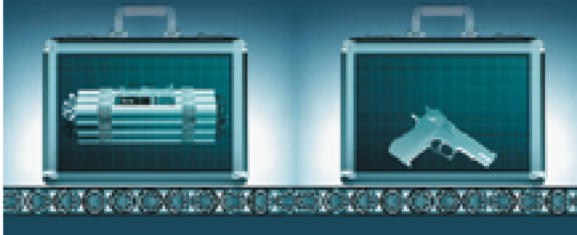

Do you know who's
inside your network?
Spyware

information STORAGE+ SECURITY journal

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

Group Publisher
Jeremy Geelan jeremy@sys-con.com

Advertising

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

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

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

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

Associate Sales Managers
Dorothy Gil dorothy@sys-con.com
Kim Hughes kim@sys-con.com

Editorial

Editor-in-Chief
Patrick Hynds phynds@sys-con.com
Bruce Backa bbacka@sys-con.com

Executive Editor
Nancy Valentine nancy@sys-con.com

Associate Editor
Seta Paparzian seta@sys-con.com

Online Editor
Roger Strukhoff roger@sys-con.com

Production

Production Consultant
Jim Morgan jim@sys-con.com

Art Director
Alex Botero alex@sys-con.com

Associate Art Directors
Louis F. Cuffari louis@sys-con.com
Tami Beatty tami@sys-con.com
Andrea Boden andrea@sys-con.com

Web Services

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

Web Designers
Stephen Kilmurray stephen@sys-con.com
Percy Yip percy@sys-con.com
Vincent Santaiti vincent@sys-con.com

Accounting

Financial Analyst
Joan LaRose joan@sys-con.com

Accounts Receivable
Gail Naples gailn@sys-con.com

Accounts Payable
Betty White betty@sys-con.com

Customer Relations

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

Subscriptions

Call 888-303-5252 or 201-802-3012
www.sys-con.com or subscribe@sys-con.com

Editorial Offices

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

Worldwide Newsstand Distribution
Curtis Circulation Company, New Milford, NJ

For List Rental Information:

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

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

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

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

From the Co-editors-in-Chief

An Introduction to the EiCs

BY PATRICK HYNDS AND BRUCE BACKA

LAST MONTH WE skipped the introduction to summarize some of the things we hope to bring to you in the coming year. This month we are circling back to our backgrounds in the hopes that this will help explain why we might be suited to fulfill the agenda hinted at last month.

Bruce can be thought of as the storage side of the team, though that is a gross oversimplification. It is no exaggeration to describe Bruce as a noted business leader and consultant in the IT industry; he has acted as chief architect, technologist, and project manager for assignments involving large-scale technology and implementation strategies. He 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. Bruce has been responsible for the architecture, implementation, and management of a worldwide client/server networking infrastructure for a Fortune 10 company, with a platform of over 600 servers connecting 10,000 users across 50 cities throughout North America and offshore. In 1994, he founded NTP Software, a provider of business solutions for Windows NT and other platforms. At the World Economic Forum in Switzerland, Bruce was recognized as a technology pioneer. This follows a similar award from the National Computer Conference in 1974 where he was honored as a part of the Dartmouth College team that invented computer timesharing. In his role at the helm of NTP Software, Bruce has been a thought leader in the storage resource management space and has authored many whitepapers and articles on subjects touching on both storage and security. Knowing the trends, as well as why they are trends, in storage is key to guiding the discussions that we hope will bring you as well as help you find the answers to the harder questions.

Hard questions are where security typically enters the picture and Patrick has a long history as a security consultant and has written articles on everything from writing applications that defend themselves to ways to make the hacker pay for coming after your data. Patrick is the Microsoft Regional Director for Boston and the CTO for CriticalSites, a security consultant firm based in New England. Named by Microsoft as a regional director, he has been recognized as a leader in the technology field with an expressed specialization in the field of security. A graduate of West Point and a Gulf War veteran, Patrick brings an uncommon level of dedication to his leadership role at CriticalSites and tends to approach all situations as an infantry commander in the defense. He has experience in addressing business challenges with special emphasis on security issues involving leading-edge database, Web, and hardware systems. In spite of the demands of his management role at CriticalSites, Patrick stays technical and in the trenches, acting as project manager and/or developer/engineer on selected projects throughout the year.

We hope that this introduction shows us as qualified to guide the discussion we hope to provoke in the pages and issues that follow.

— continued on page 33

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

A “Virtual” Storage Revolution

SELF-PROTECTING STORAGE SYSTEMS

BY DAVE THERRIEN

VIRTUALIZATION IS ARGUABLY one of the most prominent “buzz” technologies in the computing industry today – with more than $4.2 billion in reported deal values for M&A activity since December 2000. It’s a key component in just about any next-generation architecture due to its ability to break down the sheet metal bonds between systems and base resources and applications on logical instead of physical parameters. It has also signaled a major shift in the industry toward software value and hardware commoditization.

Virtualization when used in a storage context is assumed to be virtualization of primary storage capacity for the purposes of sharing. But there’s a trend afoot that’s taking virtualization to the next level. Along with advances in disk technology/performance, component cost reductions, and emerging technologies in grid computing, virtualization is helping to drive some traditional independent data protection applications such as backup, disaster recovery, and hierarchical storage management (HSM) toward evolutionary obsolescence.

This article will examine the mechanics – and more importantly the business potential – of virtualizing data protection and integrating it more tightly with primary storage.

Virtual Standstill Leading to Virtual Extinction

Over the last 50 years, many virtualization concepts have been applied to data storage systems. So why is it that IT organizations are still waiting and searching for data storage and data protection systems that are more shareable, less manually intensive to install, configure, and operate, easier to scale, and easier to store and protect data with?

Because the storage vendors have been at a “virtual standstill.”

Consider today’s reality.

Storage Area Network (SAN) administrators must still manually create RAID Logical Unit Numbers (LUNs) out of groups of individual disk drives, and manually allocate these to specific servers through arcane Fibre Channel (FC) switch and host bus adapter (HBA) command line scripts. System administrators need to manually configure host-based FC failover systems, manually create host-based volumes, and create or expand file systems. And all this just to deliver primary disk storage to servers.

Virtualization in the context of data protection is almost non-existent. When new volumes or file systems are created, the backup administrator must manually create a backup configuration that includes specifying client schedules, file filters, media sets, retention and rotation schedules, etc. Backup administrators must manually eject multiple tapes from jukeboxes everyday and send them by truck off-site for disaster recovery protection and safekeeping. No virtualization here.

For companies that have invested in cold-site disaster recovery schemes, a Disaster Recovery Administrator must manage two or more sites of similar or identical equipment and keep them up-to-date with each other for years, even decades. The likelihood that the disaster recovery systems in a cold site are properly prepared for disaster over this time span is extremely low. For warm or hot-site failover, more complex replication systems and networks need to be in place to support fast failover. Replication of current data doesn't obviate the need for performing regular nightly and weekend tape backups, so IT organizations have to do both.

When servers run out of available storage capacity, the SAN administrator and a system administrator must come to the rescue to immediately allocate more storage to the system, whatever the time of day. Alternatively, the archive administrator would have to get involved with determining which collection of files would be considered stale so that these could be written to tapes and deleted from the server, thereby freeing up space for new data.

Archive administrators occasionally inadvertently archive data that is live and very active, causing live applications to fail without access their data. Archive tapes have to be shipped to off-site storage, just like backup tapes. HSM systems and software are costly and complex to manage. And when HSM is combined with magnetic tape as a lower-cost tier of storage, it creates data accessibility/availability exposure because of tape's inherent lower reliability and poor access times, especially when tapes are stored off-site.

What happens when backup, archive, or HSM tapes are lost or stolen? Most tapes are written in standard formats that encourage readability among multiple backups and archiving applications, making it even easier to gain unauthorized access to data with almost any application.

In the data management software space, backup, replication, archiving, and HSM packages are often unaware of each other's tape sets and replicas, which leads to inefficiencies in IT operations time, as well as capital equipment costs related to as much as 15x over-replication of the same files. This hasn't changed in 50 years.

So, have we achieved the goals of a virtualized data storage environment yet? With the technology that's been available to date, we're not even close!

Virtual Tour of the Ideal Storage System

Let's stop for a moment and just visualize the elements of the ideal storage system:

Imagine a system that provides primary storage to your clients and applications. In a virtual environment, clients and applications would be completely unaware of the specific storage system that gives them storage capacity – and the location of their data could transparently change from time to time to react to I/O performance bottlenecks, hardware failures, and even site disasters!

Imagine a system where all of your data is continually and transparently protected, and its history perfectly maintained over months, years, even decades, both locally and at an off-site location – all without operator intervention; tapes and trucks; independent backup servers, backup software licenses, tape drives, tape jukeboxes; and stacks and stacks of tapes.

Imagine a system where complete recovery from a site disaster could be initiated from anywhere with access to a Web browser. Within minutes of a complete site disaster, all clients and applications regain access to their critical data from live systems at a surviving site.

Imagine a system where clients and applications never get an "out of disk space" error message. In addition, the system automatically knows which files are inactive, and it would automatically migrate inactive files from a more costly, high-performance disk storage tier to a more cost-effective disk storage tier. All data, regardless of its location among virtual tiers of disk storage, would remain transparently accessible to all clients and applications.

Imagine a system that can expand its storage and protection capacity just by plugging more servers with integrated disk storage into a standard gigabit Ethernet switch. There would be no need for complex SAN LUN allocation, HBA and FC switch zoning, or volume and file system management.

Imagine a system that automatically checks and corrects every version and replica of every backup file it retains. And as the amount of data to backup grows, both the processing power and storage capacity of a storage grid expands, making it as easy to check and correct 100TB of data as it is to check and correct the first terabyte of data. This automated checking and correcting feature equates to checking and correcting all of the tapes in your tape vault on a regular basis! Checking all of your tapes regularly is time-prohibitive, and actually decreases the reliability of tape drive heads and the long-term readability of the data on these tapes. Even if you know which tapes had a data integrity problem, you'd still have no means of correcting the damaged or lost files.

Imagine a system where the integration of data storage and data protection, and the corresponding simplicity of a single management interface allows tasks that today are performed by one or more SAN/Network Attached Storage (NAS) administrators, server administrators, backup administrators, archive administrators, and disaster recovery administrators to be performed by a single Storage Administrator.

Virtualization is a key component to realizing the ideal storage system.

Virtual Reality – Self-Protecting Storage

The ideal storage system is a reality. It exists in a revolutionary breed of storage systems called Self-Protecting Storage systems. These systems are only now becoming available because of major advances in technology and research. CPU power is virtually free, disk storage is taking a more prominent role in protecting data, and IP MAN/WAN network cost/performance is now affordable to even mid-market companies.

Architectures that have been conceptualized and prototyped at major universities and corporate research labs worldwide over the past 10 years are being commercialized by innovative emerging data storage companies. Advances in grid computing are being leveraged to create systems that are self-discovering, self-configuring, more shareable, and efficiently scaled. But it's virtualization that is the key technology at the core of Self-Protecting Storage systems. Consider the the sidebar *Virtualization Concepts and How They Map to Tangible Self-protecting Storage Benefits*.

Virtual Monopoly – Game Over

Disruptive technologies that unseat established products will continually be developed by innovative emerging companies. IT organizations that are forward thinking and looking for relief from the significant cost and complexity in adequately storing and protecting data today are already reaping the benefits of Self-Protecting Storage systems.

Simply stated, virtualized Self-Protecting Storage systems are architected from the outset to reduce the cost, operational complexity, and associated risk of storing, protecting, restoring, and recovering data while increasing the availability and access to all data.

From a business impact, virtualization can achieve the following benefits in the storage and protection of data:

- Greatly reduced operational and capital costs
- Seconds, not hours, to restore lost/deleted data
- Minutes, not days, to recover from site disaster
- Transparent access by clients and applications to all data, all the time

From an IT operational impact, virtualization can eliminate the following data management tasks and their associated operational costs:

- No more weekend full-tape backups
- No more slow, unreliable restores from tape
- No more tapes to rotate on-site and off-site
- No more early morning "file system full" alerts
- No more guessing about which files to archive and when
- No more complex systems and procedures for managing a site disaster
- No more complex procedures to manage dozens of independent storage, backup, and archiving products
- No more tape drives or jukeboxes to purchase and repair
- No more tapes to verify and refresh

Virtualization Concepts and How They Map to Tangible Self-Protecting Storage Benefits

Virtualization Concept	Without Virtualization...	Self-Protecting Storage Virtualization Benefits
Automatic pooling of disk storage resources	LUNs to allocate, volumes to create, or file systems to expand Tape media, media sets, etc.	Repositories of backup/archive data are automatically expanded as new servers and storage are added to the pool
Automatic migration of data with two tiers of disk storage	Storage administrator guesswork about what files to archive to tape Storage administrators must react immediately to "out of space" conditions	Inactive data is identified and migrated to lower-cost storage File systems never run out of space
Distributed File System and Global Namespace	Filer share migration is disruptive to potentially hundred of clients	Clients access NAS shares without actually having to be concerned about the specific filer that stores the file
Self-healing backup/archive data	Impossible to check all backup data with existing tape-based storage systems.	The data from any failed server/storage component is automatically re-created on similar surviving components in the "storage grid" Continual integrity checking and correction of all backup/archive data
Virtualized backup processes	Inefficient weekend full backups Resource-consuming synthetic full backups Cost and complexity of D2D2T backup solutions Tapes to duplicate and send off-site by truck	Incremental-only backups are efficient Continual and transparent synthetic full backups virtualize all backup data in preparation for a full and complete system restore. All backup data transparently replicated off-site for fast site disaster recovery

Self-Protecting Storage Is Virtualization

The storage product industry has been tackling challenging data storage issues by applying incremental changes to each data storage and data protection component. This has been ineffective in significantly reducing the operational cost of managing data storage systems. Self-Protecting Storage systems represent the first technology in the last 50 years to leverage virtualization for all aspects of data storage and data protection.

About the Author

Dave Therrien is the founder and CTO of ExaGrid Systems (www.exagrid.com). He is also the author of "Self-Protecting Storage – Simplifying Your Data Storage Infrastructure" http://www.exagrid.com/pdfs/Self-Protecting_Storage.pdf.
dtherrien@exagrid.com

Fighting Spyware and Adware in the Enterprise

HOW TO KEEP YOUR IT SYSTEM PROPERLY "SCRUBBED"

BY SARAH GORDON

ADWARE AND SPYWARE – they may be as hard to define as they are to eradicate. But there's one thing just about everyone can agree on: what started as a minor annoyance has ballooned into a full-blown corporate headache.

How big of a headache? According to the most recent edition of the Symantec Internet Security Threat Report, adware is a growing concern. Between January 1 and June 30, 2004, adware made up 4 percent of the top 50 malicious code reports to Symantec. Between July 1 and Dec. 31, it made up 5 percent of the top 50 reports. As for spyware, the most common program during the second half of 2004 was Webhancer, which alone represented 38 percent of the top 10 spyware programs reported.

This growing concern about adware and spyware has put enterprises at greater risk for decreased productivity, more help desk calls, loss of privacy, and potential legal liability. Analyst firm META Group estimates that cleaning infected clients can represent 20 percent or more of IT help desk efforts.

A 2005 Forrester Research Inc. survey of IT decision-makers found that 40 percent of respondents didn't know how many systems in their organization were infected with spyware. Those who could measure the number of systems infected with spyware found that about 20 percent of systems were infected, and the number is growing rapidly.

Small wonder, then, that adware and spyware have surpassed spam and identity theft as the threats that security managers are most concerned about, according to Forrester. The research firm predicts that 65 percent of companies will either purchase or upgrade anti-spyware software this year, making it the most popular security technology of 2005.

Methods of Installation

Some organizations justify the use of adware as a way of providing services while lowering costs to customers. This is particularly true of software that is made available for users to download for free. These "freeware" programs usually require the user to agree to a EULA (end user license agreement). But some EULAs can be complicated and confusing – to the point that the user is unable or unwilling to read and understand the terms and conditions before agreeing to it. As a result, adware that is bundled with the desired software gets installed without the user's knowledge.

Adware is also often installed through the user's Web browser. This can be done through pop-up ads offering free software to download. The pop-up offers the user a choice of clicking "Yes" or "No" to accept or reject the offer. In reality, though, clicking anywhere on the ad results in the download of adware. Browser-installed

Some Definitions

Adware:
Consists of programs that display advertising content on a user's monitor, often without the user's prior consent or explicit knowledge. It is usually, but not always, presented in the form of pop-up windows or bars that appear on the screen. Adware is not always a security risk. In some cases, it simply delivers an advertising message, but this is not always the case. While much adware is benign, some forms of adware can compromise data. If attributes of a security risk include the compromise of the confidentiality, availability, or integrity of data on a computing system, some forms of adware qualify.

Spyware:
Refers to stand-alone programs that can secretly monitor system activity and relay the information back to another computer. In some cases, spyware may be legitimate programs that are employed by corporations to monitor employee Internet usage. However, it may also represent less legitimate applications. Spyware programs can be surreptitiously placed on users' systems in order to gather confidential information such as passwords, login details, and credit card details. This can be done through keystroke logging and by capturing e-mail and instant messaging traffic. Because spyware can capture sensitive information before it is encrypted for transmission, it can bypass security measures such as firewalls, secure connections, and VPNs. Spyware is a particular concern because of its potential use in identity theft and fraud.

The dividing line between adware and spyware, experts say, is intent. Programs that install themselves on a user's system without permission, avoid being detected and removed, and capture and transmit personal information without a user's permission or knowledge have crossed the line into spyware.

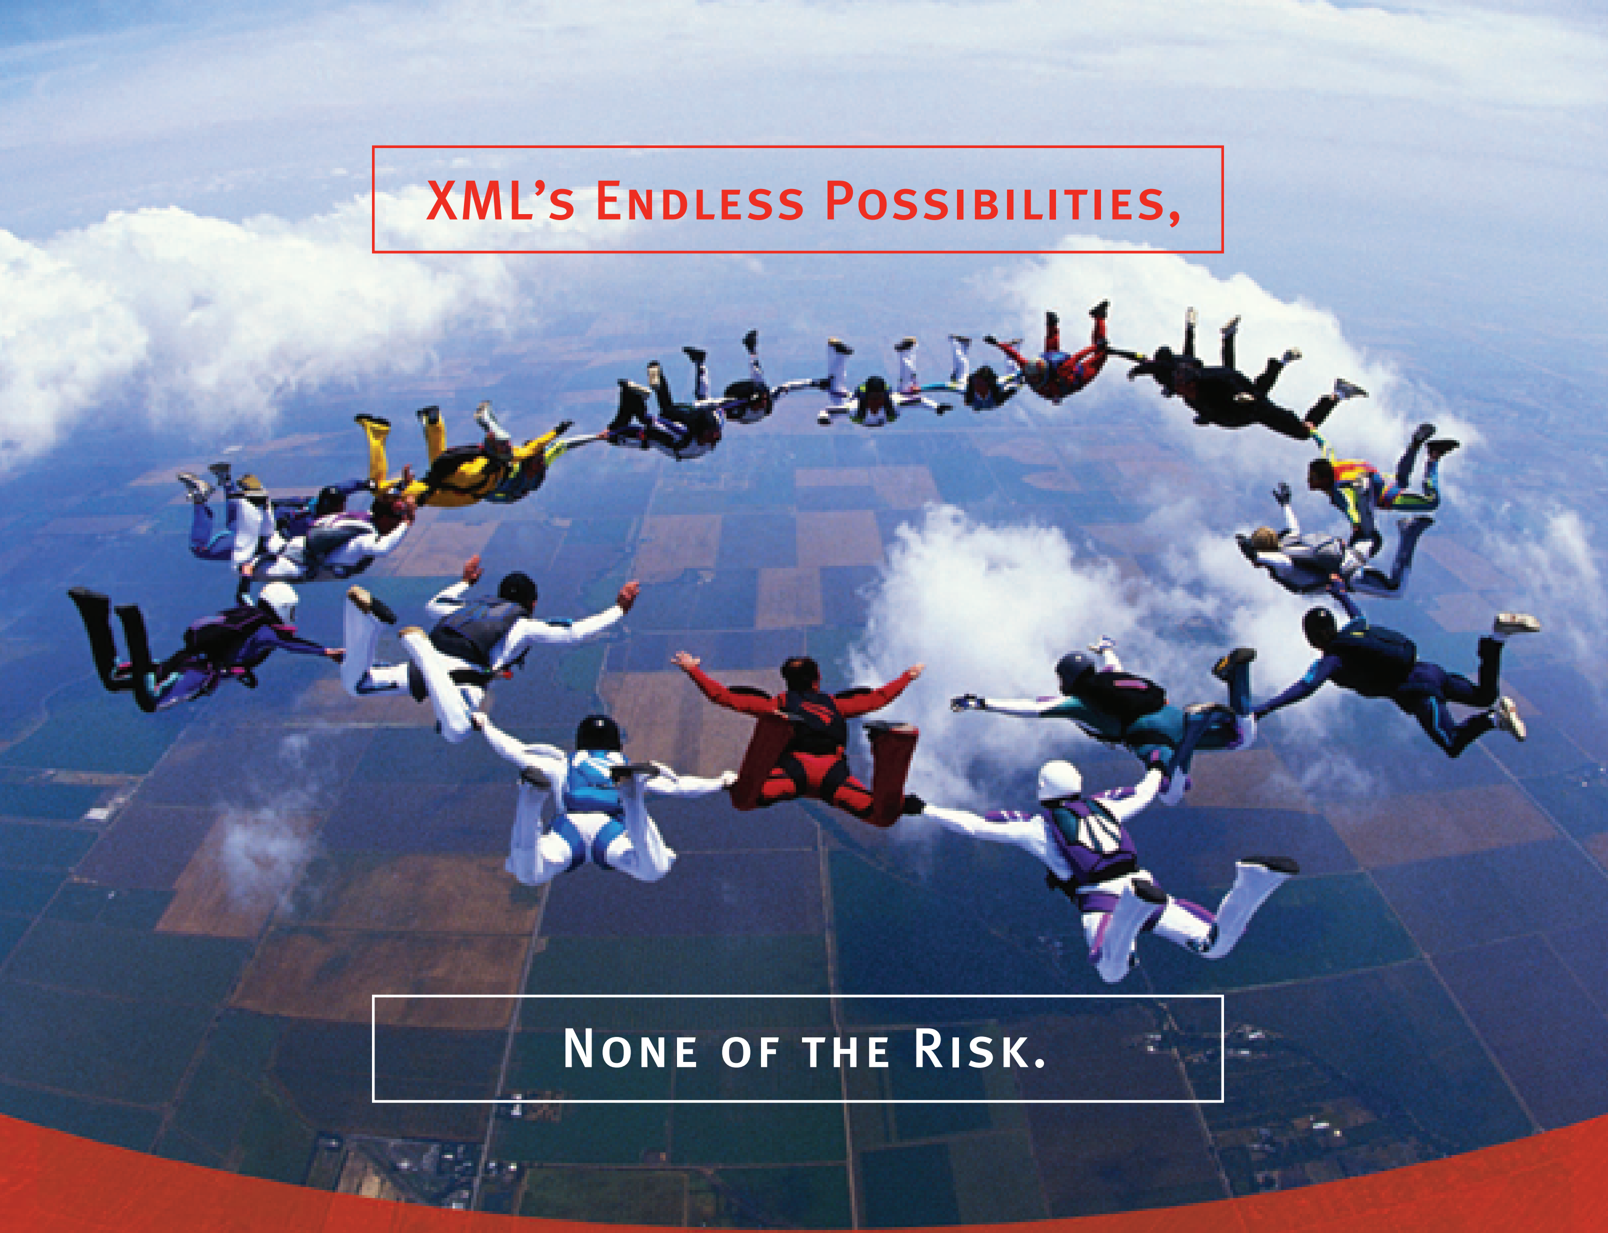

Forum XWall™ Web Services Firewall - Reinventing Security

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 adware may also be installed through ActiveX controls or browser helper objects (BHOs). BHOs can provide spyware with a wide range of functionality, including the ability to download program updates, or log and export confidential data. During the last six months of 2004, three of the top 10 reported spyware programs used BHOs.

Some adware programs hijack a user's browser and redirect searches. A program may redirect a search by replacing the default search engine or by replacing "404 page not found" messages with internal search queries. This is not only misleading for the user but also represents a security risk, as the redirection may result in the user downloading malicious code from the new page. Five of the top 10 adware programs reported in the last six months of 2004 hijacked browsers. Spyware can also hijack browsers.

If users' browsers are enabled to accept cookies and ActiveX files, as many are, unwanted code can be installed in the background without their permission or knowledge. Spyware also travels on fake messages telling users their systems need to be tuned up, or similar instant message screens that appear to be sent by a system administrator.

Keeping Trouble Out

Like viruses and worms, adware and spyware are moving targets, and enterprises can best protect themselves by deploying multiple defenses – at the desktop, the gateway, and across the enterprise – and by educating users on what behaviors will best keep the spies where they belong: out in the cold. The most effective way to reduce risks from programs such as spyware and adware is to use a complete security solution that deals with a wide range of threats. In particular, enterprises need a solution that categorizes programs according to their functionality and allows them to choose an acceptable risk level. Integrated technologies (antivirus, firewall, and intrusion protection) should work together to provide defense in depth. For example, while an antivirus solution works to protect a system against spyware, a firewall allows an organization to create a list of recipients of personal information and to block unwanted advertisements. Furthermore, when a firewall detects that an application is trying to establish an outbound network communication (as a spyware program would to relay information to the outside world) it should automatically close the port and prevent the transmission.

Combating spyware and adware, like combating viruses and malicious code, requires a solid solution and a dedicated research and response mechanism to track new spyware risks and provide timely updates as the threat landscape evolves.

> **"Adware and spyware have surpassed spam and identity theft as the threats that security managers are most concerned about**, according to Forrester...enterprises need a solution that categorizes programs according to their functionality and allows them to choose an acceptable risk level"

Other issues to consider: the number of spyware definitions supported by a particular solution, the process used for finding new spyware programs, and how the definitions are updated.

To strengthen their defenses, businesses should also consider implementing additional security precautions like securing encrypted Internet connections, implementing more restrictive Web browser settings, and disabling the acceptance of third party cookies.

In addition to the use of strong technologies, there are policy measures that can help organizations reduce their risks. For example, make sure that you know and trust the authenticity of any software before you download it and install it. Read the EULAs of software programs to make sure you know what you are getting, and make sure that you understand, and agree with, the program's functionality. Examine EULAs carefully to make sure they are in agreement with your security policy. Also, as some spyware is installed using ActiveX controls, consider requiring a prompt for ActiveX to execute within Web browsers.

The Federal Trade Commission warns: "Before using a file-sharing program, you may want to buy software that can prevent the downloading of spyware or help detect it on your hard drive." Due to the breadth of security threats and risks, it is vital that organizations heed this warning and use security products that can not only deal with spyware and adware, but the entire breadth of Internet security threats. Antivirus and firewall products allow users to protect themselves from malicious code such as viruses and Trojans, as well as expanded threats, which include spyware and adware.

Summary

Spyware and adware infections have become a top concern for IT operations as well as security managers. While much of this code is benign, some is not. Even if a security risk isn't present, cookies and pop-ups can cause significant performance and productivity problems. Enterprises are encouraged to follow the recommendations in this article to keep their systems properly "scrubbed." ■

About the Author

Sarah Gordon, senior principal engineer of Symantec Security Response, is an expert on the psychology of virus writers. She serves as a director for The WildList Organization International, an independent source of information for "in the wild" viruses. She is also technical director at EICAR, an international non-profit group that combines numerous organizations to unite efforts against the writing and proliferation of malicious code, computer crime and fraud, and exploitation of personnel data on computers or networks.

sarah_gordon@symantec.com

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.™

I Have to Show Them What?!

E-MAIL AND THE PROCESS OF ELECTRONIC DISCOVERY

BY DENISE REIER

WHILE E-MAIL MAY be a killer app, poorly archived e-mail can kill a business. During a recent prescription drug antitrust case, the plaintiff demanded a discovery search of 30 million pages of e-mail stored on the defendant's backup tapes for names of particular individuals. The defendant suggested the plaintiff shoulder the cost of compiling, formatting, searching, eliminating duplicates, and retrieving the requested e-mail. Sadly, the defendant lost the argument and the court found the burdensome and expensive discovery process was the defendant's problem because of a bad e-mail retrieval process. The defendant paid through the nose.

Court cases routinely approve discovery motions to sift through electronic documents, especially ubiquitous e-mail. The consequences of NOT having the information available or being able to access it in a reasonable amount of time are severe. Despite this, few companies have enforceable records retention policies and fewer still have the technology tools needed to support it.

Another problem is keeping e-mail you don't legally have to keep, as Microsoft found out. The smoking gun e-mail surfaced during the discovery phase of Microsoft's antitrust trial, even though it was AOL's e-mail, not Microsoft's. The Justice Department found one of Gate's own e-mails with the undying line: "We have to make sure that we don't allow them to promote Netscape."

Electronic discovery isn't just for enterprise corporations. Mid-sized companies frequently experience legal discovery, so if it hasn't happened to an individual business it almost certainly will. So how can organizations best prepare for electronic discovery? By balancing risk against cost – in essence; it's establishing policies and capabilities for efficiently accessing secure archives without breaking the bank.

Managing the Archive for Discovery

Companies should begin by establishing and enforcing retention policies, including policies against destroying or altering data potentially relevant to discovery motions. This goes double for destroying or altering data after discovery or litigation starts. Seems obvious, but over the last few years we've heard about executives being indicted for deleting messages pointing to insider trading. Not only did the federal investigators recover the deleted messages, they tacked on additional serious charges. Archiving procedures must support evidentiary measures and record retention policies must be in writing and enforced with a method to prove regular enforcement.

This is a tall order and its success depends on a cost-effective technology to support electronic discovery for messaging files. This strategy hinges on two major elements: managing cost and managing risk. A balance between the two yields a cost-effective technology for managing messaging archives and enforcing records retention.

Managing Cost

Managing cost includes reining in storage costs, improving operational efficiency and company productivity, and decreasing retrieval/discovery costs.

Until recently companies were limited to first-generation archiving applications. These applications backed up incremental or full copies of data to backup media. With no way to manage duplicate copies of data and an awkward and time-consuming retrieval process, archiving would complicate discovery procedures and significantly increase cost. In fact, storing e-mail alone often represents over 40% of an organization's storage costs because of both the sheer volume of e-mails and the multiple copies of messages that are retained.

New archive applications are engineered to compare e-mail messages, record and validate the original, and eliminate duplicates across multiple messaging servers. E-mail archives can now take up a fraction of the storage space previously used, allowing companies to shrink the amount of backup media, backup windows, and retrieval time. These archiving apps are also capable of rapid searches based on a number of parameters, letting organizations quickly retrieve detailed e-mail subsets in response to a discovery demand.

Managing Risk

Managing risk includes keeping a complete archive, enforcing retention policies, proving authenticity and evidentiary weight, and maintaining security and privacy.

Poor archive retrieval systems are extremely time-consuming and costly due to their big-dump approach – archive

BARRACUDA
NETWORKS

Elements of Managing Costs

Element	Capability
Storage costs	• Single-instance message storage reduces storage costs up to 80% • Consolidates e-mail from multiple platforms like Exchange and Notes
Operational efficiency	• Centralized archive ensures uniform retention practices • Simplified e-mail management including smaller backup windows, higher availability, and automated disposal
Organizational productivity	• Controlled user access to archive allows for quick search and retrieval • IT spends less time recovering user mailbox data • Moving personal message stores to the network frees up hard drive space and lowers liability
Retrieval/discovery costs	• Reduced time and cost in fulfilling discovery requests • Eliminates duplicate messages to reduce retrieval costs • Exports results to a portable format for review by an outside party

Elements of Managing Risk

Element	Capability
Complete archive	• Real-time message capture for a complete and unaltered archive • Audits and reports guarantee archival integrity • Archives all relevant messages including e-mail, attachments, and IMs • Retains all user information, such as aliases and distribution lists
Reinforcing policy	• Organization controls record retention and disposal • Monitors and audits to ensure internal compliance
Authenticity and of evidentiary weight	• Tamper-proof archive aligns record authenticity to rules evidence • Audits and reports on all archive views and operations on all levels of access rights • Fully indexed records return complete and accurate discovery results • Records archived in original form meets "legal and true" copy requirements
Security and privacy	• User can only view individual archived e-mail • Audits administrator archive deletions

everything and hope for the best. This results in huge volumes of badly indexed messages, an awkward and labor-intensive retrieval process, and no way to prove the archived media's integrity.

Next-generation e-mail archive management can audit and report on all archive access and operations, automatically run retention schedules, and capture and index all messaging data, including attachments and IMs. It should also allow for real-time data capture, monitor user information and support privacy by protecting user access.

Real-Life Discovery

Is it worth deploying a next-generation archiving product? Previously, a mid-sized company with 1,500 employees and $350 million in annual revenues had to retain 95% of its e-mails. The company fielded three to five discovery requests a month and spent a cool $15,000 to $20,000 per discovery request. The grand total of meeting discovery demands was approximately $1.2 million a year.

After deploying next-generation archiving tools, the company took that $1.2 million down to just $57,000 a year to fulfill discovery demands. The archiving software speeded discovery retrievals, decreased storage costs, centralized e-mail archive management, and provided audits and reports to prove that archives met evidentiary requirements.

Being unprepared for electronic discovery can be disastrous, resulting in thousands of hours of employee labor and millions of dollars in consulting and legal fees. It's vital that companies develop, audit, and enforce electronic discovery policies and invest in supporting technology from companies like EMC. Protecting, auditing, and producing information on-demand not only allows organizations to protect themselves during litigation, but also supports strategic business goals by cost-effectively managing critical business information.

About the Author

Denise Reier is vice-president of product marketing for messaging solutions in the EMC Software Group at EMC Corporation.

BLACKHAT

InteliTrac Delivers Integrated Airport Security System

ITS CHILEAN "SECURITY-IN-THREE-SECONDS" TEST BED

BY GARY SIMPSON

WITH HEIGHTENED SECURITY requirements since 9/11, airports have a major responsibility to distinguish and identify law-abiding travelers from wanted criminals and legitimate security threats. As millions of commuters, tourists, businessmen and women, and students fly across the globe everyday, security officials are faced with a daunting challenge. Modern security systems must be quick and efficient enough to be non-obtrusive, but powerful enough to sort through a huge database of records and images effectively to identify criminals – those on the "no fly" list – all while protecting privacy rights.

With immigration continuing to be such a priority in the U.S. inbound international flights must be as secure as possible. In some ways, airport security has become the centerpiece of global security, whether in the U.S. or at international airports abroad. Take for instance the Chilean National Airport in Santiago, an international airport in a foreign country that has adopted stiff security policies and implemented a strong security solution.

InteliTrac has led a major initiative in airport security at the airport. In partnership with AssureTecSystems and Cognitec Systems, the technology vendors have come together to develop a complete biometric and database management solution that addresses forgery problems with travel documents and movement of criminals across borders, while speeding up airport queues. The system is supported by a powerful Unisys Itanium 2-based ES7000, running SuSE Linux Enterprise AS, and is capable of performing biometric searches in less than 1.5 seconds according to tests involving database and memory response and retrieval times. It beat expectations resulting in a 250% increase in rates over the initial estimate. First deployed at the Chilean National Airport, the partners are working to bring the solution to a worldwide market.

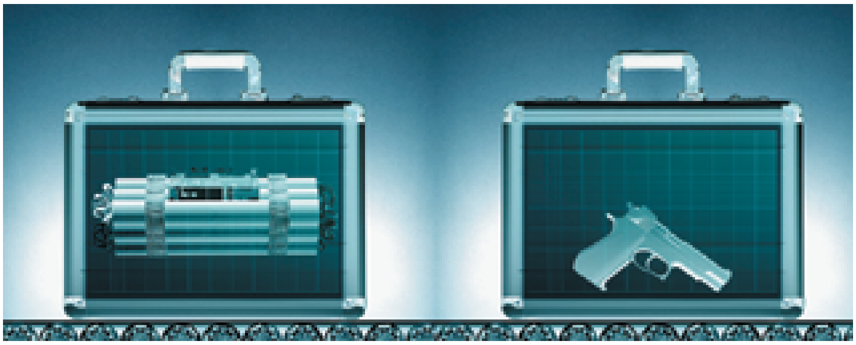

InteliTrac creates integrated hardware and software solutions for standard computer networks. These networks process data 10 to 1,000 times faster than previously possible for most applications on the same computer networks. This is made possible by manipulating all available data in Random Access Memory through an indexation process in a neural network system.

Built on InteliTrac's groundbreaking high-speed computing system, the Policia Investigaciones de Chile (PICH), the Chilean equivalent of the FBI, implemented IdentiPort, a system at the airport that does real-time image verification, fingerprint verification, and facial recognition functions.

The three technological components of the Chilean airport security solution include InteliTrac's Biometric Memory Resident Database Management System and AssureTec Systems's advanced document authentication firmware fully integrated within the IdentiPort system. Cognitec Systems provides IdentiPort's baseline facial-recognition technology. The solutions suite, which is currently live in Santiago, is organized into different conceptual security rings: document authentication, face-recognition delinquency check, and fingerprint check.

Travelers in line only need to follow a few easy steps during InteliTrac's IdentiPort security passage. First, the traveler's passport is scanned to determine authenticity by checking it for security marks. Next, the traveler's face is scanned and compared to the image on the passport. The facial-recognition technology uses both 2D and morphed 3D facial images. IdentiPort employs a live 2D video camera to capture images and compare them to an 80 million-image database. With a search rate of more than 20 million images per second, the system combines speed with a high rate of accuracy.

Simultaneously, the image is crossed with a Criminal Pictures Database (CPD) using the face-recognition algorithm. If a traveler is flagged and the previous checks haven't been conclusive, a fingerprint is scanned and cross-referenced with civil and criminal databases. It uses Identiport's fingerprint verification process, with ultrasound-based fingerprint scanning and standard AFIS lookup. Each IdentiPort unit is fully scalable. It will operate with data from 4,000 to

36,000 individual images, with up to 500,000 parsed images, for default standalone operation.

How long does the check-in process take? With the biometric security solution running on the Unisys ES7000 server, it takes just a half-second to do a one to 750,000 search (1:750,000). The real-time response makes the kiosk queues shorter and is more convenient for travelers. Most importantly, the system helps keep travelers safe. If you aren't a known criminal flagged for a fingerprint check, you can expect to get through the kiosk and on your way in less than three seconds.

Because biometrics is only the first application of its technology, InteliTrac sought a cost-effective platform that supported its ambitious plans for expansion. Its core technology depends on getting the maximum amount of addressable memory and processes such as facial recognition require a great deal of CPU power. A server platform with large amounts of memory and processing power and very high reliability was a must. After evaluating several competing platforms, InteliTrac selected the Itanium 2-based Unisys ES7000 running Novell's SuSE operating system. The solution includes Unisys's Server Sentinel system management software, which provides manages the environment from a single point-of-control.

Officials at the airport have been pleased by InteliTrac's efficiency and air travelers report feeling an added sense of safety and security when traveling through the InteliTrac airport security suite. InteliTrac CEO Marc Gunderson has praised the implementation and the real-time response that the solution offers. Gunderson said, "We're proud of the successful introduction of the IdentiPort in Chile, and we look forward to working together with our Security Alliance partners to support PICH in the upcoming phases of the project. By implementing the system, the Chilean government has enhanced safety at Santiago Airport and placed itself at the leading edge of international airport security."

So, where does InteliTrac expect this technology to go in the future? The company is hoping that the Chilean implementation spurs interest from airports around the world. However it is not ruling out other avenues of distribution. Corporations and organizations of every description require fast, precise, cost-effective, real-time data processing for all sorts of applications, security-related and otherwise. InteliTrac knows this, and supports supercomputing at the cost of today's network systems: speed, stability, instant data fusion, instant disaster recovery, built-in expanded security – with little use for a hard drive. For organizations looking for this type of real-time capability, InteliTrac can deliver.

About the Author

Gary Simpson serves as chief operating officer of InteliTrac, a computer industry innovator and pioneer in the practical application of high-speed, memory-resident databases. Founded in 2002, InteliTrac created a breakthrough system of hardware and software solutions that enable computer systems to work from resident memory instead of the hard drive. The result is data processing at speeds once reserved for supercomputing, and the new ability to share virtual memory over a server. The technologies have countless applications across a diverse range of markets, including security, finance, medical, and business intelligence.

Key Operational Issues to Consider for Application Firewalls

EASING THE ADOPTION OF NEW TECHNOLOGY

BY MARK KRAYNAK

THE OPERATIONAL IMPACT of deployment is a primary inhibitor to the adoption of new technology in many companies.

Application and database security is a new topic on the minds of many security groups. A key challenge in evaluating alternative solutions is estimating the cost and time to deploy and manage them.

However, some issues are difficult to anticipate because they emerge only in a broad deployment, while most evaluations are done on a smaller scale than that of the actual deployment. With application security products, the impact of these issues is heightened since many Web and database applications directly affect business operations and revenues. In fact, operational pressures are the top reason cited by managers of unsuccessful firewall deployments.

The following list describes what key deployment and operational questions you should ask your vendor and your project team to help anticipate the issues that might emerge only in a broad deployment, but which affect the ultimate success of your application firewall project.

Does the application firewall protect everything you are trying to protect?

Secondary questions:

What are the key elements of the system you are trying to protect? Does the product provide protection for all components?

Most business applications comprise at a minimum a "front-end" application (typically a Web application) and a "back-end" database. Increasingly, Service Oriented Architectures (SOA, also referred to as "Web Services") are being used, primarily for integration between applications and application components. Further, each of these parts typically runs on a standard server platform (i.e., an Apache or IIS Web server) and operating system (e.g., Linux or Windows).

All of these elements of the business application represent a path for attackers to compromise the system. However, many of the products in the application security space are focused on protecting only one element (i.e., only Web applications or only databases). In such cases, solving the whole problem, protecting the business application and its associated proprietary information, can require multiple different security devices. From an operational perspective, this implies the need to train personnel on multiple systems and increases the administrative burden by requiring separate management of the different security systems.

Does the application firewall require changes to the network infrastructure as part of its deployment?

Secondary Questions:

What's the impact of deployment on your IP addressing scheme? On the routing scheme? Are any DNS changes required?

The first task in deploying any application firewall is setting up the network connectivity. This isn't always as easy as it sounds.

Some application firewalls terminate user sessions to get access to the application traffic for inspection. These devices then open separate connections to the destination server. In these cases, traffic must be redirected to the application firewall (implying re-configuration of the network routers or switches).

Often, application firewalls function in the network as routers, implying changes to the routing design on the network.

Depending on the specifics of the environment, the addition of a new layer of routing can have implications on the IP addressing scheme in use.

Finally, many application firewalls rewrite, or translate, application URLs as part of their operation. This often implies a need to change the configuration of DNS inside the organization, or to propagate new DNS entries to the external DNS servers.

Does the application firewall require changes to SSL certificates?

Secondary Questions:

How does the product "see" into encrypted traffic? Will I have to replace any existing SSL termination products?

Most Internet facing applications use SSL encryption for some or all of the interaction with the users. As such, it's critical that your application firewall is able to inspect the encrypted traffic. Sometimes providing this access can simply be a matter of placing the application firewall behind a separate third-party SSL termination device. But in many cases this isn't an option, so the application firewall must have a mechanism to inspect encrypted SSL sessions.

There are two different strategies for accomplishing this: decryption and termination.

Decryption devices will load the SSL keys from the server application and use this information to decrypt the information without actively participating in the session. Generally, decryption devices require no changes to your SSL infrastructure.

Termination devices have their own SSL certificates and act as the endpoint for the encrypted session, reopening a new session (usually with the option of re-encrypting or sending clear text) to the application server.

The operational implication of this is a need to issue new certificates for the application firewalls and ensure that the certificate matches the correct domains on the user side. If the certificates aren't matched, users will get notice of such when they attempt to login to the application. It's possible to train users to ignore these prompts, but it greatly increases the likelihood they'll ignore the warning signs of spoofed sites (used, for example, in "phishing" attacks). The result is that new certificates need to be generated, and typically the management overhead for SSL certificates increases.

Does the application firewall require changes to the application code?

Secondary Questions:

Does the product change application URLs? Does the application firewall insert content into application data streams?

Some application firewalls make changes to the application data stream, such as rewriting URLs, signing or encrypting cookies, or even inserting their own "tokens" into the pages of the application. These techniques often imply a need to re-code applications to replace hard-coded IP addresses, or to rewrite JavaScript that accesses cookie information on the client. If the application security product is removed, even temporarily, these changes may need to be reversed to maintain continuous operation.

How much security administrator time and training is required for deployment?

Secondary Questions:

How much knowledge does the administrator need regarding application design? Does the firewall require that you create rules manually? How much does the product assist administrators in developing these rules?

Many application firewalls require a detailed understanding of the application to build and/or tune the security policy by hand. This implies that the security team must work with the developers to understand how to build the rules base and then the security team must communicate these changes to the operations team.

For a new deployment, the level of application understanding required, as well as the complexity of the application firewall, will dictate how much security administrator time and effort is required. For legacy applications, this level of understanding often doesn't exist in the organization, making it difficult to deploy products that require a detailed knowledge of the applications they protect.

For application changes, how much security administrator time and effort is required for re-configuring and re-testing?

Secondary Questions:

How much does the product automate or simplify the update process?

While most of the administrative effort for changes stems from the same requirements as those for initial deployment, the effort associated with application changes is often overlooked during evaluations. Unfortunately, it's also probably the most common reason for application firewall deployments failing.

In the test lab, or even for initial deployment, vendors or consultants can help with the initial configuration, essentially eliminating deployment issues as a concern, but what many security operations groups don't realize is that applications change far more often than networks and network protocols...how hard is it going to be to keep up?

Application changes may involve changes in modules, functions, URLs, parameter values and lengths, cookies queries, and scripts. Some application security products require manual intervention to account for these changes.

About the Author

Mark Kraynak is the director of product marketing at Imperva. Before joining Imperva, he held marketing and consulting positions at Check Point, CacheFlow (now BlueCoat Systems), and Ernst & Young's Center for Technology Enablement. Mark is a regular speaker on application and database security and participates in industry efforts to define the role of application firewalls in security architectures.
mark@imperva.com

Demand Vigilance from IT Security

ANTI-VIRUS IS NOT ENOUGH

BY GENE MANYAK

VERSION UPGRADES FROM software, infrastructure, and security vendors give businesses the impression that their enterprises are protected from new threats — but is it a false sense of security? The answer is yes if your security deployment doesn't address the elements that comprise today's threat landscape.

Hackers Don't Wait for Patches

To stay ahead of hackers, security software vendors release version upgrades on a regular basis. These upgrades typically include new defenses against the most recent attacks. However, major upgrades take time to implement, and to maximize operational efficiency, many organizations install upgrades once a year or less. Some larger global organizations upgrade their security mechanisms only once every several years. Even in a best-case scenario, when an organization immediately upgrades to new security software versions as soon as they become available, months can go by between installed upgrades. Hackers are acutely aware of this lag time between availability and installation, and are increasingly looking to exploit it. Without real-time security updates, businesses are powerless to stop them.

An April 4 InternetWeek article reports that "more than 70% of virus writers are now writing spyware under contract." Daily news items, such as a New Zealand Press Association report on March 9 that an Internet cafe attack made $500,000 of New Zealand Bank funds available to hackers, or the theft of $200,000 from Internet users through a fake auction site in Romania reported two days later in the Financial Times, indicate that today's hackers are increasingly motivated by real financial return. They're a more pernicious bunch than those of yesteryear who seemed motivated by the simple "challenge" of breaking in. This is precisely the type of hacker aiming to exploit the window between the availability and installation of security upgrades.

Keeping Up with Emerging Protocols

Networks are constantly supporting new protocols — like VoIP or 802.11x — before their security products do. New protocols mean new vulnerabilities, but what happens between upgrades? The answer would probably alarm most executives.

Anti-virus vendors provide ongoing virus signature updates. Intrusion-protection vendors provide ongoing protocol anomaly signatures. But few network and Web security product vendors offer analogous defense updates for new protocols, applications, and defense techniques. In other words, an ideal solution should provide ongoing updates not only for existing protocol and application defenses, but also dynamically add completely new defenses and defense techniques for protocols and applications as soon as they are supported. So if a completely new kind of vulnerability is discovered, or a previously uncommon protocol becomes popular, new defenses can be added dynamically to the security product's arsenal without requiring a complete product upgrade.

Remote Access: Another Can of Worms

No discussion of ongoing defense updates for network and Web security would be complete without mentioning remote access security. Often overlooked, remote access opens holes in network defenses because remote access traffic is often not subject to the latest available protections like other network and Web traffic.

SSL VPNs, in particular, contribute to the insecure nature of remote access. Most organizations think of SSL VPNs as secure connectivity, but security issues have prevented many SSL VPN pilots from expanding into full production environments. Spyware is a prime example of the vulnerability of SSL VPNs. While core defenses against spyware are provided by some Web security gateways, hackers are constantly creating new spyware programs and techniques. In many ways, the current spyware explosion is similar to the virus proliferation of previous years, and like their virus counterparts, spyware defense requires constant updates.

Preventative Medicine: AV Isn't Enough – But What Else Is Out There?

There's a misconception in the marketplace when it comes to upgrades and patches, and the antivirus software industry is the unwitting culprit. Many enterprises believe that their entire network is being protected once AV patches are installed. While it does much good, anti-virus software distribution isn't enough to protect all of the vulnerabilities in your network.

Viruses get a lot of press, but many network and Web attacks aren't, in fact, viruses, and aren't prevented by AV software. They're actually more complicated attacks that exploit protocol and application vulnerabilities. Consider Microsoft's monthly "Security Bulletin." Most exploits targeting the vulnerabilities in the bulletin will take the form of worms, and targeted protocol and application attacks. While most security software provides basic protection against such exploits, few protect against the most recent threats.

The bottom line in today's threat environment is that to obtain the highest level of defense, organizations simply can't rely on the next upgrade of their core security products. Achieving a truly secure network requires getting real-time, ongoing, dynamic defense updates for all types of network and web vulnerabilities, not just computer viruses. While you'll still have to do the heavy lifting involved in occasional product upgrades and patch management, a service that provides ongoing updates for defenses and security policies can save your business from the danger that lurks in-between upgrades.

About the Author

Gene Manyak manages Check Point's SmartDefense Services and has expertise across Check Point product and technology offerings. SmartDefense Services provide preemptive, ongoing, and real-time updates on exploits and vulnerabilities, and new attack protection capabilities to all existing Check Point customers and configuration advisories for Check Point defenses and security policies. Before joining Check Point, Manyak led product marketing activities at Valicert and was a general management consultant with A.T. Kearney, where he specialized in IT-related management challenges. Gene holds a BS in computer science from Cal Poly and an MBA from the University of Chicago.

gmanyak@us.checkpoint.com

Algorithm Agility & OATH

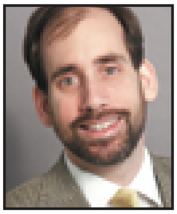

WHY A SINGLE STANDARD ALGORITHM FOR ONE-TIME PASSWORDS WON'T FULFILL THE PROMISE OF STRONGER USER AUTHENTICATION

BY BURT KALISKI

THE RECENT STARTLING announcement that the SHA-1 hash function wasn't as secure as believed raised interesting questions in the world of one-time password technology, since the newly proposed HOTP algorithm is based on SHA-1: Should the industry standardize around a single one-time password (OTP) algorithm? And what role should algorithm agility have in the future of one-time passwords?

HOTP, the HMAC-based One-Time Password algorithm, was introduced by OATH, a consortium organized last year to promote OTP technology. HOTP is based on the HMAC-SHA-1 algorithm (HMAC itself stands for Hash-Based Message Authentication Code), which in turn is based on SHA-1. In HOTP, a one-time password is computed as a function of a token secret and a counter value:

one-time password = HMAC-SHA-1 (token secret, counter).

Although HOTP is new, HMAC-SHA-1 itself is fairly widely standardized as a method for assuring message integrity, and it's also often recommended for additional purposes such as key derivation.

As it turns out, recent research results, which only affect the collision resistance of SHA-1 — the difficulty of finding two new messages with the same hash value — don't directly affect HMAC-SHA-1, which primarily depends on the one-wayness of SHA-1. Since HOTP depends on the strength of HMAC-SHA-1, not the collision-resistance of SHA-1, the results don't directly affect HOTP either.

Nevertheless, there's still good reason to question whether HOTP is suitable as a standard algorithm for one-time password generation — and, more generally, whether such a standard algorithm is necessary at all.

When an algorithm supports a protocol that's employed on a one-to-many basis, standardization can be quite important because the "many" may reflect multiple different implementations from a variety of vendors. For instance, code signing and digital certificates need standard algorithms to ensure that the signatures generated by one party can be verified by many others.

OTP algorithms that are based on a shared token secret, however, are inherently one-to-one: one token generates a one-time password, and one authentication authority verifies it — namely, the one that shares the token secret. Other parties may transport the one-time password (a desktop client, an application server, etc.), but they don't need to know how to generate or verify it. (Although the authentication authority might be implemented across multiple servers, these servers act in concert, being under the same administrative control.)

If a single standard OTP algorithm isn't necessary, one might ask if there's any harm in establishing a single standard. There are two major reasons, in fact, that it would be counterproductive to do so.

First, algorithms come and go over the years. SHA-1 itself was already on course for replacement by the next decade simply because of its originally expected underlying security level for collisions. The 80-bit security level, following recommendations by NIST and ANSI X9F1, has a "best before date" of 2010. It's not that the algorithm will become insecure at that point; it's that a conservative design suggests planning for gradual upgrades to higher security levels, and these take a long time in practice. A system based on HMAC-SHA-1 would need to accommodate stronger algorithms over time anyway, just to keep up with these recommendations.

Second, application requirements change over time, and innovation in OTP algorithms is needed in anticipation. At RSA Security, we've been developing a number of enhancements to our traditional time-based algorithm that offer a variety of interesting new features. Other token vendors likely have their own extensions to offer as well. A single standard algorithm would make this kind of innovation difficult.

Both these reasons are instances of the principle of algorithm agility: a system should be flexible in its choice of algorithm, where possible, both to maintain security and meet application requirements over the long-term.

The principle of algorithm agility is evident in many of the security specifications in wide use today. X.509 certificates, for instance, can convey any type of public key and can be signed with any digital signature algorithm — even algorithms that weren't envisioned when X.509 was first proposed two decades ago. Meanwhile, the industry has been able to transition from one hash function to another (e.g., MD5 to SHA-1), and to support multiple public-key algorithms (RSA, DSA, ECC), without any change in the certificate structure itself (although such changes would happen as improvements for other reasons).

The SSL and TLS protocols likewise support multiple alternative algorithms through the concept of "cipher suites," a process that has facilitated algorithm innovation for new applications. In addition, the PKCS #11 interface for cryptographic tokens works with a large variety of algorithms so that a one-time investment in the interface can return value in multiple algorithm environments.

Even if one could argue that the industry should move toward a single standard OTP algorithm, it's not clear that HOTP would be the best one. Counter-based algorithms are fine for many applications, but they suffer from the potential danger that user error may require significant resynchronization — what if I accidentally click the token too far ahead? — and they also provide no assurance of the actual time that the one-time password was generated. Both time-based algorithms and challenge-response algorithms address these concerns, and applications benefit from their availability as well.

Standardizing the generation of OTPs rather than their use doesn't really enable innovating in this space, except perhaps among the multiple vendors that will compete in terms of the implementation of an HOTP token. Competition around the features enabled by different algorithms, across a common framework, is a lot more interesting — and more robust for the long-term. This is why RSA Security has focused on building out that framework through the One-Time Password Specifications (OTPS) that include techniques for provisioning token secrets, retrieving one-time-passwords from tokens, transporting them to applications and authentication servers, and validating them — but, notably, not for generating them. Accordingly, the OTPS framework can work with any OTP algorithm, including HOTP. This framework will encourage more widespread use of many kinds of stronger user authentication, which will benefit the industry as a whole.

Ultimately, the most important issue is what is in the best interests of users. For the reasons just explained, standardizing on a single OTP algorithm doesn't fulfill the promise of stronger authentication for users. Industry collaboration toward a standard framework for integrating a wide range of OTP algorithms can. By ensuring that users and the organizations they interact with can leverage the OTP algorithms that best meet their needs — within whatever context they need to authenticate — the industry will be encouraged to make necessary long-term investments in stronger user authentication. ■

About the Author

Burt Kaliski is vice-president of research at RSA Security and chief scientist of its research center, RSA Laboratories. Burt received his bachelor's, master's, and PhD degrees in computer science from MIT, where his research focused on cryptography.

bkaliski@rsasecurity.com

›Linux ›Java ›Web Services ›.NET ›XML ›Wireless ›Storage ›Security
The Premier Resource for Today's Corporate & IT Decision Makers
IT solutions
›GUIDE
WWW.SYS-CON.COM/IT
TECHNOLOGIES YOU NEED NOW!
How to Manage Your Ideas Using Today's i-Technologies
Top IT Solution Providers!
› Candle
› ClearNova
› Computer Associates
› Creative Science
› DralaSoft
› GraphOn
› InterSystems
› ITKO
› Merant
› O'Reilly
› Parasoft
› QuadBase
› SAP
› Willy Technology
› Zero G

Security and Storage Granularity

WHICH DATABASE STORAGE SYSTEM IS RIGHT FOR YOUR COMPANY?

BY WINN SCHWARTAU

WHEN MY COMPANY was designing its data center, we had to make a choice: What kind of database storage system was going to be the backbone of our operations? As in most things IT, the options were seemingly endless, and there are many criteria to consider before investing time or money into development and deployment.

1. Price

Some database storage approaches can be very expensive, often requiring recurring license fees and specific hardware. Others are virtually free and can function perfectly well on generic platforms.

2. Scalability

How much expansion is going to be needed over what period of time? How many users are going to be accessing how many files distributed over how many locations connected by what infrastructure?

From my point of view, everything else about the database storage decision process was security oriented.

Reliability (Availability)

As a security guy, I am intensely aware of the need to ensure that all systems are up and running at all times. This is especially true in my storage network, without which I am out of business.

High degrees of reliability are achieved using several complementary approaches. Redundancy suggests having more than one of everything, including the engine and production system. Fault tolerance is a popular method to protect against errors in real time and backup is an absolute must for all storage systems. Real-time backup is great for time-sensitive information where there is no room for error; less expensive batch backup is fine for less critical applications.

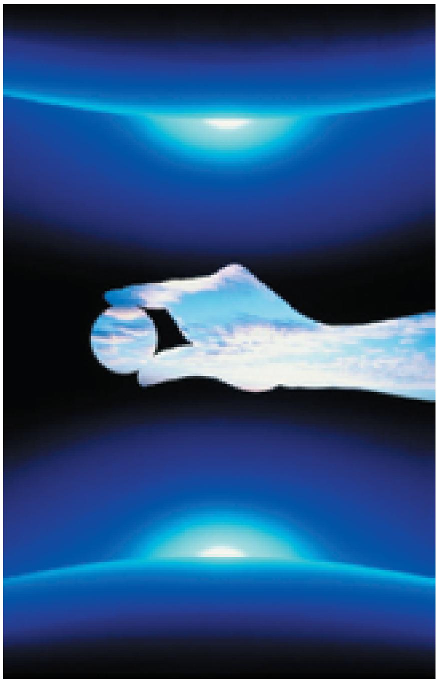

Physical Separation of Assets

Don't put your backup system close to your primary storage facility. A physical attack or Act of God should not take down everything. Resiliency in the physical domain is just as important as in the logical.

If your storage systems are really mission-critical, check out the physical pathing of the logical connections to your data centers. I've seen way too many times that the effectiveness of expensive redundant servers is diminished when both sets of equipment are serviced by the same hunks of copper or glass. Add the extra cable and make sure that they are laid in different places.

Encryption (Integrity and Confidentiality)

You don't want your compliance and governance data hitting the net nor do you want the prices changed in your shopping cart. Privacy is s must these days for customer confidence.

Access Control

Who is able to read, write, modify or delete the data in your storage facilities. Pretty important I should say! Security in the storage arena can be very complex, but there are ways to minimize the complexity by thinking out the problems and solutions in advance.

Think Granularity

All pieces of your network are not created equal, but all pieces should be treated equally from a security standpoint.

In enterprise network security, does the intensity of security at the branch office need to be same as that of the main office? Does the accounting storage server need to be more or less protected than the public marketing materials files?

If all objects in your enterprise are not of the same value, do you place an inordinate amount of effort (and money) into protecting assets that may in fact need minimal security?

The issue of granularity is critical to all security and storage at all levels:

Network
Sub-Network 1…n
Sub-sub-networks 1…n
Servers 1…n
O/S 1…n
Applications 1…n
Data Bases 1…n
Files 1…n
Cells 1a…nz
Attributes

Designers need to ask themselves: "Do I really need to exercise so much control to make sure Mary and Bob cannot access cells X, Y, and Z in DB34.DB on MYDB.EXE?"

Is it easier, perhaps, to create logical isolation through access control at a grosser degree of granularity? Consider: Is it easier to control group access to subnets, servers, and applications or…is it easier to control access to specific files and cell calls within the database?

Our answer, for our purposes, was based on the real-world access control rules from a business operations standpoint. And cost. We have thousands of users...and lots of accessible objects all over a range of servers. We architected our systems with security in mind, at the level of granularity dictated by our needs.

One possible approach was picking up a single licensed database storage solution with an absolute infinity of options and security granularity down the attribute level. There are some great products that will do this – add fault tolerance and all of that nice stuff. However, it is pricey and the administration can be incredibly cumbersome. Complexity breeds weakness, and we did not want to hire a cadre of application-specific experts to manage such a huge system.

We like the concepts of distributed management, simplicity, appliances, and low-cost with an overriding measure of security, so we chose an open source approach for our storage needs.

I really like appliance models. If one fails, I'll lose only a small piece of my operation. Sure, we have FT and all the rest, but cost-effective scaling is questionable. I really like simplicity. There's a whole lot less room for error and failure. I don't want to have to send a cadre of managers to database management school when the open source knowledge is a common skill set.

> "before spending a dime, consider the true costs of ownership, management, training, staffing, licensing, and maintenance **before deciding a less granular open source approach won't work"**

That being said, we function in a fairly benign environment and we are not guided by overarching compliances and governance. For organizations that have massive centralized repositories of common data sets, with legions of distributed "clerks" who must interface with the data regularly, a highly structured and intensely granular design for the storage data base may be in order.

But before spending a dime, consider the true costs of ownership, management, training, staffing, licensing, and maintenance before deciding a less granular open source approach won't work. You might be surprised.

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

NAS and SAN: The Waiter and the Chef

COMPLEMENTARY TECHNOLOGIES

BY JOHN MEDASKA

EACH YEAR, ONE of the most eagerly awaited events for food aficionados is the publication of the new Zagat Guides for restaurants. Within the pages of various editions are listings and reviews of hundreds of the top dining experiences around the world – each designed to delight the palate and rejuvenate the soul. Earning top marks in the Guide is not only a source of pride; it's essential to the success of these top-tier establishments.

While there are a great many factors that go into a great dining experience, essentially they fall into one of two categories: the quality of the food and the quality of the service. Both must be exceptional in order to make the top of the list. If the food is great and the service is poor, or service is great but food is humdrum, the restaurant falls down the list.

Today's storage world is much the same way. Much is being made about network attached storage (NAS) and storage area networks (SANs) as options for the high-volume data storage needs of modern enterprises. Yet when you look closely, these are not competing solutions, but rather complementary technologies that are best suited to different tasks.

Just as you probably don't want the chef waiting on your table, or your waiter cooking your duck al'orange, it's important to make sure your storage technologies are doing what they do best, and not trying to cover a function best left to the other. Let's take a look at the best functionality of each (the technologies, not the restaurant staff), and see how they fit into an overall information lifecycle management (ILM) strategy.

NAS – The Waiter

In the storage world, NAS serves the function of the waiter. It works best for file or block-level data access, acting as a gateway between the SAN and workgroups or users. In other words, it brings the information out from the kitchen and sends it to the appropriate table. This is a function it performs very well.

NAS is attractive because it is generally plug and play, with a low cost of acquisition and management. There's no need to carve out logical units (LUNs) the way you do with SAN since the RAID array, tape, hard disk, or other device is attached directly to each server or group of servers. This method ensures it is up and running quickly. NAS is also very agile, serving up data quickly as needed since there's a one-to-one relationship between the network and the storage unit.

From a technical standpoint, NAS uses an IP protocol to serve files to clients. In effect it acts like a giant network server, only providing access to a larger pool of files.

Where enterprises run into trouble with NAS is when they try to make it their primary method of high-volume bulk storage. Usually they are comfortable with the NAS design they already have in place and continue to add to it. That strategy seems logical on the surface, but in practice it doesn't work as well as you might hope.

The problem is that while NAS has some scalability, that scalability is not linear. At some point the curve flattens out and NAS is no longer capable of handling the workload. Depending on the size of the organization and the topology of the network, having individual NAS servers for various workgroups also tends to work against its native simplicity, requiring more resources rather than fewer to manage the organization's storage needs.

In a small organization, NAS can serve both functions – just as one person can cook and serve the food in a small restaurant. But as the enterprise grows and becomes more sophisticated, the needs change and it's time for a separation of responsibilities.

SAN – The Chef

Where NAS is more of a device-oriented strategy, SANs are really an architecture or method of providing storage. They incorporate a wide variety of storage devices and storage spaces that sit at a higher level than a typical NAS device. They serve up data blocks to servers over a Fibre connection rather than directly serving files to clients. A server taps into the SAN when a request comes in, then provides the files out of that data block.

SANs are designed to help improve throughput and file sharing by centralizing data rather than dividing it by workgroups. This arrangement also helps speed and simplify critical backups in large organizations. In short, it is the lynchpin in an effective ILM strategy.

Going back to our restaurant analogy, the SAN is the kitchen where all of the food is prepared. It doesn't matter if the diners order beef, fish, poultry, or even vegetarian. Everything needed to give them what they want is there, and it is then routed out on demand via the waiters. Using NAS for the same task would require separate kitchens for each of those types of dishes, or one kitchen for the tables covered by each of the waiters. And the waiters would have to take the order, then go in back to cook the meal. This is not the function of that employee. With that in mind, NAS solutions don't fit every storage need. The convenience of replicating a NAS storage solution throughout the entire enterprise is outweighed by the fact that it's not designed for certain situations.

Putting a SAN solution in place takes the burden off of local servers, speeding delivery of the information to the user by eliminating the need for servers to search their own disks (or extensions thereof) for data. The network is not congested with an abundance of IP traffic but rather the storage fabric network handles that transport. Instead, data storage becomes more of a virtual function, a pass-through from

"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

Free iPod® shuffle

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™

Check www.e7software.com for official contest rules.

the server to the mass storage arrangement that has been optimized for this single function. Separating storage from servers simplifies storage administration; instead of having to manage multiple LAN or WAN storage arrangements, IT resources can administer a single, centralized, dedicated resource.

SAN makes storage more efficient as well. In a typical network, one server might be maxed out on storage space while another has several GB of space available. SANs efficiently pool all the storage together so each server has equal access to the total amount of space available within the organization. They also provide the ability to manage that storage centrally, which alone could justify the ROI on that investment. This method also helps reduce file redundancy since files should only be stored in one place (the SAN) in the organization rather than on multiple servers throughout the enterprise.

Another advantage of SAN is that it makes the "black box" concept for storage work. It places an umbrella over the system, allowing you to mix and match manufacturers rather than having to accept a monolithic storage solution that locks you into a single manufacturer (and that manufacturer's pricing structure). With SANs, the economics of competition come into play, allowing you to seek out the best product (and best deal) as new needs arise. That is true storage virtualization. This will allow companies to constantly analyze their IT storage portfolio to maximize their storage investment.

This concept also ensures that you are able to protect the investments you've already made rather than having to scrap one SAN in favor of another. You can add or replace storage units as necessary and easily fit them into your overall SAN strategy.

The downside of SAN is that it's not as good as NAS at working with multiple file platforms. For those used to working with NAS, there may also be some sticker shock as it can be far more costly. In addition, setting up a SAN is much more complex than installing NAS on the back end of your servers. Some SAN customers never utilize the full benefits of the SAN, instead using it as a basic backup and storage device. In those cases, customers should make sure that if they are paying for an eight-course meal they are not merely nibbling on the appetizer and going home. You have to make sure you have the time and resources available to make the jump in order to realize the business benefits.

Which to Choose

Given these factors, the question facing many organizations is whether to stay with (or add) NAS to their networks or turn over all the cards and bring in a SAN. The answer, in my opinion, is yes.

A blended solution of NAS and SAN offers the greatest flexibility and performance advantage for most organizations. The more heterogeneous your server environment, the more important NAS becomes for smooth operation between servers. And the higher the volume of data firing around your enterprise, the more important SAN becomes for working with it effectively.

Having NAS in place simplifies access to the SAN. In fact, NAS is the ideal gateway to a SAN, helping take the data blocks provided by the SAN and routing them to the proper servers in the form of files. At the same time, having a SAN in place allows NAS to work more efficiently by removing the burden of mass storage of less critical data. Important files can be stored locally on the NAS device, while those thousands of joke e-mails tying up space on the Microsoft Exchange server can be offloaded to the SAN.

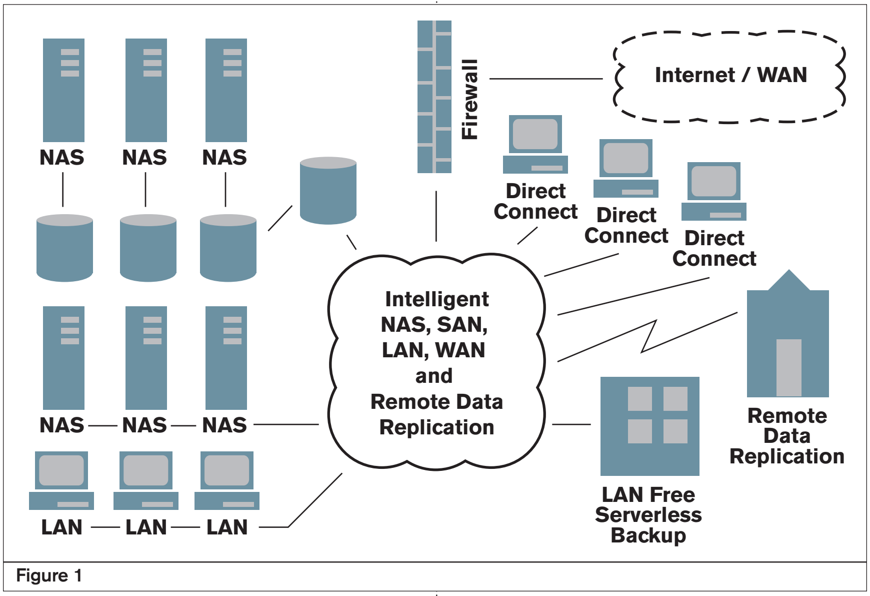

Figure 1

Getting to ILM

Establishing the right combination of storage is critical to achieving the goals of information lifecycle management. The whole purpose of ILM is to allow organizations to prioritize data and establish a hierarchy for information based on its value rather than treating all data as equal. Combining NAS and SAN makes it happen.

SAN provides a foundation for ILM by allowing you to segment storage in LUNs. Once it is segmented, backups can be set to occur at different intervals by order of importance rather than backing up everything on the network every night. It also allows you to move from physical tape to virtual tape drives – systems that use your current tape back up software to save data onto hard drives.

Virtual tape drives are proving far more reliable than tape, which studies have shown fail roughly 70% of the time during backup. They also eliminate the need to perform a time-consuming restoration since the backup is disk-based and therefore nonlinear.

NAS provides accessibility to the stored data, helping users get to critical data on the SAN quickly. It makes it easier to find data in a given LUN and provide the cross-platform access required by various applications. Together, they help the organization manage its data better, driving down costs while freeing up resources for other tasks.

Cooking Up Success

Just as a restaurant needs both top chefs and attentive waiters to earn the top ranking from Zagat's, your enterprise will benefit from a hybrid NAS and SAN solution. Together they will keep your storage optimized, your business more productive, and your enterprise data safer than an either/or approach. ■

About the Author

John Medaska is vice president of business development for Relational Technology Services, a wholly owned subsidiary of Relational, LLC. Relational Technology Services specializes in IT portfolio management, technology assessment, and installation, including storage strategies and execution. jmedaska@relationaltechnology.com.

Guided Intelligence Analytical Server

FUSIONWARE HELPS IT TRANSFORM DATA THAT MEETS COMPLIANCE REGS

BY ROBERT HOUBEN

THE FINANCIAL AND healthcare industries continue to face increasingly more stringent compliance requirements. Small and medium-sized enterprises (SMEs) are turning to business intelligence (BI) tools to help make sense of the mountains of data and to meet the stiff compliance regulations. Often though the most difficult, expensive, and complex part of using such systems is the challenge of organizing information into a standardized format and building the correct processes to get the right data through the correct BI tool.

Guided Intelligence has implemented a solution, the Guided Intelligence Analytical Server (GIAS) that uses the FusionWare Integration Server to provide secure point-and-click ASP-based BI on-demand to transform data into a powerful analytical application to meet compliance mandates. Such applications were typically only available to large companies with million-dollar data warehousing systems before GIAS was introduced.

Automating Systems in Accordance with the New Standards

Following the important movement that's underway to standardize information into one format for financial reporting as initiated, for example, by the XBRL consortium (http://xbrl.org/), Guided Intelligence worked on building a viable solution to sell to end users via channel partners. The company needed to find an XML server that would help build the framework for transforming its quasi-automated process into a completely automated, push-button system. The end result is the Guided Intelligence Analytical Server (GIAS), providing a service-oriented architecture as the basis of a Business Intelligence on-demand or BI as Web Service solution.

One of the company's target markets is the financial industry, a business that's heavily involved in compliance and always looking for ways to exchange data per Sarbanes Oxley.

Building a Framework

Guided Intelligence chose FusionWare (www.fusionware.net) technology as the framework for its automated system. The FusionWare Integration Server enables companies to build and deploy robust integration solutions rapidly and cost-effectively because of its built-in code generators, wizards, and query builders through an easy-to-use point-and-click interface.

The FusionWare Integration Server provides the complete runtime infrastructure for the Guided Intelligence ASP model, requiring no extra servers or databases. It also provides the development environment that lets Guided Intelligence build the core Web Services that enable files in an application to be transformed into analyticals cube on-demand without having to pre-define them. The processes created help define automated scripts to transform data from a user's application into a standardized format that drives the cubes, dashboards, and other analytical presentation models. Essentially, the company is enabling existing reporting models or forms to be transformed into OLAP (Online Analytical Processing) cubes that let people understand better what's going on with their data. OLAP is a metaphor for a data storage model in a BI environment modeled after a Rubik's cube.

FusionWare met all of Guided Intelligence's criteria from a technology perspective. Guided Intelligence was also impressed with FusionWare's unique willingness and ability to understand what it was trying to accomplish and, therefore, deliver results.

"We had been following the industry trends, working very closely with several players in the XML world, and evaluated the leading vendors in this space," said Guided Intelligence president Warren Richman. "We chose the FusionWare Integration Server because the product met all of our criteria from a technology perspective, but what really drove our decision was the FusionWare team – they are sensitive to the needs of their channel partners and really deliver on what they promise in their marketing."

Transforming Data into Sophisticated Analytics

Using its ASP-based usage model, Guided Intelligence developed a solution targeted at the financial services market called Audit Intelligence. For roughly what a CPA would charge a client an hour, a company can get the benefit of the entire Guided Intelligence Analytical Server. Normally, it could take days or weeks to set up a sophisticated OLAP cube and accomplish the same thing. With FusionWare technology driving the Guided Intelligence Analytical Server, the customer doesn't have to worry about hiring expensive .NET or Java consultants to customize its solution or add complexity to its existing infrastructure with additional systems such as application servers, Web servers or databases.

With pushbutton simplicity, Audit Intelligence integrates a customer's data, takes financial statement information, and transforms it into a forensic application for analyzing a business's financials over multiple periods of time. Any organization that needs to validate the integrity of a company can use this powerful, yet easy-to-use tool.

CPAs, financial planners, CFOs, internal auditors, and certified fraud examiners will have an effective way of getting at all the information. Instead of having to worry about how to gather information from diverse data sources, they can just focus on analyzing it. The FusionWare Integration Server's powerful data and application connectivity capabilities simplify the integration of the customer's back-end systems.

Another benefit of providing an "expert" system is that it helps address the critical shortage of trained CPAs for firms and their clients that have to meet guidelines. As the article "Accounting Firms Scramble To Find Experienced CPAs" by Suzanne McGee published in *Career Journal* points out:

> *Public-accounting firms are ramping up efforts to hire and retain senior-level talent as new auditing and accounting rules continue to fuel clients' need for their services* [*http://www.careerjournal.com/salaryhiring/industries/accounting/20050322-mcgee.html*](http://www.careerjournal.com/salaryhiring/industries/accounting/20050322-mcgee.html)

Similarly in healthcare the requirements for better management of information for more efficiency in managing a medical practice is both facilitated and encumbered by the rigid HIPAA requirements. In much the same way that Guided Intelligence transforms financial data, its applications for healthcare can transform medical billing data into a rich source of analytics including practice analysis, financial analysis and documentation compliance analysis.

Opening the Channel

The Guided Intelligence Analytical Server can be delivered both as an ASP solution or installed in an enterprise. Guided Intelligence was able to create a "frictionless" process that requires virtually no manual data manipulation. This was important because it wanted to drive its price point to a level the SME market it was calling on could afford. Guided Intelligence delivers its solutions through channel partners who focus on that market.

Guided Intelligence VARs can offer Audit Intelligence and any of its other applications to their partners as branded applications. This indirect model is also ideal for application portals or transaction-based models to include an analytics model.

The Future

The next phase of Audit Intelligence running on the Guided Intelligence Analytical Server will let people analyze any publicly traded stock through a portal. The pushbutton automated system will allow the automated transformation of publicly filed financial statements through the SEC's Edgar database into Audit Intelligence, which will analyze multiple years worth of financial data and help identify trends in ways that relate to fraud, misstatements, or other financial shenanigans. These are tools that have previously been unavailable to small companies.

With this increased functionality, any of Audit Intelligence's users can compare private trial balance information and public companies. These companies will be able to look at their own financial statements and put them in the same format, model, and terminology as they would industry leaders. This opens potential areas for benchmarking, mergers, and acquisitions.

Using FusionWare's HIPAA-compliant transformation capabilities, Guided Intelligence will be targeting other applications including healthcare, CRM and expense analysis to help companies connect and make sense of exhaustive amounts of unstructured and structured data, and comply with strict regulations. ■

About the Author

Robert Houben is CTO of FusionWare and the chief architect of the FusionWare Integration Server. His 20 years of expertise working with technology-driven companies includes an in-depth knowledge of middleware and applications, database technologies, and Web Services. Robert was a driving force in the development of the ODBC standards in 1992 and actively participates in promoting open standards technologies. He has authored a number of patent filings for the FusionWare Integration Server and published numerous articles related to database technologies and legacy integration solutions.

Virtualizing NAS?

KEY CRITERIA FOR NETWORK FILE VIRTUALIZATION

BY JACK NORRIS

MANAGING THE SURGE of file-based data has become increasingly difficult and complex. Virtualizing NAS through Network File Virtualization (NFV) simplifies storage management and enables administrators to easily address management and utilization challenges without affecting data access.

With NFV, storage administrators need not be concerned about the impacts on end-user data access. Administrators can keep data always accessible and online. This eliminates a major constraint on the storage administration process. NFV lets end-users retain full read/write access to data as it is being dynamically relocated within networked storage. This dramatically changes unstructured data management and enables administrators to increase capacity utilization, improve performance, leverage tiered storage and ease consolidations all while end-users continue to access and update the data.

Not all Network File Virtualization solutions are equal, however. The key to successfully adopting NFV in your environment is to ask the right questions. There are three key questions to consider:

- What problem does it solve?
- Does it create new problems?
- How does it leverage my existing environment?

First, you need to look closely at what is actually being virtualized. Is it only the location/namespace, or does it include active files? The ability to perform active data management drives many benefits. Performance management applications require an NFV solution that virtualizes active content and can dynamically relocate open files across devices to relieve hotspots. Similarly, managing tiered storage requires a solution that goes beyond "moving dead data." Simply identifying content that hasn't been accessed in a period of time and moving this infrequently accessed data, is a small part of Tiered Storage Management. Actively matching content with the appropriate underlying storage device requires the real-time relocation of content and active data management.

The second major question is "Does the solution create new problems?" You will need to know if the NFV product will create worse problems than it solves. For instance, an NFV product's high management overhead might cancel out other productivity gains, its persistent metadata may or may not be secure, or it may cause a serious performance bottleneck. Look at the product's total potential benefits in light of all the associated costs and likely impacts. You'll want to know details on issues like fault tolerance, potential downtime, high availability and performance levels, and restore performance and procedures in case of failure.

Third, you should determine how well a Network File Virtualization solution works with your existing infrastructure. Does it support your existing file systems and all of the software and management tool investments you've already made? This is particularly important when it comes to existing backup and recovery procedures. You need to know the specifics of how a given NFV technology acts in the real world with your snapshots and backup applications in your environment. Other key questions include downtime issues, integration with primary storage platforms, support issues with software and virtualization vendors, namespace schema, and protocol functionality.

These questions can be used to segment virtualization approaches and to understand the differences among them.

Approaches to Virtualizing NAS

One popular approach is to combine a proprietary namespace into a virtualization appliance. This delivers many of the benefits of NAS virtualization. It provides file movement with no end-user disruption, but there are several issues to consider. This approach raises many issues with existing backup and recovery procedures to ensure reliable restores. How

Checklist Sidebar

✓ Transparency
✓ Transparent to install, no mount point changes
✓ No client or server software to install
✓ No proprietary namespace
✓ No high availability, DR tasks
✓ No data access risk
✓ No performance exposure
✓ Leverage standard namespace
✓ Supports existing environment
✓ Management tools, data protection policies
✓ Standards support
✓ Vendor certifications

do you perform a reliable restore of an individual volume or file server?

With this approach, all data access has to flow through the in-band appliance. The scalability concern is not limited to the network hop. In fact that's the easy part. Scalability and performance concerns stem from the amount of processing and file information required. The larger the environment the more likely the file level information will be written and retrieved from disk greatly impacting latency.

Last, the proprietary namespace introduces risk. The namespace and associated meta data needs to have data protection procedures and high availability to ensure it is always available. If the appliance were to go down, all of the look-up information is contained on the box and all of the clients are mounted to the failed box. How do you recover? And how long will that take?

Another approach is the out-of-band only approach. This does not have the performance or data access risk of the proprietary namespace approach. But without an in-band only capability there is no way to provide read/write access to active data. This solution cannot handle open files, or stale mounts to ensure continuous access. Without the ability to handle open files there is no performance management support or tiered storage applications. While the embracing of an out-of-band namespace has advantages, this solution is also incomplete.

There are solutions on the market that combine the best of in-band and out-of-band approaches without the downsides. When looking at solutions to virtualize NAS, here's a checklist to go with some of the detailed questions. First, look for a solution that provides transparency. Not just transparency of file access, but a solution that is transparent to the environment – a solution that does NOT require mount point changes or the deployments of agents on clients or servers.

Second, look for a solution that doesn't require a persistent namespace. There are many advantages with this approach. It limits risk, limits performance concerns, but also leverages continuing investments being made by large vendors and standards body whether your namespace is Microsoft DFS, Automount, or in the future NFS V4.

Third, virtualization should leverage the investments you've already made in your storage infrastructure and management tools. Check for standards support, vendor certifications, and make sure your existing management tools and data protection policies are not adversely impacted.

NFV dramatically changes NAS management. With it you can dynamically relocate data for capacity, performance, or cost reasons without disruption to end-users or applications.

Virtualization approaches differ greatly. The overall management savings, risk exposure, and scalability can range greatly. Virtualization should be a transparent layer in your environment that simplifies – not interferes with your existing environment.

About the Author

Jack Norris is the VP of Marketing at Rainfinity (www.rainfinity.com), the first company to optimize IP-based storage (NAS and file servers) with its Network File Virtualization Platform.

jnorris@rainfinity.com

From the Co-editors-in-Chief

— continued from page 3

In parting, we would like to emphasize that there is a special relationship between security and data as the whole point of security is most often all about protecting the data. As a periodical with both "Information Storage" and "Security" in the title, it makes sense if we start with how these two relate. Data, in all its various forms, is the only thing that matters to a business organization. If a company loses hardware, facilities, even staff, as long as the data still exists to track who customers are and what is owed then the pieces can be picked up. On the other hand, if the data is gone, you'll find that more often than not, there is no longer a business. Without knowledge of inventory and invoices and payroll, etc., all you really have is infrastructure. You most likely already know this, but it often bears repeating since data provisioning and security often get short shrift due to their illusive nature. The data looks safe and secure mainly because it can't be seen completely. The best way to avoid this bad habit is to treat data like the asset that it truly is. On the day of this writing, Patrick encountered yet another incident of a company that had set up backups for vital data, but allowed the process to be left unrun for months because no one was checking the details. Examples like this from the real world are compelling and in this issue you should find your fill of solid examples and statistics based on experience. We hope you enjoy this issue and those going forward as we work to bring you the best information for your hardest decisions!

ISSJ | Advertiser Index

Advertiser	URL	Contact	Page
Barracuda Networks	www.barracudanetworks.com	408-342-5400	13
Black Hat	www.blackhat.com		15
E7Software	www.e7software.com/risk	800-824-4717	27
Forum Systems	www.forumsys.com	866-333-0210	9
Imation	www.imation.com		7
ISSJ	www.issjournal.com	888-303-5282	17
IT Solutions Guide	www.sys-con.com	201-802-3021	23
NTP Software	www.ntpsoftware.com/learn	800-226-2755	29
SafeNet	www.safenet-inc.com/hse/15	800-697-1316	Cover IV
SurfControl	www.surfcontrol.com/go/threatshield	800-368-3366	Cover II
SYS-CON e-Newsletters	www.sys-con.com	888-303-5282	21
Tacit Networks	www.tacitnetworks.com/ROI	888-757-TACIT	11
Tenable Network Security	www.tenablesecurity.com	877-448-0489	Cover III
The Security Awareness Company	www.thesecurityawarenesscompany.com	727-393-6600	19, 25

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

Application Security Announces Alliance with Internet Security Systems

(Washington, D.C.) – Application Security, Inc. (AppSecInc), a provider of proactive database security solutions for the enterprise, has announced a strategic relationship with Internet Security Systems (ISS) that highlights the advantages of a layered approach to securing enterprises.

Under terms of the ISS and AppSecInc agreement, the companies will engage in joint sales and marketing efforts in support of the AppDetective database vulnerability assessment scanner. AppDetective discovers applications within an organization's infrastructure and assesses their security strength. It locates, examines, reports, and fixes security holes and misconfigurations. www.appsecinc.com

Tenable Executive Named 2005 Techno-Security Professional of the Year

(Columbia, MD) – The TrainingCo. has announced that Marcus J. Ranum, chief security officer of Tenable Network Security, has been named 2005 Techno-Security Professional of the Year.

Marcus Ranum, an expert on security system design and implementation, is recognized as the inventor of the proxy firewall, as well as the implementer of the first commercial firewall product. Since the late 1980s, Ranum has designed a number of groundbreaking security products including the DEC SEAL, the TIS firewall toolkit, the Gauntlet firewall, and NFR's Network Flight Recorder intrusion detection system. He has been involved in every level of operations of the IT security industry, from solutions developer to CEO. Ranum has also served as a technology advisor to a number of start-ups, established concerns, and venture capital groups.

The TrainingCo. is a provider of specialized security, forensics, and cyber crime training and conference production. www.tenablesecurity.com

SafeNet Completes Acquisition of MediaSentry

(Baltimore) – SafeNet, Inc., a provider of information security, has completed the acquisition of MediaSentry, Inc., a global provider of anti-piracy and business management services for the recording and motion picture industries.

The acquisition will expand SafeNet's anti-piracy offering to include the protection of copyrighted content on peer-to-peer networks, while gaining a new competency in managed services with a stellar customer list.

SafeNet acquired all the issued shares of MediaSentry, Inc., for a cash and stock consideration of $20 million, complemented by an earnout schedule. The mix of cash and stock is 70 percent and 30 percent, respectively. www.safenet-inc.com

Fifty Percent of Storage Professionals Surveyed Cite Backup Reliability as Top Concern

(Westborough, MA) – ExaGrid Systems, providers of self-protecting storage, have surveyed storage IT professionals attending recent storage conferences. The survey showed that 50 percent of the respondents are concerned that their data may not be reliably protected by current backup processes. In addition, an even greater number don't have confidence that they could adequately restore backed up data in an appropriate time frame if needed.

The informal poll of more than 100 professionals was conducted at two conferences for storage IT professionals – Storage Decisions and Storage World Conference 2005. Additional data from the survey showed 40 percent of participants expressed concern about the effectiveness of their disaster recovery abilities, with 30 percent of total respondents not even having a disaster recovery solution in place for reasons that included lack of budget, organizational priorities, and general confusion about the technology options available. www.exagrid.com

Barracuda Networks Adds Two Enterprise Models to Spyware Firewall Line

(Cupertino, CA) – Barracuda Networks, Inc., a provider of enterprise-class spam firewall solutions, has expanded the Barracuda Spyware Firewall line to include two new enterprise models: the Barracuda Spyware Firewall 610 and 810. Featuring greater throughput and higher capacity, the Barracuda Spyware Firewall 610 and 810 models are the ideal solution for large enterprise environments that demand greater bandwidth for Web communications.

The Barracuda Spyware Firewall 610 and 810 provide over 200MB/sec throughput and both models include hardware redundancy features including RAID disk storage. Both models also include dual gigabit Ethernet ports in addition to the powerful features offered with the Barracuda Spyware Firewall 210, 310, and 410 models. www.barracudanetworks.com

Brocade and Tacit Networks Announce Strategic Investment, Licensing, and Development Agreement

(San Jose, CA) – Brocade Communications Systems, Inc., a provider of infrastructure solutions for Storage Area Networks (SANs), and Tacit Networks, a provider of enterprise-wide remote office IT solutions, have announced a strategic relationship to deliver Wide Area File Services (WAFS) to enterprise customers on Microsoft's Windows Server 2003 platform. Brocade will invest up to $7.5 million for a minority ownership share in privately held Tacit Networks, and will immediately add Tacit Networks' WAFS solution built on Microsoft Windows Storage Server 2003 to its product portfolio.

The Tacit Networks' WAFS solution includes a data center server and remote office appliances that work together to give all connected locations access to stored data center files and applications, while maintaining centralized management and backup capabilities. Under agreements entered into in connection with Brocade's investment, Brocade will market the solution to its partners and customers worldwide, and the two companies will partner in customer support, and on product development programs.

www.tacitnetworks.com
www.brocade.com

Forum Systems Announces 64-Bit Support to Deliver System Performance and Capacity

(Las Vegas) – Forum Systems, a provider of Web services security for threat protection and trust management, has announced the immediate availability of a 64-bit version of its Web Services Firewall. Forum XWall X5.0 x64 Edition provides unlimited virtual memory and workload processing capabilities. Since memory is a primary resource for scalable computing, increased physical and virtual address space allows larger amounts of data to be processed. The handling of large data volumes such as binary attachments within SOAP (Simple Object Access Protocol) messages is becoming a mandatory requirement for financial services organizations and government agencies. The 64-bit technology also enhances I/O (Input and Output) to raise the ceiling on transactions processing rates and offers better cache management to eliminate delays in servicing requests.

www.forumsys.com

TENABLE
Network Security
NeVO

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/hse14*

Call 1-800-697-1316 to be SafeNet sure.
www.safenet-inc.com/hse15

SafeNet®
The Foundation of Information Security

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