Tokenization: The Building Blocks of Spam ⟨20

information STORAGE+ SECURITY journal

www.ISSJournal.com

In This Issue:

VOLUME: 2 ISSUE: 5 2005

DEPLOYING A SAN ⟨4

TO CENTRALIZE STORAGE ACROSS THE ENTERPRISE

Fighting the cost and complexity of storage

Are You Ready?

Yesterday's point-solution is no match for today's blended threat—and you can't expect your enterprise IT security experts to be a 24/7 clean-up crew. But you can count on SurfControl's Enterprise Protection Suite to deliver unequaled protection against every threat—traveling through every entry point—every time.

It doesn't matter whether it's spam, spyware, phishing, viruses or a specialized day-zero hybrid. Nor does it matter whether it comes from inside your organization, or from outside company walls. The SurfControl Enterprise Threat Protection Suite delivers a powerful unified threat management solution, securing Web, e-mail and IM/P2P traffic—from the network gateway to the user desktop. Plus, it's backed by SurfControl's 24/7 Adaptive Threat Intelligence Service.™ Now you're ready.

FREE 30-day trial **www.surfcontrol.com/go/blended** | 1 800.368.3366

Enterprise Protection Suite
Web, E-mail, IM/P2P, Mobile

Enhance Security
Manage Usage Policies & Compliance
Increase Productivity
Reduce Costs & Administration

SurfControl Adaptive Threat Intelligence™
First to Detect™ First to Protect™

SurfControl®
Enterprise Threat Protection

information STORAGE+ SECURITY journal

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

Security – No Longer the Perennial Afterthought

BY PATRICK HYNDS AND BRUCE BACKA

STORAGE ALWAYS SEEMS to come first in technical discussions and security seems to be the perennial afterthought. This can be considered reasonable given how we shop for things in general, namely finding the thing that meets our expectations and then ensure it has all the bells and whistles. The good news is that this seems to be changing bit by bit as our industry realizes that security is no longer a nice-to-have feature, but is actually a core requirement. This movement was brought into focus recently when Patrick was involved in a meeting with Senator John Sununu of New Hampshire in which they discussed current technical challenges. A year or so ago the mention of security in discussing wired versus wireless infrastructure would likely sound like a non sequitur to many, but now it was something that had already been considered. Decisions in public policy and those in corporate board rooms are finally gaining the correct perspective.

Senator Sununu commented that he felt that we are currently at the low point of security in technology in general. While this sounds like a glum declaration, it's actually a hopeful prophecy. If we can make everyone think about data and the security of that data in the same way they currently think about data and the accessibility of that data, then we may dare to hope that things will be better in the future. Consider us here at ISSJ as being among the hopeful!

In this issue we are dealing with some of the all time heavy-duty subjects of information storage, namely SAN, NAS and backup. These are the topics where the word Terabyte is currently most likely to be heard. In this issue, you should expect to find valuable information on the best topics we think we can find to help you meld the worlds of data storage and security. An example of such an article is the one titled "Optimizing Storage with Network File Virtualization" by Jack Norris or "New Backup Software Migration Approach" by Kelly Harriman-Polanski. As you will see even for technologies that have been around for a while, best practices are not always practiced.

We hope our treatment of these topics are found to be useful to you and expect us to continue to refine our content over the next months and years to always strive to bring you the very best to help you manage the challenge of delivering information securely and predictably.

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

Deploying a SAN to Centralize Storage Across the Enterprise

FIGHTING THE COST AND COMPLEXITY OF STORAGE

BY MICHAEL McNAMARA

THE GROWTH OF business data continues to explode along with the need to store it. Workers generate more and more e-mail messages and file attachments, users demand instant access to data like never before, IT managers install more storage-hungry applications, and aging paper-based data continues to be converted into digital form. Information growth is so intense, in fact, that spending on data storage is expected to outstrip server spending.

However, with IT managers facing flat or shrinking budgets, the pressing challenge for them is to do more with less – to squeeze the most data storage out of every IT dollar. To achieve this objective, they must start by assessing all data storage costs – those tied to initial equipment acquisition, as well as those for resource management, capacity use, and most importantly, system downtime.

Three options exist today for managing data: Direct Attach Storage (DAS), Network Attach Storage (NAS), and Storage Area Networks (SAN).

Direct Attached Storage (DAS) represents the status quo in many organizations that aren't aware of the hidden costs or technology limitations related to this form of implementation:

> ***Difficult to Manage*** – Data is dispersed over many servers, which increases the personnel cost for supporting the organization with online configuration management and backup/restore capabilities.
> ***Limited Asset Utilization*** – Since each server owns the storage connected to it, DAS makes it almost impossible to share storage assets across multiple servers.

> ***Low Scalability*** – Server scalability is limited by the number of I/O buses supported and the SCSI bus maximum of 15 devices.
> ***Limited Distance*** – SCSI implementations typically have a 12-meter limit, which doesn't provide flexibility or let storage assets be located in secure locations in a facility or on a campus.

Network Attached Storage (NAS) is an attractive alternative to general-purpose computers, but has limitations that constrain customer configurations:

> ***Performance Constraints*** – Based on workload, the NAS box's performance can be constrained by CPU power, network throughput, and storage I/O bottlenecks.
> ***Bandwidth Requirements*** – Network bandwidth for the NAS server can compete with the other computer resources on the network.
> ***OTLP/Database Bottlenecks*** – NAS excels at file-based access but can be bottlenecked on OLTP applications and database block-level driven applications.

Storage Area Networks (SANs) represent a topology for connecting storage assets directly to the network and establishing a peer-to-peer server/storage implementation. SANs have historically been based on Fibre Channel and can now also incorporate iSCSI as a method of server/storage communication. SANs solve multiple issues for large enterprises with data centers and remote facilities and meet the IT requirements of SMB environments.

For years adding storage meant buying additional servers, tape libraries, and disk enclosures to attach to the server – a costly and inefficient tactic that left large amounts of storage capacity and computing power unused. Today, SANs – high-speed networks that connect multiple storage devices so they can be accessed on all servers in a local area network (LAN) or wide area network (WAN) – have proven to reduce management costs as a percentage of overall storage costs. Other benefits include:

> Increased disk utilization
> Reduced data center/rack floor space
> Improved data availability
> Improved LAN/WAN performance
> Reduced storage maintenance costs
> Improved protection of critical data
> Reduced CPU loads on servers freeing up computing power

The iSCSI (Internet SCSI) protocol extends the cost benefits of SANs by letting users create storage networks using existing Ethernet technology, eliminating the need for costly proprietary alternatives such as Fibre Channel (FC). With iSCSI, expanding storage to keep pace with data growth is as simple and economical as buying a disk array or adding drives to an existing disk array.

One way to combat the increasing cost and complexity of storage is to consolidate it in a single pool with fewer storage devices shared among multiple servers. By consolidating storage in a SAN you can:

> Reduce the number of physical devices to manage
> Reduce complexity
> Centralize storage management tasks
> Simplify growth and expansion
> Maximize storage utilization and return on investment

Large enterprises have adopted SANs for these reasons, but smaller enterprises and departmental IT organizations have waited to move to SAN-based storage because of

—continued on page 25

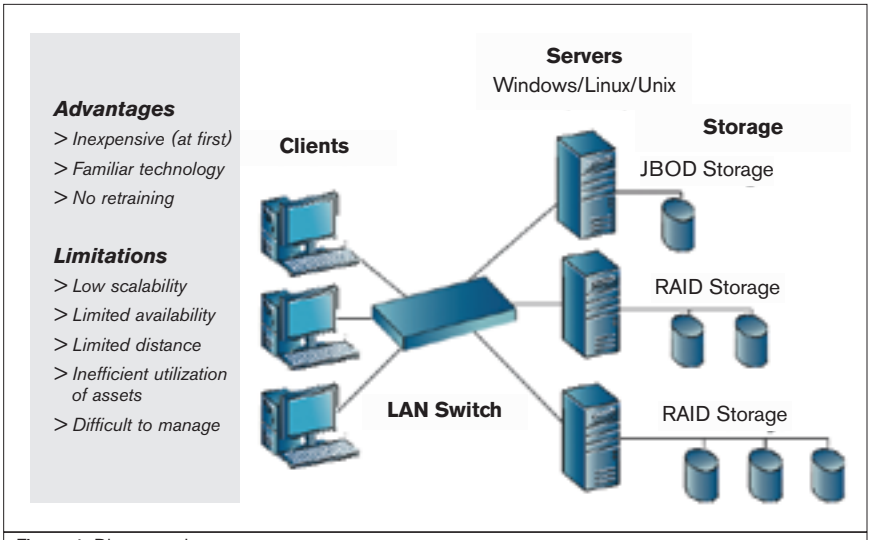

Figure 1: Direct attache storage

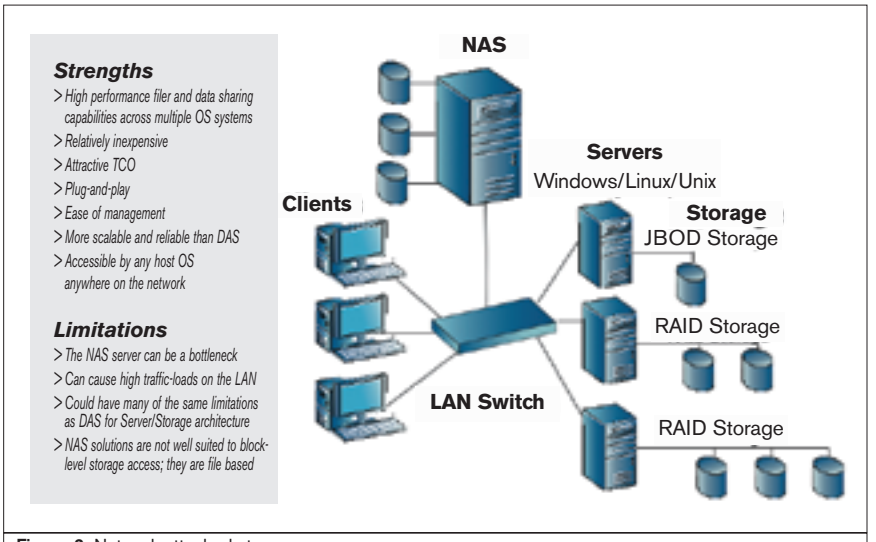

Figure 2: Network attached storage

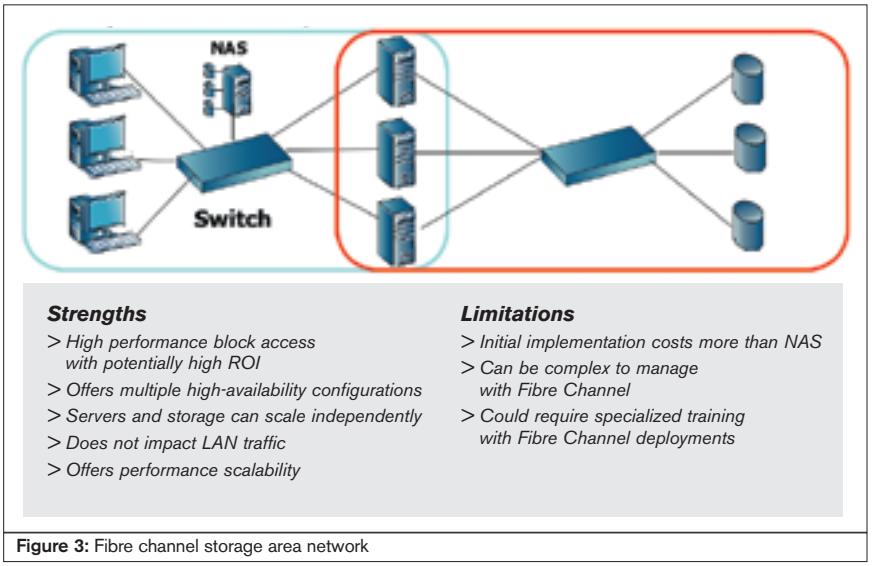

Figure 3: Fibre channel storage area network

Filtering Out Spam and Scams

A BELT AND SUSPENDERS APPROACH

BY GARY CANNON

IN SPITE OF legislation and the first conviction of a spammer under that law, it appears spammers will keep spamming as long as there's money to be made.

According to Symantec's September Internet Security Threat Report, one of the most comprehensive analyses of trends in cyber security activity, spam made up more than 60% of all e-mail traffic during the first half of 2004. And Jupiter Research estimates that the average consumer will get 2,000 spam hits a day in 2005, up from 40 in 1999.

Spam is no longer simply a time-consuming irritant. Today's spam is blended with malicious threats such as viruses, worms, spyware, and phishing scams. Now accidentally clicking on a spam message can open a Pandora's box of trouble, from activating a Trojan horse to turning your PC into a spam-sending machine.

For business, the economic impact of spam and spyware is all too clear. Not only do these threats impact productivity, network bandwidth, hardware resource, and support, they introduce serious legal liability issues and undermine hard-earned corporate brands and reputations.

In the face of such a threat, what's a concerned business to do? Problems such as spam and spyware threaten to undermine the integrity of its information. While corporate information has to remain secure and reliable, it must also remain available. And because spam and spyware use the same vehicle — the Internet — as legitimate business-critical communications, the challenge is to ensure that necessary information exchange continues while unwanted activity is halted.

Keeping spam, spyware, and other threats out of the workplace requires a powerful combination of information

security technologies, including anti-spam, anti-virus, firewalls, and policy management.

Today's Spam Attacks

Spammers now use a number of tactics to evade detection by anti-spam solutions with only limited filtering abilities. As a result, the most effective anti-spam solutions use a variety of filtering techniques to stop complex spam attacks in real-time — without compromising accuracy. Essential filtering technologies in an anti-spam solution include:

- ***Reputation Filtering:*** Reputation filtering vets the quality or reputation of the sending source or mail server of a message. This kind of filtering can identify Internet protocol addresses of suspect servers or the open proxies spammers use as well as servers that don't send spam.
- ***URL Filters:*** URL filters, in turn, identify spam URLs in messages and remove characters that conceal a Web site address in a message. This kind of filtering is effective against disguised URLs, extreme randomization, and short messages.
- ***Heuristics Capabilities:*** Heuristic capabilities are characterized by programs that are self-learning. In other words, they get better with experience. Heuristics offer a effective defense against new spam by analyzing the header, body, and envelope information of incoming messages looking for distinct spam characteristics such as excessive exclamation marks or capital letters. While poor heuristics do little more than create an administrative burden by producing countless false positives, the best heuristics can result in near-perfect accuracy.
- ***Signature Technology:*** Signature technology also plays an important role in filtering out spam. The most advanced signature technology actually strips random HTML from spam and counteracts the variations that spammers often insert, which can be a potent answer to today's highly randomized, HTML-based spam attacks. Similar signature technology is also used to identify embedded images, executables, zip files, and other message attachments through which spammers entice recipients.
- ***Foreign Language Identification:*** Foreign language identification is another essential spam filtering technique that can identify the 10%-20% of global spam not sent in English.

Mixing It Up

Effective protection against today's complex threat landscape, where spam is blended with malicious threats, requires that organizations employ a combination of information security solutions.

Anti-virus technology works to identify viruses, worms, and spyware, which are often distributed through spam. When updated regularly and configured appropriately, anti-virus solutions can automat-

imation

Selling used tape can be like giving your confidential data to a competitor

What does this mean for your data center?

Imation recommends secure tape destruction as a best practice method for managing excess tape inventory, tape format migrations or obsolete/retired tape cartridges. Companies that consider selling their used cartridges may face IT control risks that could lead to a Sarbanes-Oxley related internal control deficiency or violate data privacy regulations.

Adopt a "Don't Get Used" policy

Buying used tape media can look like a compelling financial offer, but keep in mind that the history of used tape cartridges is unknown. You don't know what you could get. Don't introduce an unknown variable into your data center environment—always specify "new" on your purchase orders. If you suspect a recent new tape puchase was fulfilled with used cartridges by your dealer, contact Imation —we'll analyze your tape(s) free of charge to determine if they have been used before.

Selling used tape can:

- Expose confidential financial records
- Disclose personnel information
- Reveal corporate IP information

Buying used tape provides:

- No manufacturer's warranty
- Unknown cartridge history
- Potential for dust and debris to be introduced into your data center

Not only may your corporate data be at risk, but the practice of selling used media may expose your company to compliance violations.

Find out how you can qualify for free tape destruction through Imation's Alliance program. Provide your contact information and we will have an Imation Account Manager contact you with details.

* Available to Alliance program members.

1) Learn more about the risks of selling or buying used tape cartridges

2) Find out how you can qualify for free tape destruction as part of Imation's Alliance program

3) Contact Imation to analyze suspect cartridges free of charge to determine if the new cartridges you purchased have actually been used

Visit www.imation.com/usedtape to learn more

ically delete or clean malicious messages, including mass-mailing worms that can result in hundreds of spam messages.

Firewalls that are configured to allow only authorized outbound traffic can also reduce the threat of spyware and malicious code that attempts to phone home over the Internet without the user's knowledge or permission or tries to launch fraudulent applications. Firewall rules can be created to block access to known spyware sources.

Corporate information security policies can be updated to ensure that file-sharing and other software is correctly implemented and that appropriate usage policies are in place and are followed. Many of the best Internet firewalls and advanced anti-virus applications are circumvented by careless or uninformed employees who haven't been trained to recognize and respond to Internet threats. In developing and disseminating a solid up-to-date information security policy, employees are educated and reminded of their role in fighting invading threats. A number of policy management tools are available to streamline this ongoing process, making it easier and less time-consuming to achieve and demonstrate company-wide compliance.

"Keeping spam, spyware, and other threats out of the workplace requires a powerful combination of information security"

Information security technologies provide a sophisticated and effectual deterrent of information security attacks that threaten to undermine the integrity of business-critical information. By using the most innovative and powerful anti-spam filtering techniques together with anti-virus, firewalls, and other security technologies, organizations can protect the security and availability of their business information while new generations of Internet threats emerge.

About the Author

Gary Cannon is president and co-founder of AIS and has over 32 years of technical and managerial experience in computer and communication systems, networks, and security. He is a Certified Information Systems Security Professional and a Symantec Certified Security Practitioner. Gary has an MS in software engineering from Colorado Technical University and an MBA in information systems from the University of Colorado. He is a member of the Symantec North American Partner Advisory Council, the Information Systems Security Association, and the Armed Forces Communications-Electronics Association.

Storage

—continued from page 5

their concerns about SAN cost and the skills required. But several factors are now falling into place to make SANs a viable option for smaller enterprises and departmental units with limited IT resources.

- ***Cost-effective SAN Bundles*** – The cost of the switches and host bus adapters (HBAs) required to build a SAN have been dropping and vendors have been creating bundled solutions specifically geared to smaller SAN implementations.
- ***Advances in Disk Technologies*** – SATA disk drives are becoming more and more common in storage arrays, increasing in capacity (400GB today) and incorporating more enterprise-class features in next-generation releases. SATA drives offer significant cost savings over high-performance Fibre Channel drives. In addition, systems designed with a SAS midplane can support both SATA and higher-performing, more reliable Serial Attached SCSI (SAS) drives. For example, an array could be configured with six SAS drives partitioned in one storage pool or LUN and assigned to a server with a transaction-intensive application such as a reservation system, and the other six drives in the array could be SATA and partitioned in a storage pool or LUN and assigned to a different server running a reference application such as medical imaging. This flexibility isn't possible with Fibre Channel and Parallel SCSI subsystems and provides customers with the best of both worlds.
- ***Evolving Management Standards*** – As the industry moves to management standards such as the Storage Networking Industry Association's Storage Management Initiative Specification (SNIA SMI-S) and Microsoft's Virtual Disk Service (VDS), tools are evolving to simplify the management of SAN environments. With a SAN in place, organizations can consolidate the storage existing on multiple storage devices on to a few larger devices shared by many servers.

In summary, to achieve the benefits of a SAN, organizations need:

- Tested and validated SAN configurations that are easy to install
- Scalability to address growing storage requirements
- Simplified management of both systems and components
- Flexible solutions that can be tuned for specific markets and applications

The benefits of a SAN far outweigh the alternatives, and as their cost drops and complexity lessens with advances in technology, SAN adoption in the SMB market will increase.

About the Author

Mike McNamara is manager of product marketing at Adaptec, the provider of end-to-end storage solutions. He has over 16 years of marketing experience in the computer industry, with over 10 years in storage. Mike was graduated from Fairfield University with a bachelor's degree in business management and has an MBA from the Clark University Graduate School of Management.
Michael_McNamara@adaptec.com

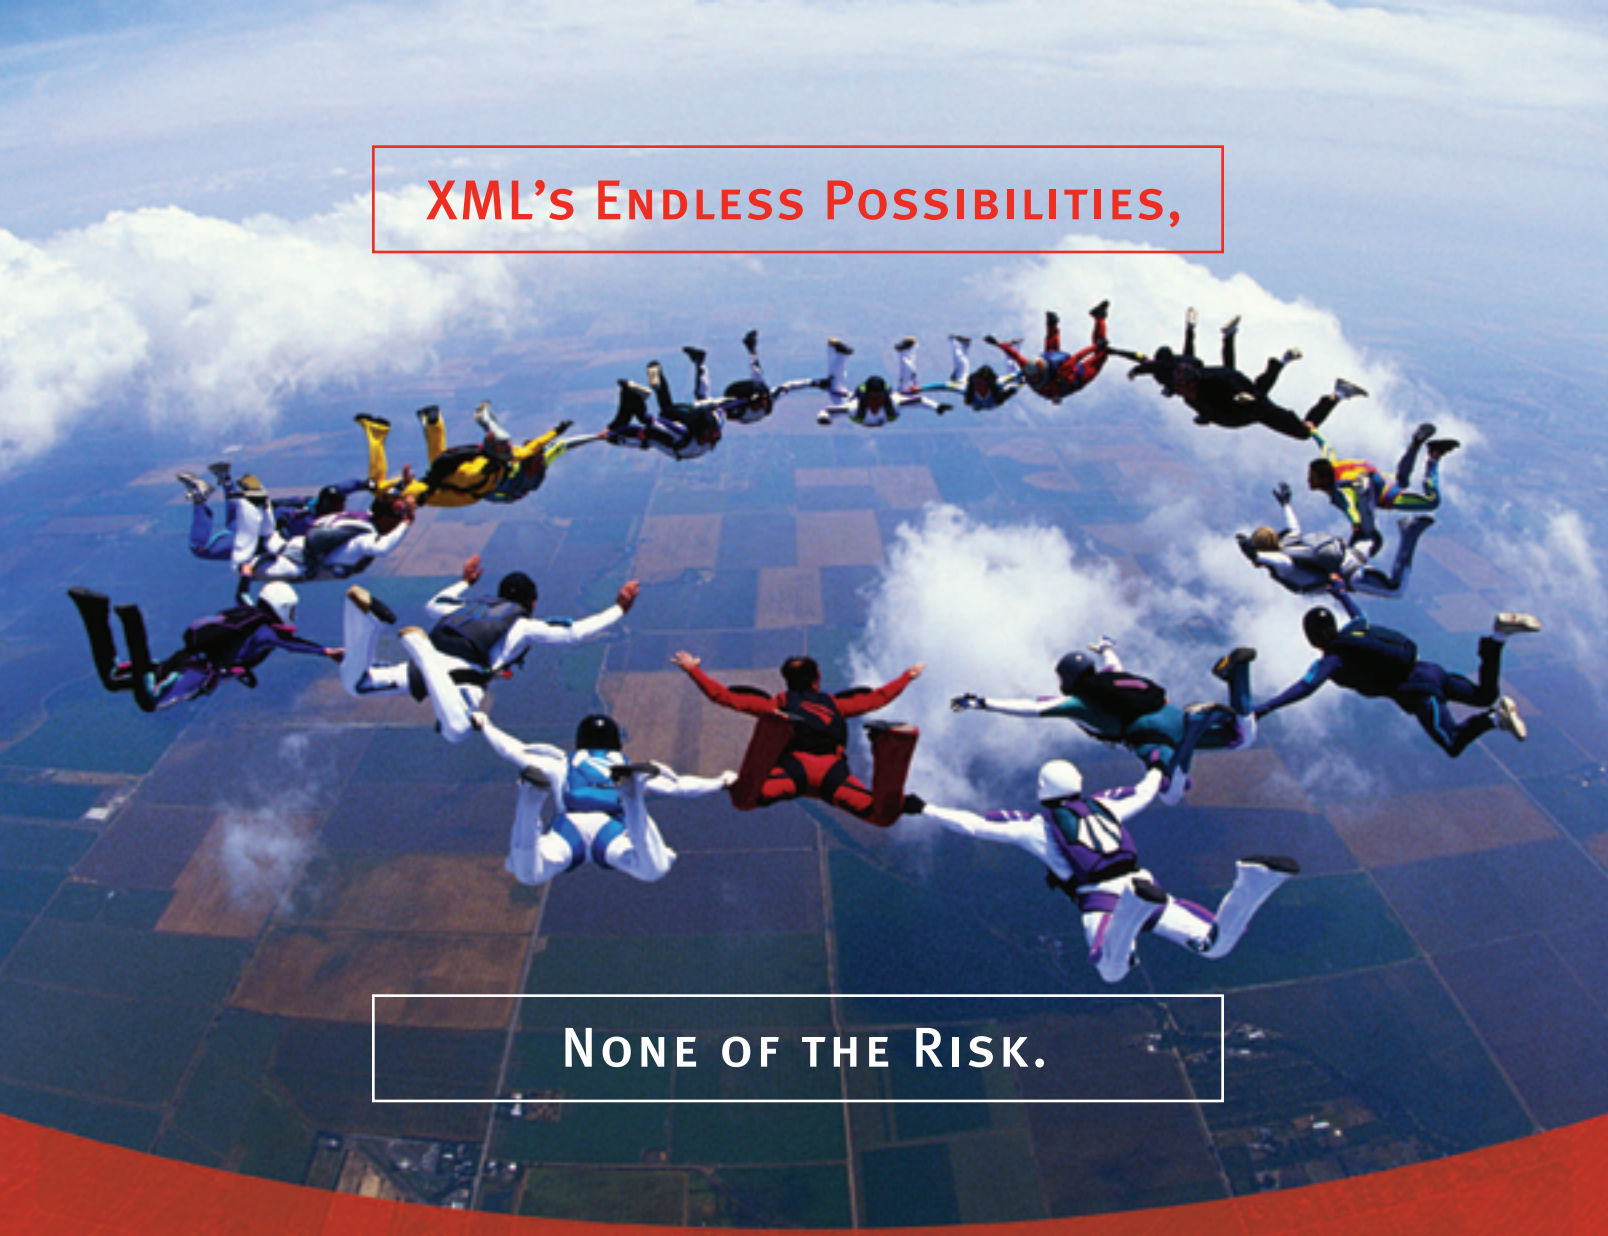

Forum XWall™ Web Services Firewall - Reinventing Security

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

Trusted Computing: Flip the Switch and Help Your PC Protect Itself

YOUR DEFENSES ARE BUILT-IN

BY STEVEN SPRAGUE

PERSONAL COMPUTERS HOLD treasure troves of confidential and personal information ripe for the picking by hackers, thieves, and scammers. Patient records, consumer credit card information, invaluable R&D data, personal finance...we've become increasingly reliant on computers, and need powerful security to protect the confidential data, hard work, and critical information contained in our PCs. Despite major advancements in systems security over the past several years, analysts and industry experts quantify global economic damage from digital risks exceeding a record-breaking $500 billion in 2004.

The PC industry has observed this pain, and introduced powerful new tools to enhance the security and privacy of your network. The barrier standing between your crucial assets and malicious intruders is about to get better.

Since 1999, a core group of leading PC manufacturers, hardware, and software vendors have been hard at work creating a hardware-enabled standard for improving the security of every type of computer — from desktop and laptop PCs to handhelds and other devices. This group, the Trusted Computing Group (TCG), has combined expertise from more than 100 companies including Dell, Intel, AMD, Microsoft, HP, and Wave Systems. The resulting breakthrough is a hardware security chip called the Trusted Platform Module (TPM), which helps ensure that your computer, no matter where you're using it, is more secure... even if lost or stolen.

Tens of millions of TPM chips have already quietly shipped, and "Trusted Computing" capabilities are now embedded in computers worldwide. Observers say that TPM deployment is on the verge of exploding, with IDC estimating that by 2007, up to 55% of computers shipping worldwide will contain TPMs. When leveraged with appropriate software, Trusted Computing offers protection from identity theft, information leakage, sensitive data exposure and other security risks, making your computer — and your business — more secure.

Today, most computers rely solely on software to shield their data — passwords, data encryption, firewalls — but, the software is inherently insecure, as seen through almost constant attacks, providing ample room for theft, hacking, and data loss. The Trusted Computing model allows for the standards of software security to be amplified by the newly intrinsic secure hardware.

A common and very real threat is that unauthorized persons access data stored on a PC. The consequences of unauthorized access can include legal penalties (the exposure of a customer's personally identifiable information), competitive disadvantage, embarrassment, fraud, and extortion. Managing platform data is a responsibility of the business. The data security solutions provided by TPM and the right software enable owners of data and applications to impose strict controls on who can access and use those assets.

It's critical that enhancing data security not compromise functional integrity. The new wave of encryption appearing through Trusted Computing ensures that data in any format is both accessible and more secure. This includes transparency for the end user — the data remains encrypted without constant action from the end user – and authenticated access.

Authentication via passwords is the standard model used today for everything from Web site access to transaction authorization. But passwords are only as secure as a hacker's ability to guess, record keystrokes, or fraudulently get them. Experts blame weak and insecure passwords for unauthorized financial transfers, privacy breaches, identity theft, and even the hacking of corporate networks.

Trusted Computing eliminates this threat by adding a second factor of authentication that strengthens the entry point to the PC, application, network, or data being accessed. If the password is stolen, it's useless — the password alone isn't enough to gain access to a Trusted Computer's valuable data.

Included in the valuable data is information that lets hackers steal your customers' identities. As identity theft and unauthorized access reach unprecedented levels, businesses and consumers

Figure 1

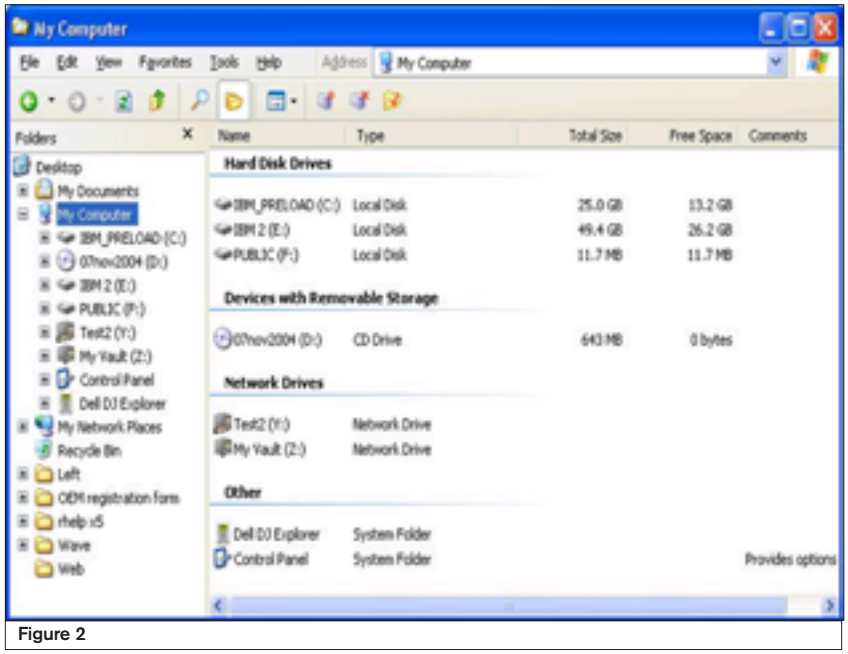

Figure 2

are devising stronger means to safeguard personal identities, specifically to combat the great vulnerability that lies with electronic identities.

Digital certificates are commonly used as proof of identity for access to networks, data, and services. The keys tied to certificates are also the basis for digital signatures. Theft of a digital certificate offers substantial opportunity for crimes of fraud and unauthorized access. Fraud or forgery using a stolen digital signature isn't easy to prove. Since a digital certificate could be stolen by making a copy of it, it could take the owner some time to realize a theft had occurred. It's

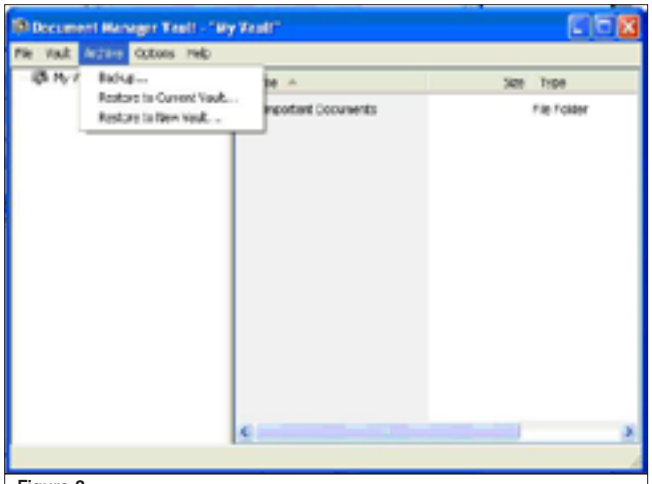

Figure 3

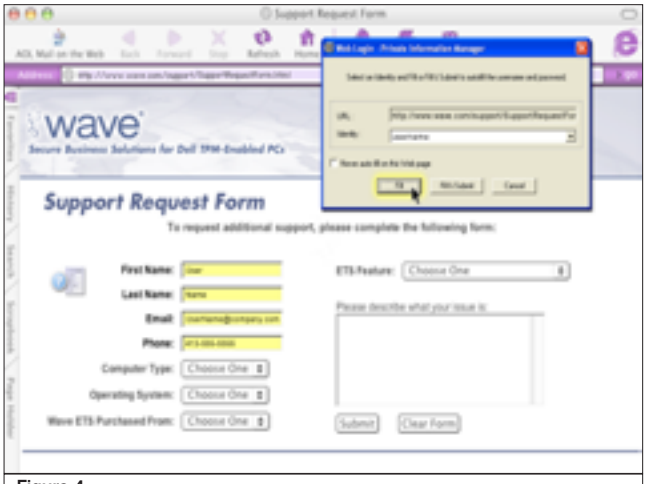

Figure 4

"Leveraged with the right software, Trusted Computing protect against identity theft, information leakage, sensitive data exposure and other security risks, making your computer – and your business – more secure"

extremely important to provide the best possible security around the storage and use of digital certificates. Using Trusted Computing standards for hardware-protected digital certificates provides a safeguard against theft by storing cryptographic keys and other data securely and away from traditional storage.

Perhaps one of the most critical features of the TPM is the flexibility it offers businesses and their mobile employees. Secure authentication makes sure that you — and no unauthorized users — have access to your system and its services. Through the integration of a trusted platform into a corporate network, a company can ensure better and stronger platform identity. For example, when an employee works externally or accesses a corporation's network remotely, the company can control the access from outside sources, monitoring the identity of every platform and only allowing valid users logging in from valid platforms to sign in. Secure authentication means continued productivity without sacrificing security.

Finally, a core function of the TPM is to be able to measure the key software components such as the operating

system and security software running on the PC to determine if they're still in a known and trusted state. This will enable better detection of viruses, trojans, and systems that have been compromised by attacks. Knowing a platform's trustworthiness is a key requirement for letting remote systems into corporate networks and participating in high-value transactions and sensitive Web Services.

The vision of an industry standard for security has been forming for years. We're now on the cusp of its fruition, with shipments of TPM-enabled computers having reached a critical mass where users can "flip the switch" and recognize the benefits that come from cutting-edge security hardware welded right to the motherboard of their computers. Powerful software is now available to leverage the increased protection, introducing a whole new era of security. The hacker's job is about to get a lot harder — with improved security built in directly, computers can finally secure and protect themselves more effectively.

This is the compelling case for suggesting that all PC purchases going forward should be Trusted Computers.

Resource

www.trustedcomputinggroup.org

About the Author

Steven Sprague is president and CEO of Wave Systems Corp. A pioneer of the Trusted PC movement, Steven has spoken and presented at more than 50 industry events, sharing his expertise on trusted computing applications and services. Wave Systems has a portfolio of fundamental patents in security and e-commerce applications and employs some of the world's leading security systems architects and engineers. For more information on Wave Systems and trusted computing solutions, visit www.wave.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.™

ILM Is Happening – Is Your SAN Infrastructure Ready for It?

BEGINNING THE TRANSFORMATION TO INTELLIGENT STORAGE NETWORKS

BY RANGA BAKTHAVATHSALAM

THE DEMAND FOR storage will continue to grow. Endless amounts of data are being created driving greater storage capacity requirements and price improvements. That same data must be classified and moved into various tiers of storage to facilitate cost-effective implementations. Information lifecycle management (ILM) offers a set of practices and tools for managing the classification and movement of data in alignment with service-level and cost-of-ownership objectives. For ILM to deliver real end-user value, the storage infrastructure has to provide a foundation that can host the necessary tools and processes. Present day storage infrastructures fall short of providing this foundation because of inherent limitations that include decentralized management, disconnected SAN islands, and inefficient use of storage resources. Intelligent storage infrastructures address these limitations and provide the foundation for a successful deployment for ILM.

Phased Approach to ILM Implementation

Information lifecycle management offers a set of policies and practices that let users apply values and rules to business information. Information is classified at its source based on its business value and stored (or discarded) on a device matching its asset value. For ILM practices to be implemented in a real environment, a number of steps have to occur. The core requirement of the ILM strategy is that an enterprise must understand the relative value of its information and how that value changes over time. This understanding provides the ability to classify and store information and achieve the service-level objectives (SLAs) established.

From the deployment perspective, ILM is all about the classification and movement of data from one storage medium to the other based on its asset value. Initial implementations of ILM are being carried out with tools that exist today but for ILM for deliver its full benefits, the deployment has to be well planned and carried out in multiple phases. To enable a streamlined deployment, Storage Networking Industry Association's (SNIA) Data Management Forum (DMF) is suggesting a multi-phased approach as shown in Figure 1.

The first phase of ILM deployment is to deploy storage on the network and provide a centralized management scheme for the storage services. While deploying networked storage is a common practice in large enterprises, current infrastructures fall short of providing the centralized management of storage resources and services.

Storage Infrastructure Requirements

ILM requires the storage infrastructure to support classification and, more importantly, the movement of data from one storage medium to the other. The infrastructure should also enable storage resources to be allocated on-demand and support non-disruptive data migration to meet service-level objectives. To facilitate these ILM requirements, the storage network should support for the following:

- ***Multi-tiered storage*** – The infrastructure must support storage devices ranging from high-end arrays and tape libraries to low-cost storage disks. This enables the movement of data from one tier to another based on the business value of the information.
- ***Heterogeneous storage*** – The infrastructure must support storage devices from multiple vendors. Interoperability across vendors enables the migration of data from one vendor's storage device to another.
- ***Multiple protocols*** – The infrastructure must support connectivity across storage networks that have implemented protocols such as Fibre Channel and iSCSI.
- ***Multiple applications*** – The infrastructure must enable storage and data management applications to be implemented on the network.
- ***Efficient utilization of resources*** – The infrastructure must enable efficient use of resources by supporting dynamic resource allocation and reallocation.

Drawbacks of Present Day Infrastructures

Current infrastructures fall short of meeting these requirements because of

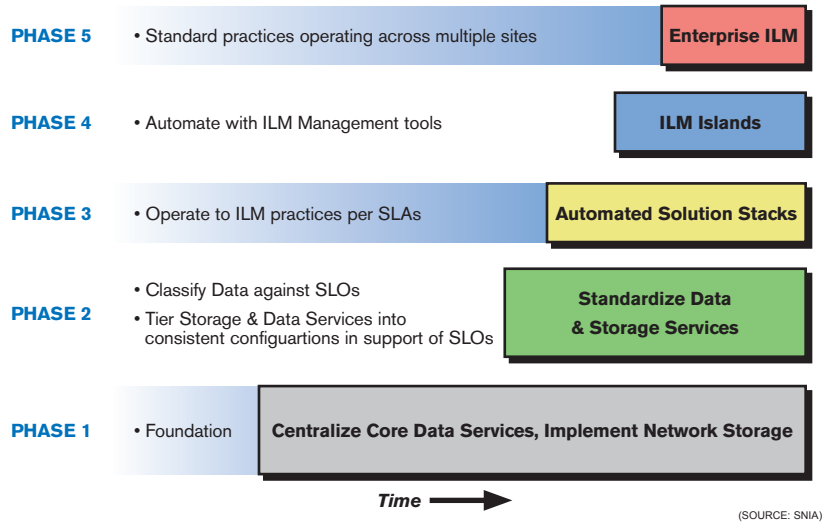

Figure 1: Phased implementation of ILM practices

inherent limitations. Storage arrays from different vendors don't interoperate, locking in customers to a single vendor. Large deployments of SAN infrastructures have led to multiple disconnected SAN islands forcing customers to deploy high-cost resources, tape libraries, for example, in multiple SANs resulting in the underutilization of expensive resources. In cases where SAN islands are connected by a simple switch, the resulting "merged SAN" gives rise to reliability concerns because of changes in network configurations and the limits of the infrastructure's scalability.

Though SMI-S, when implemented, will address the interoperability issues between management application and storage resources, as a management interface it doesn't address the drawbacks related to efficient resource utilization, seamless data migration, or protocol connectivity.

Intelligent Storage Networking

Intelligent storage networking addresses all the current limitations of the infrastructure and lays the foundation for a successful rollout of ILM. Intelligent storage networking involves installing and administering intelligent SAN platforms in the storage network to optimize the transfer of data between servers and storage elements and the transfer of data among storage elements.

Intelligent SAN platforms address the infrastructure requirements of ILM by enabling applications such as network-based virtualization, data movement, and data replication. Network-based virtualization and data management applications enable highly efficient storage resource utilization and enable the movement or replication of data in a manner that's transparent to the applications. Thus virtual volumes (storage) used by the applications can be moved transparently from expensive arrays to inexpensive storage based on their asset value, or replicated to a remote location to meet service-level objectives.

Intelligent SAN platforms also address network connectivity requirements by enabling routing across SAN islands. This results in network infrastructures that are highly scalable and independent of the underlying protocols such as Fibre Channel or IP.

By offering highly optimized storage I/O processing, intelligent SAN platforms such as intelligent switches, directors, and storage appliances are enabling the core requirements of ILM that include centralized management, high resource utilization, and high availability. Successful implementation of ILM practices depend on having this foundation layer in place.

Migration to Intelligent Storage Networks

Storage administrators will have to start with getting the right storage infrastructure in place when considering implementing ILM. Administrators are often presented with two broad choices for deploying intelligent storage infrastructures, i.e., storage appliances and intelligent switches/directors. Storage appliances offer storage applications running on general-purpose servers or custom-built hardware that enable I/O acceleration. Intelligent switches or directors supported by the co-existing storage application offer I/O acceleration and enable more scalable deployments. Administrators will have to understand the choices and ensure that their infrastructure upgrades are in sync with their ILM objectives.

Conclusion

Executing ILM practices brings about storage optimization, efficient data protection, increased management efficiency, and overall cost reduction. These advantages are driving end-user implementations. Intelligent storage infrastructures provide the necessary foundation for implementing ILM tools and practices. Intelligent SAN platform vendors are proposing multiple alternatives for deploying intelligent storage infrastructures. Storage administrators should understand the choices and plan their infrastructure upgrades accordingly to begin the transformation to intelligent storage networks. ■

"Current infrastructures fall short of providing the centralized management of storage resources and services"

About the Author

Ranga Bakthavathsalam is product manager at Aarohi Communications.

ranga@aarohi.net

Optimizing Storage with Network File Virtualization

UNSTRUCTURED DATA MANAGEMENT WITHOUT THE HASSLE

BY JACK NORRIS

EXISTING STORAGE MANAGEMENT methods and tools can't keep pace with exploding storage requirements. Application expansion, digital media formats, and regulatory compliance have all contributed to the fast-growing demand. According to IDC, storage administrator productivity has to increase 60% a year just to keep up with the anticipated growth in storage capacity. To make matters worse, 24x7 data access requirements are closing the management windows available to administrators to do management tasks.

The Challenge

A new approach is needed to dramatically simplify network storage management and drive improvements in capacity, performance, and tiered storage management. Any potential solution, however, mustn't introduce additional risk into storage environments. Administrators don't want to risk data integrity problems, disaster recovery issues, or performance bottlenecks.

A related concern for any viable solution is how end-user access is managed. A solution that requires changing end-user mount points or installing special software on each server or client can significantly outweigh the benefits of uninterrupted access during data movement. Organizations also have to understand how a solution impacts other management tasks. How difficult is the initial deployment? How does it do data retention? Will a potential solution create management headaches elsewhere or in the future? For example, does a solution take a proprietary approach that conflicts with upcoming industry standards?

Network File Virtualization: Changing Storage Management

Virtualization is key to managing demanding file storage requirements. Rainfinity is the first company to optimize IP-based storage with network file virtualization (NFV) that enables unstructured data management without disrupting end-user or application access. Its patented Network File Virtualization Platform optimizes Networked Attached Storage (NAS), eases storage management overhead, simplifies end-user access, and enables additional storage management functionality. NFV lets administrators support heterogeneous storage environments and optimize networked storage across different vendor platforms increasing flexibility and lowering TCO.

Strategic approaches like ILM, Storage Grid, and Utility Computing require data to move freely across the environment without disrupting end-user or application access. Standards-based NFV offers this required non-disruptive data movement across heterogeneous environments, which represents a significant change in an administrator's ability to manage NAS and file server environments effectively.

In traditional file server environments, adding data storage greatly increases management burdens and data exposure, and impacts performance and availability. NFV lets administrators efficiently manage their NAS networks across large-scale multi-vendor storage environments regardless of size, regulatory requirements, data volumes, service levels, or high-availability requirements.

Optimizing Storage with Network File Virtualization

- Management Applications Combined with Network File Virtualization – One product, for example, RainStorage from Rainfinity, combines network file virtualization with purpose-built applications to simplify storage management, increase flexibility, and lower cost. It uniquely optimizes networked storage with applications that identify, analyze, and resolve capacity, performance, and tiered storage issues. Existing tools on the market provide monitoring capabilities and can provide a great deal of information about the status of a network storage environment, but these tools can't take action and can't actively manage active data. Active data management is the focal point. The applications described below let administrators optimize storage instead of simply checking status.
- ***Capacity Management*** – Automatically identifies over-allo-

"I felt like grabbing him by the throat."

I was delivering the results of a security review, essentially "white-hat hacking", and my client was trying to justify why he hadn't changed his service account passwords in two years. Too busy, no cost justification, etc…

I then asked him how many people had left the IT group in the past two years. "About 5," he said. I about jumped out of my skin. "And why haven't you changed the passwords", I asked. "I didn't see the need", he responded. I felt like grabbing him by the throat and shaking him as hard as I could. No auditor would accept this; it could easily get him fired. I just had to leave the room.

"Hands down, without a doubt the single most common mistake I see is not managing Service Accounts properly."

Do me a favor. Take a couple of minutes and download this white paper at www.e7software.com/risk and start to understand the importance of having properly managed service accounts.

Your friend,

The Cyberspace Samurai

Ps: Hey, you might even win a free gift when you download the white paper.

Pps: Check out my blog. www.cyberspacesamurai.com

e7software™

cations at the file server, volume, or quota tree level and takes corrective action. RainStorage presents the top-capacity issues from left to right so administrators quickly understand the environment and easily identify capacity problems at the file server, volume, and directory level. A simple click lets them analyze issues in more detail and immediately resolve issues on-demand.

- ***Tiered Storage Management*** – Analyzes access and puts content on the most appropriate storage tier; identifies the least accessed directories on the online storage tier, and identifies the most accessed near-line directories to determine which content can benefit from relocation to the online tier; transparently supports service levels without disrupting storage systems.
- ***Performance Management*** — Identifies file server CPU, volume, and directory process bottlenecks and resolves issues by dynamically distributing content to alternate locations to balance performance better. RainStorage presents the top capacity issues so administrators can quickly understand the environment and easily identify capacity problems at the file server, volume, and directory level.

Capacity Issues

> Average utilization is 35-50%
> Management cost per TB isn't improving
> Over-provisioning is too costly

Performance Issues

> User productivity impacted by poor response time
> Application throughput limited by I/O bandwidth
> New devices don't address the bottlenecks

Tiered Storage

> Nearline storage represents a huge cap-tal expenditure advantage over online
> Limitation to dynamically manage data between online and nearline

Storage Consolidation

> Consolidation projects slip because of the organizational impact
> Consolidation requires data relocation along with security and access settings

Increased Utilization

A large ISP increased storage utilization to 90% and is able to respond to capacity issues without disrupting end-user data access

Productivity Impact

One of the largest US financial services firms saved over a half-million dollars in labor alone through network file virtualization

CAPEX Savings

A major semiconductor company cut storage CAPEX by 50% by using SATA enabled by NFV

Business Efficiency

A Fortune 500 firm completed a 30TB consolidation project in a tenth the estimated time with a four-week payback

Network File Virtualization – The Benefits

NFV not only helps administrators meet current storage management challenges, NFV drives tremendous benefits.

> ***Ease of Deployment*** – RainStorage plugs easily into existing networks without requiring configuration changes to NAS servers, filer servers, clients, application servers, or storage management tools and utilities. Additionally, no proprietary client, server, or filer hardware or software has to be deployed.
> ***Storage Consolidation*** — Pools multiple file servers so they look like one file server either permanently or temporarily as part of a migration project. It also handles the associated complexity surrounding security, permissions, and domains and automates security ID translations and access control settings to ensure that the data, security, and metadata are moved correctly, quickly, and transparently.
> ***Synchronous Mirroring*** — Synchronously mirrors file data across multi-vendor IP-based storage environments. RainStorage creates a synchronous mirror across NAS and file servers to protect production data that can't tolerate any loss. Mirroring data to a remote site provides for rapid recovery so that business can continue in the event of a disaster.

Enterprise Scalability through NFV Network Processing Layer

NFV architecture should include advanced multiplexing and de-multiplexing capabilities and network fastpath processing to process selective traffic efficiently. During data transfer, network processing also moves RainStorage in and out of the data path to provide high-speed throughput and transparent redirection. RainStorage's architecture enables it to move completely out-of-band when no optimization transactions are executing, maximizing data throughput and response time. Layer supports all devices that are accessed via industry standard NFS or CIFS protocols and enables both protocols' access from a single appliance.

Complete Data Integrity

Data safety is the most important factor for storage administrators. It's not only how scalable and heterogeneous a solution is but how well data is protected. With RainStorage's unique transaction-based processing there's complete data integrity with no persistent metadata, no single point of failure, and no disaster recovery exposure. RainStorage manages open files and open locks and guarantees data integrity at all times. To ensure the ultimate safety, RainStorage adopts a transaction model in managing data reorganization so that any system failure in the middle of data reorganization doesn't affect data integrity.

Summary

Network File Virtualization is changing network storage management. Rainfinity is the first company to combine a patented NFV platform with purpose-built applications so organizations can easily simplify management, increase utilization, decrease over-provisioning, resolve performance bottlenecks, leverage tiered storage, and lower TCO. The Network File Virtualization platform is a key building block for utility computing, storage grid strategies, and ILM.

About the Author

Jack Norris is vice-president of marketing at Rainfinity.

It's not a fantasy, it's real!

Universal User Based Storage Management (UUSM) is providing corporations with the platform independence they've been wishing for.

Storage management is an enterprise problem. Compliance regulations, ILM, and the RIAA have caused senior executives to become focused on storage management. NTP Software's family of products utilizing "Universal User Based Storage Management" is the solution to enterprise storage – it manages Network Appliance, EMC, HP, IBM, and Dell machines, whether it is a SAN, NAS or RAID configuration.

No more duplication of efforts, no more partial information.

You've got plenty to do. You don't have time to stress about users saving illegal files or consuming so much storage space that it could cause the server to crash. These just shouldn't be on your radar. No one in their right mind wants to do that kind of work anyway, it's boring, tedious and you're always having to battle with users. Take control and get back to spending time on the projects that are more fulfilling....remember, the ones you were hired to do.

Take complete control of your all your storage!

Download our free report at **www.ntpsoftware.com/learn**, and learn about what most of the global 2000 is already doing.

You may also qualify for a $10,000 Storage Review that will give you just the information you need to cure those storage headaches.

NTP Software™

Tokenization: The Building Blocks of Spam

HEURISTIC COMPONENTS OF A STATISTICAL SPAM FILTER

BY JONATHAN A. ZDZIARSKI

This article is an excerpt from Ending Spam: Bayesian Content Filtering and the Art of Statistical Language Classification. *Printed with permission from No Starch Press. Copyright 2005.*

UNLIKE OLDER SPAM filters, in which the author programs the characteristics of spam, statistical filtering automatically chooses the characteristics (or "features") of spam and nonspam directly from each e-mail. Two years from now, when spam has evolved in content, statistical filters will have learned enough to continue doing their job. This is because unlike older spam filters, in which the author programmed rules to identify spam, statistical filters automatically identify damning features of a spam based on message content.

Tokenization is the process of reducing a message to its colloquial components. These components can be individual words, word pairs, or other small chunks of text. Data generated by the tokenizer is ultimately passed to the analysis engine, where it is interpreted. How the data is interpreted is important, but not necessarily as important as the quality of the data being passed. In other words, the way that a message is tokenized is more important than what we do with it later; even a simple change in tokenization can affect the accuracy of the filter. From a philosophical point of view, this raises the question, "What is content?" If content were just words on a page, then tokenizing only complete alphabetical words should be sufficient –but content is much more than that, as we'll see throughout this article.

Tokenizing a Heuristic Function

The one heuristic aspect of statistical filtering is tokenization. Even though the process of identifying features is dynamic, the way those features are initially established – how they are parsed out of an e-mail – is programmed by a human. Fortunately, languages change slowly, and only a few minor tweaks are necessary to adapt the tokenization process to handle some of the wrenches thrown at it by spammers. Tokenization is the type of heuristic process that is usually defined once at build time and rarely requires further maintenance. In light of its simplicity, many attempts are still being made to establish tokenization through artificial intelligence, to remove all sense of heuristic programming from the equation. Within a few years, filters should be able to efficiently perform their own type of "DNA sequencing" on messages, determining the best possible way to extract data. In fact, this is already being researched as a solution to filtering some foreign languages that don't use spaces or any other type of word delimiter.

Basic Delimiters

Besides deciding how best to break apart a message, there are many other issues to consider when tokenizing. For example, we need to determine what constitutes a delimiter (token separator) and what constitutes a constituent character (part of the token). Do we break apart some pieces of a message differently than others? What data do we ignore (if any)?

The fundamental goal of tokenization is to separate and identify specific features of a text sample. This starts with separating the message into smaller components, which are usually plain old words. So our first delimiter would be a space, since spaces commonly separate words in most languages. This makes it very easy to tokenize a phrase like the following:

```
For A Confidential Phone Interview, Please Complete Form & Submit.
```

which can be broken up into the following words:

```
For     A          Confidential  Phone  Interview
Please  Complete   Form          &      Submit.
```

As we've learned, each word typically is assigned one of two primary dispositions: spam or nonspam. The example above will cover a lot of text, but we're left with a few punctuation issues. For example, is the word "submit" on its own likely to have a different disposition from the word "submit." with a period after it? How about "interview" and "interview," containing a comma? In these cases, it makes sense to add some types of punctuation to the set of delimiters, as punctuation suggests a break in most languages. The following are some widely accepted punctuation delimiters:

- period (.)
- comma (,)
- semicolon (;)
- quotation marks (")
- colon (:)

Some other punctuation, such as the question mark, is a bit more controversial. Some authors believe that "warts" and "warts?" should be treated the same, in most cases as spammy tokens.

Including too much punctuation in the makeup of tokens could result in five or 10 different permutations of a single word in the database. This can very rapidly diminish their usefulness. On the other hand, not having enough tokens can cause the tokens to become so common among both classes of e-mail that they become uninteresting. The trick is to end up with tokens that would stick out in one particular corpus. If there were 100 spams about warts in the user's corpus, but only one posing a question in which "warts?" was used, the filter is likely to overlook this feature in the one message.

Note: *I've found that treating a question mark as a delimiter results in slightly better accuracy (on the order of a few messages) in my corpus testing, as opposed to treating it as a constituent character. This could likely change in the future, however.*

Redundancy

Some types of punctuation are very useful; for example, the exclamation point makes a remarkable difference between "free" and "free!" and so you want to use some punctuation marks as con-

stituent characters. One of the problems a filter author might run into when allowing these types of characters, however, is redundancy. Most would agree that there's no real difference between "free!" and "free!!!!" in a message, as both are equally condemning characteristics of spam. On the other hand, messages in which symbols are used to b!r!e!a!k up a word may behave a bit differently.

Some authors will view punctuation as part of a token only if it appears at the end of the token. If an exclamation point appears elsewhere, it will be treated as a delimiter in most cases. For those punctuation marks that are permitted, we should consider working some method of de-duplication into our tokenizer, where only the first occurrence of the punctuation is used. We essentially look at "free!!!", "free!!!!!!!!!!", and "free!" as the same token by truncating the extra chaff. I've found that using the exclamation point as a constituent character slightly improves accuracy, which is the opposite effect that question marks appeared to have. This is probably because more spams use an obnoxiously loud used-car-salesman type of pretense rather than actually posing questions. Perhaps one day, spammers will become more philosophical, and then question marks will become just as useful as exclamation points.

Some filters permit a certain window size before the token is truncated; for example, tokens may be allowed to have up to three exclamation points before being truncated, giving the filter three different meanings for "free!", "free!!", and the extremely guilty and shameless "free!!!" One of the advan¬tages to doing this, other than measuring the three levels of unbridled fervor, is that it allows a really obnoxious message that uses all three tokens to fill up more slots in the decision matrix.

It's important to truncate extraneous characters at some level because spammers could easily use not truncating them as a way to hide very spammy tokens; for example, a spammer wanting to hide the word "porn" could send "porn!!!" in the first spam and "porn!?!?!" the next time, so that in both cases the token would be considered a new token. Truncating will reduce both of these tokens to "porn!" or even "porn" if exclamation points are ignored all together. Tokens should generally be limited to only one acceptable punctu¬ation mark at the end, or to an N-sized window of homogeneous punctuations at the most.

Other Delimiters

Other delimiters used by many applications include the following:

- brackets [ ]
- braces { }
- parentheses ( )
- mathematical operators + - / * = < >
- special characters | & ~ `
- the at (@) sign
- underscores and other rare characters

These delimiters frequently prevent the duplication of several different permutations of tokens, such as "when" and "(when". Other characters, such as the new line character, are also treated as delimiters. The nice thing about the way text is delimited is that it's going to result in unique tokens, even if the tokenization isn't perfect. This can be good or bad, but most of the time it's good. Even a token that isn't in human-readable format may be machine-readable and may occur with enough frequency to be a good identifier. In fact, Bayesian antivirus filtering uses an entirely different set of delimiters, because antivirus analysis involves the cataloging and analysis of several different binary sequences.

Exceptions

Some exceptions to the basic delimiters we've mentioned involve one-off instances where we actually want to preserve certain complete tokens. For example, IP addresses make for good spam markers, as do certain HTML characters like © and . If you're reading this book, there is most likely no shortage of spam in your inbox (or quarantine). Often the best way to discover new approaches to tokenization is to take a look at some of the text spammers are using in their samples. It's very important that the token¬izing approaches being used aren't biased against present-day spam.

The tokenizing algorithm should be generic in such a way that it can easily break down any kind of natural language or new type of message style, but it shouldn't be so plain vanilla that the features it generates are likely to appear as common in all e-mail. It would be relatively easy to tokenize a message into individual characters, but that wouldn't be very useful, since the token "v" could occur in "viagra" or "violin". All-numeric tokens are generally not very useful on their own, but when combined with the proper punctuation (such as a dollar sign or exclamation mark) can make a signifi¬cant distinction between "19" and "$19" or between "95" and "95!". Provide enough information to allow the token to be set apart from the rest, but not so much that it is unlikely to show up only a handful of times.

To some degree, this anal-retentive exercise is overrated. Any reasonable level of tokenization will most likely yield levels of accuracy above 99 percent, but making a mistake could cost a few misclassifications on occasion. I've found that using the question mark as a constituent character in my tests resulted in approximately three additional errors per 5,000. Experimentation and thorough testing is one of the best ways to decide on the tokenization approach that works best for the filter.

Token Reassembly

Occasionally, tokens will turn out to be a little too small due to attempts by spammers to obfuscate them. When this happens, reassembling individual letters into a token can help improve accuracy. Let's look at an example of obfuscated text:

```
C/A/L/L/ N-O-W - I/T/S F_R_E_E
```

If the tokenizer we're using considers underscores, dashes, and slashes to be token delimiters, then instead of ending up with four one-word tokens, we'll end up with 14 single-character tokens. Many filter authors believe it's healthy to allow these individual characters to tokenize, while others believe that the resulting information is too generalized to be a good indicator of anything, at least without the risk of false positives.

Filter authors who share the latter philosophy can use token reassembly to join the original tokens back together. Token reassembly isn't a perfect science, but it provides more useful tokens to work with. The tokens "VIA" and "GRA" are much more useful than individual characters and are defi¬nitely more indicative of spam. Token reassembly basically concatenates single-character tokens that are adjacent to one another, looking for larger amounts of white space amidst the slicing and dicing to make an educated guess about what words go together. Since statistical filtering involves machine learning and not human learning, tokens like this are very useful to the computer, even though they may not make much sense to us. For example, the token "VIA" really doesn't mean much, which is exactly why it makes a great indicator of

spam – you'd rarely see the word "VIA" in a legitimate message unless you were talking about motherboards. The word "GRA" is even more rare in legitimate mail. The fact that these tokens aren't neces¬sarily comprehensible to a human makes it easier to identify them in spams. My dataset considers some of these fractional words to be extreme indicators of spam:

```
Agra      S: 00030  I: 00000  P: 0.9999
Eacute    S: 00021  I: 00000  P: 0.9999
Prematur  S: 00020  I: 00000  P: 0.9999
```

Degeneration

Another solution Graham introduced into tokenization is called degeneration. Degeneration allows a token that hasn't been seen before to be reduced in complexity (location, case, and punctuation) until it matches a simpler token. If no tokens match a given token, we make it simpler until we find a match. For example, consider the use of the word "FREE!!!" in the subject. If it has never been seen before in the subject, degeneration has us reduce the phrase until it matches something we have seen before.

```
Subject*Free!!!
Subject*free!!!
Subject*FREE!
Subject*Free!
Subject*free!
Subject*FREE
Subject*Free
Subject*free
FREE!!!
Free!!!
free!!!
FREE!
Free!
free!
FREE
Free
Free
```

Degeneration has a lot of room for customization, including the order in which the tokens decrease in complexity. At the very least, degeneration of punctuation is a wise move. If the word "free!" doesn't exist in the dataset yet, it makes good sense to use the value from a similar token.

Header Optimizations

Most filter authors agree that a token in the subject header is very different from a token in the message body, and that a token that appears in two different headers is unique enough to warrant keeping track of. Header tokens are usually processed differently from body tokens in order to main¬tain the origin of each token. Let's look at an example of an e-mail with a lot of useful header information.

```
From: bazz@xum2.xumx.com
To: bazz@xum2.xumx.com
Reply-To: mort239o@xum2.xumx.com
Subject: ADV: FREE Mortgage Rate Quote - Save THOUSANDS! kplxl X-
Keywords:
Save thousands by refinancing now. Apply from the privacy of your
home and receive a FREE no-obligation loan quote.
http://211.78.96.11/acct/morquote/

Rates are Down. YOU Win!
Self-Employed or Poor Credit is OK!
Get CASH out or money for Home Improvements, Debt Consolidation and
more. Interest rates are at the lowest point in years-right now!
This is the perfect time for you to get a FREE quote and find out
how much you can save!
```

In the spam shown here, several different tokens stand out. First, if my e-mail address happened to be bazz@xum2.xumx.com, I wouldn't expect to be seeing it in the From: header, but it would be very normal in the To: header. Seeing my own e-mail address in the From: header would be a clear indicator of spam, since most people don't usually send e-mail to themselves unless they've had too much to drink.

Second, the word "Save" appears in both the subject line and the mes¬sage body. I would expect to see it in the message body more frequently in legitimate mail – for example, "Save your files in the blue folder" or "Save me from this dreaded cubicle." Seeing the word "Save" in the subject header is much more suspicious, though, and it makes sense for me to have a different entry in the dataset for each of them.

The word "FREE" also shows up in both the subject line and message body but, in this case, they're both very guilty indicators of spam. The filter still benefits here because the tokens "FREE" and "Subject*FREE" now have the ability to take up two slots in my decision matrix, further con¬demning the spam. Header tokens are extremely useful for identifying both spam and legitimate mail.

Other types of header tokens are frequently found to be useful, and the set of delimiters used in the headers is usually slightly different from those used in the message body. For example, if I want to catch all of the IP addresses in the Received: headers, I would treat a period as a constituent character (part of the token) instead of a separator. If I wanted to tokenize the message-id, I'd also include the @ sign as a delimiter, as it is used to separate some pieces of the message-id.

Another advantage of including the header as part of the token is that it helps to create a virtual "whitelist" of users you trust. If I exchange a lot of corre¬spondence with bobsmith@somedomain.com, tokens like "From*bobsmith" and "From*yourcompany.com" will start to appear in the dataset, usually with very innocent values. This works equally well in identifying the hostnames of trusted mail servers in the Received: header too.

URL Optimizations

Everyday innocent-sounding words like "order" and "cgi" often appear in the body of messages I receive from legitimate mailing lists. Seeing them appear in a URL, however, is much more suspicious. URLs are the spammers' pre¬ferred means of contact. It's much easier to run a scam using a Web site as your point of contact than it is to pay for the overhead of a phone system or mail processing department. Spammers also like their privacy, since the rest of the free world hates them, and they prefer that even customers not know how to contact them or the companies they spam for. Whether it's a link to click to visit a site or the URL of an image inside the message, URLs provide a lot of useful information specific to their own kind. Even non-sensible

INFITECH

X5 NAS

empower your data network

High Performance
Rack Mount Servers and Storage Solutions

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

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

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

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

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

Powered by NetEngine

Visit Us www.infi-tech.com

or Call 1-800-560-6550

to Find Out More

numbers will frequently stand out in URLs. This makes really good data for identifying not only spam but some legitimate mailing lists that use URLs in their unsubscribe tag lines. Users who are subscribed to some mailing lists that frequently include embedded advertisements (such as Yahoo Groups) will notice some specific characteristics of the URLs used in these advertise¬ments that help the filter distinguish between advertising and real spam.

URLs are frequently tokenized differently than the rest of a message. The only delimiters usually used when tokenizing a URL are the slash, question mark, equal sign, period, and colon, although some filter authors perform the same basic type of token separation as they do in the rest of the message body. Tokenizing using URL-specific delimiters is done because the individual tokens are more frequently found based on their path in the URL, rather than on a specific context inside the URL. Regardless of how they are tokenized, URLs, when analyzed, can yield a lot of useful information. They can be categorized as places you want to go and places you don't want to go. A spam containing places you don't want to go is just as informative as a legitimate message containing places you do.

```
Url*getitrightnowwholesale     S: 00026  I: 00000  P: 0.9999
Url*thesedealzwontlast         S: 00026  I: 00000  P: 0.9999
Url*biz    S: 00008  I: 00000  P: 0.9998
Url*us     S: 00000  I: 00050  P: 0.0001
Url*java   S: 00018  I: 00000  P: 0.9999
Url*www    S: 00000  I: 00030  P: 0.0001
Url*com    S: 00000  I: 00033  P: 0.0001
Url*img    S: 00066  I: 00000  P: 0.9999
```

Ironically, legitimate URLs seem to be rare among spammers, while the wild and obnoxious names always pop up, with the exception of "java," of course, which appeared as spammy only because this user doesn't use Java (not because Java programmers were spamming). The appearance of certain naming conventions, such as the extensive use of "img," makes the task of identifying malicious URLs pretty easy. If we wanted to, we could probably determine the disposition of the message based on the URL information alone.

Ironically, URLs containing well-known Web addresses are likely to appear as innocent or hapaxes. Not a single URL token containing the following words has ever appeared in my corpus as spammy:

- Url*microsoft
- Url*whitehouse
- Url*sco
- Url*linux
- Url*quicken
- Url*intuit
- Url*_amazon
- Url*fbi

HTML Tokenization

One area that has plagued many filter authors is the decision as to what HTML to include and what other parts of the message to ignore – for example, should we ignore JavaScript? What about font tags? Most filters pay attention to all HTML tags except those on an exclusionary list, namely, a specific set of tokens that are common to all types of e-mail. This approach works quite well, but there's still room for improvement. Ignoring data is always something to be concerned about, and you shouldn't do it unless you have good reason. The justification for ignoring some HTML data is that many people normally converse only with senders who do not use HTML. This could cause any type of message with embedded HTML to be rejected as spam, which could be bad for the recipient if their boss suddenly started using an HTML-enabled mail client. The tags most filters ignore include

- td
- !doctype
- blockquote
- table
- tr
- div
- p
- body
- Short tags, with fewer than N characters of content
- Tags whose content contains no spaces

It is probably better to use an exclusionary list rather than an inclusion¬ary one. You're more likely to miss a few tags or possibly to fail to name certain tags you never thought could be used in spam (for example, the object tag has recently become popular). If this happens, at worst the tag will sit and collect dust in the dataset with some neutral value or will fill up a decision matrix slot in error. If you fail to add a tag to an inclusive list, though, you're bound to ignore an important data point and may not even realize it.

Some of the HTML tags commonly used by spammers (which a filter should definitely be looking at) include the following:

```
APPLET    BGSOUND   FRAME     IFRAME
ILAYER    IMG       INPUT     LAYER
LINK      SCRIPT    A         AREA
BASE      DIV       LINK      SPAN
OBJEC     FONT      BODY      META
```

Some filters like to mark the tokens generated from HTML tags with an "HTML" identifier, while others go so far as to mark the particular tag the text belonged to (for example, "BODY: BGCOLOR=#FFFFFF"). Regardless of which tags the filter decides to keep and which get discarded, it's very impor¬tant to handle HTML comments correctly. Spammers are using many tricks to obfuscate their text so that it's human readable, but not very machine readable. For example, the following may look like a complete mess in its machine-readable format:

```
Received: from 64.202.131.2 (h0007e9075130.ne.client2.attbi.com
[24.218.222.43])
Message-ID: <cp6-mh-rn-w$4pa2o965rl84@jn4y0hq1bcy>
From: "patsy stamm" <arthropathology71255@earthlink.net>
Reply-To: "patsy stamm" <arthropathology71255@earthlink.net>
Subject: Giving this to you
Date: Fri, 08 Aug 03 07:29:02 GMTX-Mailer: MIME-tools 5.503 (Entity
5.501)
MIME-Version: 1.0
Content-Type: multipart/alternative;
  boundary="AD0E55.76_15.C" X-Priority: 3 X-MSMail-Priority: Normal
--AD0E55.76_15.C
Content-Type: text/html;
Content-Transfer-Encoding: quoted-printable

Yes you he<!lansing>ard about th<!crossbill>ese weird
<!cottony>little
pil<!domesday>ls

that are suppo<!=anabel>sed to make you bigger and of cou<!chord>rse
you think
```

```
they're b<!soften>ogus snake potion. Well, let's look at the facts:
<strong>G<!eigenspace>RX2

has be<!waldron>en sold over 1.9 Mill<!audacity>ion times within
the last 18
months</strong>...
With awe<!tapestry>some results for hun<!wield>dreds of
thous<!locale>ands of
men all over the planet!  They all enjoy a seriously enhanced ver-
sion of their
manh<!rescind>ood and <b>why shou<!seoul>ldn't you</b>?
```

But when the user clicks the message to read it, the HTML comments won't be visible and the user will see this:

```
Yes you heard about these weird little pills that are supposed to
make you bigger and of course you think they're bogus snake potion.
Well, let's look at the facts: GRX2 has been sold over 1.9 Million
times within the last 18 months... With awesome results for hundreds
of thousands of men all over the planet! They all enjoy a seriously
enhanced version of their manhood and why shouldn't you?
```

A simple way to ensure that the message is tokenized correctly is to remove the HTML comments and reassemble the message.

Word Pairs

Using word pairs, or nGrams, has recently become very popular among authors of statistical filters and adds a lot of benefits to standard single-token filtering. Pairing words together creates more specialized tokens. For example, the word "play" could be considered a very neutral word, as it could be used to describe a lot of different things. But pairing it with the word adjacent to it will give us a token that will inevitably stick out more when it occurs – for example, "play lotto." This approach helps improve the processing of HTML components by identifying the different types of generators used to create the HTML messages. Each generator, whether it's a legitimate mail client or a spam tool, has its own unique signature, which joining tokens together can help to highlight. Tokenizers that implement these types of approaches are referred to as concept-based tokenizers, because they identify concepts in addi¬tion to content.

Sparse Binary Polynomial Hashing

Bill Yerazunis originally introduced the concept known as SBPH, or sparse binary polynomial hashing. SBPH is an approach to tokenization using word pairs and phrases. If it wasn't so effective at what it does, it would probably be a terrible idea, but Yerazunis has repeatedly astonished the spam-filtering community with the leaps in accuracy made by SBPH tokenization. Graham refers to SBPH with the same mixed feelings regarding its ingenuity and need for medication.

> *Another project I heard about . . . was Bill Yerazunis' CRM114. This is the counterexample to the design principle I just mentioned. It's a straight text classifier, but such a stunningly effective one that it manages to filter spam almost perfectly without even knowing that's what it's doing.*

SBPH tokenizes entire phrases, up to five tokens across, and allows for word skipping in between. It led the way in terms of accuracy for a long period of time, but it also created an enormous amount of data, which is one of the reasons it presently functions only in a train-on-error environment. SBPH provides the benefit of using the simplest, most colloquial tokens but giving special notice to more complex tokens as well, which are usually much stronger indicators of spam when they appear.

A few filters, such as CRM114, perform this type of word skipping, which will tokenize something like "manh+<!rescind>+ood" and also help the filter "see" the original token by performing the word skipping: "manh+ood." Since tokenization is an imperfect process, approaches like this generally provide more machine-readable tokens to deal with, without necessarily requiring much work. The more permutations of machine-readable tokens are created, however, the larger and more spread out the dataset will become, possibly affecting accuracy. The amount of data generated by SBPH generally turns a lot of filter authors off to it in favor of simple functions such as HTML comment filtering.

Internationalization

The tokenization methods discussed thus far have covered only standard character sets. The issue of foreign languages will eventually require a solution. Most spam filters simply use wide characters as placeholders, such as the letter "z" or an asterisk. This functionality allows the filter to catch just about any messages written using a wide character set. Some users, however, may expect to receive e-mail from others speaking such a language, and for them this approach won't function well at all, filtering only based on header data. The rest of the body will look (to the filter) like this:

```
ZZZZZ,

ZZ ZZZZ ZZZ ZZZZZZZ ZZZ ZZZ Z ZZZZZZ Z ZZZZZZ ZZZZ Z ZZZ ZZZZ
ZZZZZZZ ZZ ZZZZZZ ZZ ZZZ ZZZZZZZ

ZZ,
ZZZZZZZZ
```

Some filters implement i18n internationalization, which lets their filter support some additional languages. To make matters more complicated, however, some languages don't use white space, making it very difficult to identify words at all. This commonly calls for more advanced solutions such as variable-length nGrams.

Final Thoughts

We've run the gamut of approaches to tokenizing in this article. Tokenizing strives to define content by defining the construct and, more important, what the root components of content are. This is a noble quest but, as with other areas of machine learning, is a function that may eventually be better left up to the computer. As new types of neural decision-making algorithms surface, the analysis of unformatted text may become one of the next forms of AI. Until this happens, tokenizing remains one of the few heuristic components of a statistical spam filter. It should therefore be respected and kept somewhat simple, so as not to require any maintenance in the years to come. ■

About the Author

Jonathan A. Zdziarski has been fighting spam for eight years, and has spent a significant portion of the past two years working on the next generation spam filter DSPAM, with up to 99.985% accuracy. Zdziarski lectures widely on the topic of spam.

Are Your Systems Too Available?

WHO HAS ACCESS TO ALL YOUR SYSTEMS?

BY WINN SCHWARTAU

I OFTEN THINK LIKE I'm paranoid. I get paid for it.

So when I think about availability, I can conjure up an amazing array of things that can go wrong. But, instead of discussing the many security-related aspects of your storage systems availability, let's talk about how your systems may be too available. That's right – too available.

When a man wearing a telephone company hard hat and a service belt comes to your offices, where is he permitted to go? Does he have free rein of your offices including your NOC (Network Operations Center)? Can he get to the executive floor and repair phones unescorted? Does he have as much or less physical access than your employees?

Just consider that the hacker magazine 2600 has their van painted almost identically to a Nynex phone truck. Can your receptionist tell the difference?

Faced with two people, both appearing to be from the telephone company, how do you know who is legitimate and who is a hacker, or perhaps from a competing company, an investigation firm...or maybe just a bad guy out to get you? What is your company's policy on letting in the phone man, the power company, or other utility employees? Where can they go? Do they require escorts? Think about how invisible people in well-recognized uniforms are. They are innocuous, in the background like waiters at a cocktail party. We don't notice them, yet there they are and most of us don't even take a second glance.

Do you let the electrician into your NOC or computer room without supervision? Can the telephone man go to the fifth floor phone room that happens to have a network computer with a floppy disk? What damage to your networks can be done from there? Could he, with the insider access he now has, install a network sniffer or install a Trojan horse?

Maybe we make our systems and NOCs a bit too available to those invisible people who are supposed to be providing those critical support services that allow our businesses to function flawlessly.

What compounds this potential for availability problems is poor physical network design. For example, too many companies put routers and other networking components into very convenient locations like telephone or electrical rooms, or basements near shipping/storage areas. Then, receptionists or other staff point the utility man to the utility door with nary a second thought – too much availability. The electrical and telco rooms of companies in industrial parks are often located for easy access from the parking lot, and some firms – I swear it's true – leave those doors unlocked for easy access. The trouble is key networking components are often located there, too.

Some of the more security-aware companies I deal with require an escort for all outsiders, no matter how official looking they may be. The only (paranoid) problem here, though, is do your physical guards understand what the technical people are doing?

Now, ask yourself the following question: What two groups of people have virtually unlimited access to your entire facility? The CEO? The chief information officer? Accounting? Think again. Most companies give unfettered access to their cleaning staffs and private security forces.

Question two: Who are the two lowest-paid groups at your company? You might think yourself, but the right answer is the cleaning staff and physical security guards again. This has always seemed to me to be an oxymoron of security policy, behavior, and attitude. Give the greatest physical access to the lowest rungs on the corporate ladder.

Sure, the cleaning crew is bonded...but what does that really mean? It means that no one on the cleaning crew has committed a crime – or more accurately, no one has been caught. And think about the amount of availability you give them to your offices, your development and technical areas, not to mention NOCs and computer centers. Unless, of course, your security awareness is such that you have them accompanied everywhere they go by...ah...are we thinking guards? Ahem. Is that double jeopardy?

Law enforcement agencies began discovering in the late '80s and early '90s that criminal organizations were getting their people hired into "bonded" maintenance and guard services. The goal was to gain total access to a company that they wanted to victimize. Now that's what I call a bit too much availability.

Solving this problem requires awareness on the part of top management, willingness to design and enforce an effective policy, and a healthy cooperative relationship with the entire company staff. There are several simple things that com-

panies can initiate to lower the risks of too much access and availability by the wrong people. Here are a few thoughts.

> Make your staff aware of the problem of the outsider problem.
> Design and publicize an enforceable policy to your entire staff, contractors and visitors.
> Use shredders for sensitive documents. Don't forget that the cleaning crews empty wastebaskets and take the contents with them. What is your staff throwing away without thinking of the consequences?
> Passwords to company systems are never to be written down on keyboards, monitors, or under desk drawers. This must be vigilantly enforced at all times.
> Rolodexes should be put away each night. They are a key source of proprietary company information.
> Desk drawers should be locked when staff are not at their desks.
> All sensitive files on proprietary company information, customers, and employees should not be left lying around. They should be stored in secure and locked file cabinets.
> For those especially mission-critical areas of the company, a trusted (and better paid) escort should accompany them on their rounds.

The ultimate answer is trust, and some companies are turning to an approach that might be considered draconian by many people: psychological profiling. The concentration is on potential hires for key staff positions and for those to whom you will give high degrees of availability to your critical areas. What are their tendencies under ethical dilemmas? How would they behave in seemingly benign, but psychologically enlightening situations? Your human resources department can coordinate with local industrial psychologists who offer this kind of service, and then with corporate counsel to make sure that employee rights are respected. For those people who resent such profiling, maybe those are some of the very people you don't want in the first place.

Too much availability to critical network components is a real-world concern today. We need to trust our systems administrators to keep our networks going, and we have to make everything available to them to do their job. This is not an issue of trusting your staff; it's an issue of hiring people who can become trusted staff members.

This overlooked aspect of availability is being put on the table of many human resource departments by upper management, as they attempt to make sure their systems availability stays high, while also giving high degrees of availability to people they know little or nothing about. The bottom line is that making critical components of your infrastructure available to too many people, without proper controls in place, can endanger the availability of your systems when you need them most.

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

Mitigating Downtime Risk When Making SAN Changes

THE TRICK IS IN THE MODELING TOOLS AND PLANNING

BY DENIS KENNELLY

MAKING CHANGES IN a Storage Area Network (SAN) is a daily chore for many enterprise IT administrators, but so is the risk of prolonged downtime associated with configuration errors or incompatibilities in hardware or software. A popular refrain heard from industry analysts and IT consultants is that the number one cause of downtime in the data center is due to change management errors. The thinking is that undisciplined IT changes often cause problems that result in downtime. And with the size and complexity of SANs growing – especially as enterprises deploy heterogeneous environments – the downtime risks loom even larger.

Unfortunately, too many IT administrators still employ the "plug and pray" method. That is, they go blindly into implementing SAN changes in a production environment and "pray" everything will work when it goes back online. If things go awry, they must go through a tedious, reactive process of troubleshooting to single out the mistake or incompatibility, often prolonging the downtime.

Over the next 12-18 months, consolidating server and storage resources to maximize use and lower costs will be a popular reason for implementing changes in the enterprise data center. Since making even slight changes to a SAN can be potentially disastrous from an availability and SLA perspective, IT managers need to employ the right techniques and tools to prevent the worst from happening.

The Devil Is in the Details

While it's rare among most enterprise IT organizations that any significant infrastructure upgrade isn't planned in advance, often this thinking doesn't maintain when it comes to making small adjustments or seemingly minor updates. Moreover, it usually isn't possible to understand the effects of a change in a SAN until after it's implemented.

For example, making a firmware update to a group of Host Bus Adapters could prove incompatible with the connecting SAN switches. How could IT know that would happen unless it was armed with the most current interoperability information from multiple hardware and software vendors? How can IT prevent fat-fingered errors when, for example, an IT administrator inputs the wrong port assignments to an existing SAN design? Such oversights and mistakes can be costly, often impacting the performance and availability of mission-critical applications.

The rate at which these change projects fail, therefore, is much higher than if careful planning and a means to audit the proposed changes were done in advance.

Some IT organizations farm out major changes and infrastructure upgrades to outsourcers who are paid a lot of money to do the planning and implementing. These outsourcers often charge a premium just for the detailed crosschecking of device and software compatibility, sometimes using lightweight or homegrown (though not always accurate) tools designed to simulate post-change SAN performance. Or worse, inadequate tools are sometimes used that aren't specific enough for the storage domain and instead they focus on planning the time and resources needed to complete individual tasks in a change process (i.e., Microsoft Project). This often leads to incomplete planning at a technical level, creating problems during or just after the implementation. The net result is that projects take more time to complete or fail the first time around and require a second phase to repair the problems left over from the first one.

Modeling Tools Simulate and Validate in a Safe Environment

Less than a handful of SAN change management tools are available on the market that, with varying capabilities, can help IT reduce the risk of errors and downtime associated with infrastructure changes. This class of tools can have multiple uses, even to simply ensuring that existing SANs are optimally configured for

interoperability, performance, and availability. For purposes of SAN change management, these tools become powerful for not only validating that the right changes were made, but for doing frequent audits down the line to make sure the optimized environment remains unchanged.

The most compelling use for SAN change management tools lies in their ability to accurately predict outcomes. That is, the best SAN change management tools enable IT organizations to simulate how proposed changes may affect and interact with other devices and software BEFORE they're physically implemented in the SAN. Modeling functions in these tools, therefore, provide a safe environment to test different design schemes.

Good SAN change management tools also use automation to quickly upload a detailed snapshot of an entire SAN environment to save and work with as a baseline. This automation not only ensures that IT starts with an optimally configured SAN baseline, it also reduces the potential for human error and eliminates tedious manual data entry by the IT staff. This detailed topology data – which includes configuration information and device data for servers, storage devices, switches, cables, and logical access paths – can then be crosschecked against a SAN compatibility and/or configuration best-practices knowledge base to do an automated analysis of the data.

The SAN snapshot provides a baseline against which SAN changes can be modeled. By deploying SAN change management software that enables multiple scenarios to be simulated and tested, it's now possible to fully understand the holistic impact of a proposed change to the SAN before implementing a change. This critical simulation step can catch those pesky fat-fingered errors and incompatibilities. Intelligent decisions can also be made as to which change scenario carries minimal risk, while enabling the IT administrator to deliver on the required service levels for availability and performance of applications and data. Optimal change plans can then be printed out and used to support smart purchase decisions. For example, a bill of materials for the needed equipment can be generated and sent with the SAN change plan to the purchasing department along with a requisition order.

Finally, SAN change management tools accurately validate changes after they're implemented, giving IT administrators a "before" and "after" snapshot of the environment. By comparing the two snapshots against the SAN redesign plan, discrepancies can be quickly identified and corrected before the project is finalized and put into full production. This post-change verification is simply not feasible with manual entry methods and non-automated crosschecking.

About the Author

Denis Kennelly is vice-president of storage management product strategy at EMC Corp. in Hopkinton, MA. He manages the overall engineering development and product strategy direction of EMC's ControlCenter storage resource management solutions, including the development of the next-generation EMC SAN Advisor SAN change management software. Before EMC, Denis was an engineering manager at Motorola. He's also held various technical and managerial roles at Telecom Ireland Software and Digital Equipment Corporation.

ISSJ | Advertiser Index

New Backup Software Migration Approach

PROVIDING MORE AND BETTER OPTIONS FOR ENTERPRISE IT

BY KELLY HARRIMAN-POLANSKI

IT GROUPS NEED to be able to consider adopting new backup software for many good reasons. New software might have features and benefits the company needs. The current vendor's maintenance costs may be too high. The company can get cost benefits from standardizing on a single-vendor shop, combining backup and other storage software and even storage systems from the same vendor. The old vendor has acquisition or financing problems or the new vendor is willing to extend pricing breaks to increase the return on investment.

However, migrating from one backup software to another is very difficult – so difficult that many enterprises have felt that they can't do it. In fact, IT people running backup operations typically feel it's not worth their time to consider backup migration because of the expense and disruption of migrating the thousands of tape volumes that must be retained and can't be thrown away along with the backup software.

Backup migration was hard when data retention periods were only 12 months. Now, when retention periods have been expanding to seven years for most data, it's harder than ever. Most IT organizations report that they now have at least some data under infinite retention – meaning that they literally intend never to retire it.

Does that mean that enterprises are locked into using their existing backup software for at least seven years, or forever?

What the storage industry has lacked is an easy, rapid, non-disruptive backup migration service for firms looking to switch backup vendors while managing their legacy tape archives. This article will address how enterprise customers are caught in this software maintenance trap

and the problems they face in migrating their backup software given today's expanding data retention periods, ongoing software lease obligations, and heightened legal, compliance, and government regulations.

Old Fixes for Legacy Tape

Legacy tape archives didn't always pose a significant problem when companies wanted to switch backup vendors, but the situation has changed. Methods that worked well enough in the past have become impractical because of increased retention periods, software leasing models, and the sheer volume of data on tape archives.

1. ***Past Strategy #1:*** Just let it expire. Past retention periods used to average three to six months and sometimes up to 12 months. If a firm decided to change backup vendors, they simply didn't worry about the tape archives – by the time the switch was made the legacy archives would be out of retention and could be retired along with the legacy backup software.

   Present Problem: Longer retention periods. Retention periods have grown dramatically longer. In most cases retention is at least seven years – not months – at a minimum. The ever-present possibility of litigation, the changing rules governing how and when firms will be held accountable as well as the continued historical value have convinced many companies to put at least part of their data under "infinite retention." With retention periods this long (or this permanent) companies have felt under pressure to keep their original backup vendor to retain access to their legacy archives.

2. ***Past Strategy #2:*** Keep the legacy application running on a single server. To maintain access to legacy archives, companies have also taken to keeping a small footprint of the legacy backup software running – perhaps just on a single server. This has been a way to provide tape management and data restore from the legacy archive.

   Present Problem: Vendor lock-in. Instead of outright purchase, many companies lease their backup software and pay an annual maintenance fee for the privilege. If a company wants to end the lease and go to another vendor, it can. But it's not allowed to run even a single instance of its legacy application without paying full maintenance fees. Even when companies own the backup software and the right to run it without paying maintenance, it's growing less

and less acceptable for enterprises to run the software without maintenance and support since they're relying on the software to meet legal, governmental, and other compliance regulations. Annual maintenance fees for backup software are typically very expensive, averaging over 20% of the software's list price every year. So, companies are left with two poor choices: either assume a financial burden by paying a hefty ongoing fee for a service they may never use or run the unacceptable legal and compliance risk of not being able to access their archives.

3. ***Past Strategy #3:*** Migrate the data. Data migration was never fun, but when data volumes weren't as large as they are today it could be cost-effective. Companies would migrate legacy archives into the new backup software format so they could continue to manage and recover them.

   Present problem: Huge data volumes. Today, companies are faced with explosive data growth and tape is extremely awkward for large-scale restores. This makes massive data migrations hideously expensive, disruptive, and time-consuming. Don't let a prospective migration specialist tell you differently. For example, a large Californian university is undergoing a massive data migration from its old PACs system. Its new PACs vendor assured the school the migration from its legacy PACs archive would take three months. Fortunately, the university's project manager knew better and prepared for the long haul. Good thing too, since the migration has taken a year already with no end in sight.

All of these factors have combined to constrain choices around backup software, and have left enterprises and IT organizations at the mercy of their existing backup vendor. If they don't like the product or the service that they're getting, there's been little compelling the backup vendor to help them out. If the vendor wants to raise its maintenance and support prices every year – charging more and more for the privilege of keeping the same backup environment running – then there's been little to stop it.

Solving the Problem

A new spin on indexing technology is showing promise for companies that don't want to be locked in to their present backup vendor. This new approach captures the backup index, which is the metadata about the tapes made with the backup software, from across all of the legacy backup servers in an enterprise deployment. Once captured, the indexing is centralized in a single repository where it can be used to manage the legacy tape archive – without needing the legacy backup software. Using the indexing, IT organizations can know what tapes they have, when they were made, what data is on them, what the retention periods for each data set are, and where each tape is located. In short, all the information known about the tape by the legacy backup software is known by the new migrated environment.

This approach also enables using the centralized index to locate and non-natively restore from legacy archives without the old backup application. This lets the business switch backup vendors while retaining legacy archives.

This backup vendor switching approach depends on twin capabilities: the ability to extract catalogs from the legacy backup software and the ability to non-natively recover data from that archive. This enables the technology to maintain retention periods and restore data from legacy archives without copies of the legacy software and without undergoing long and expensive data migrations.

Since the technology only gathers metadata and doesn't migrate any of the data in the tape archives, indexing is quick. The initial procedure averages under a week even for large environments with multiple sites and hundreds of servers.

The catalog extraction process searches every networked server or subserver containing metadata about the old archives, pulls the metadata out of the legacy application servers, and deposits it into its own central repository database. It then responds to legacy restore requests by using the stored catalogs in its database. It finds the archive's location and can non-natively restore all data types including multiplexed data, file data, and e-mail data. It can use a series of parameters for restoration just like the old backup software. This procedure has several benefits:

> It removes replacement software licenses and maintenance fees without risking access to legacy archives.
> It replaces expensive, time-consuming, and disruptive data migrations – with this new approach, it's now quick, easy, and non-disruptive to migrate to new backup software.
> It puts all the backup metadata under one server roof instead of being scattered across the enterprise – facilitating legal discovery since searches can be done once rather than multiple times across multiple systems.
> It retains retention periods by preserving associated metadata.
> It non-natively recovers data from legacy archives.
> It meets compliance and legal regulations, and preserves historical value by retaining access to legacy archives.
> It enables companies to standardize on a single storage vendor.
> And it frees up companies to pursue cost savings and strategic goals with different backup and archiving applications.

Being locked into legacy backup isn't just a financial risk; being locked in makes it impossible for companies to improve their ROI and achieve strategic business aims by having a choice of storage management offerings. For example, legacy backup systems might not be capable of supporting storage technologies such as disk backup and snapshot protection, which leaves companies in a position of using antiquated and outdated software. And current vendors who know they have companies in their pocket feel perfectly free to saddle them with expensive licenses and hefty ongoing maintenance fees. This new approach provides rapid, non-disruptive, easy migration and gives businesses the choice of change vendors if they want while retaining access to critical archives.

About the Author

Kelly Harriman-Polanski is vice-president, product marketing in the EMC Software Group Division.
polanski_kelly@emc.com.

Endpoint Compliance, Access, or Lockdown?

MATCHING THE APPROPRIATE SECURITY POLICY TO EACH ENDPOINT DEVICE

BY MITCHELL ASHLEY

MOST ENTERPRISE ORGANIZATIONS are undertaking new projects in 2005-2006 to address the issue of endpoint security. The results of the 2005 Security IT Adoption Survey showed that 74% of respondents are budgeting, doing research on, or implementing an endpoint security solution this year. (See http://www.stillsecure.com/docs/Security_adoption_survey_Jan05.pdf). Blaster and successor malware programs exposed the Achilles heel of every network: poorly secured endpoint devices. Regulatory and compliance requirements added the business justification to allocate funds and resources to solve the endpoint security problem.

Organizations need to clearly define what endpoint security problem they are trying to solve. The answer may not be obvious at the beginning of an investigation into available endpoint security options. Rushing out to buy the latest enterprise firewall or host agent technology may not solve the right problem.

Locking Down Endpoints

Securing all endpoints, i.e., locking down or hardening the security of these devices, might seem at first like the logical solution to implementing endpoint security. One of the most significant differences when considering endpoint security approaches is that unlike network infrastructure devices (routers, switches, servers, etc.) a significant number of endpoint devices connecting to the network aren't managed, configured, or controlled by the IT or network organizations. In large enterprises, 20,000-30,000 unmanaged devices might connect through the VPN alone. Applying a single corporate standard for anti-virus updates, security patches, and personal firewalls at best

only addresses the security of corporate endpoint assets to which these polices are applied. These single policies can be difficult to enforce across the enterprise.

Most early endpoint security technologies designed to lock down endpoints were created using existing security technologies or software agents. The most common were personal firewalls, software patch delivery agents, and host intrusion detection software (HIDS) agents. These single-purpose agents have been enlarged to check for software patch levels, anti-virus, and in some cases other security checks on endpoint devices.

Any enterprise endpoint security approach must allow for the fact that multiple anti-virus, software patching, personal firewalls, and other security technologies will be used on the wide range of laptops and desktop computers connecting to the network. Rather than relying on a single personal firewall technology to lock down the endpoint, polices should be established for the security posture requirements of visitors, contractors, and home users, as well as corporately managed desktop and laptop devices. While locking down the security of endpoints may be an option for some or even most enterprise-managed assets, more is needed to address the myriad other endpoints that connect to and use the network every day.

Access Control

An important part of the endpoint security equation is controlling or limiting access for endpoint devices until the security posture of the device is known. Usually the access control method has very little to do with determining the security posture of endpoint devices. The access control technology relies on other processes, other security vendors, or even requires that the enterprise security staff build all of the testing policies from scratch. Regardless, the testing process must communicate the device's security posture status to the access control system.

Many approaches are offered for solving this problem and each has its benefits, infrastructure requirements, and limitations. A few common approaches are:

- ***Device Connection*** - Determining that new devices have connected or powered up on the network can be done in a variety of ways: through port state changes on a network switch, requests for an IP address through DHCP, or detecting network traffic from a previously unseen device. These methods can usually be implemented with little impact or change to the network infrastructure configuration.
- ***User Authentication*** - Users can supply credentials through a Web-based network registration login, network OS based login (such as the Windows domain login), VPN authentication, or an 802.1X authentication process.

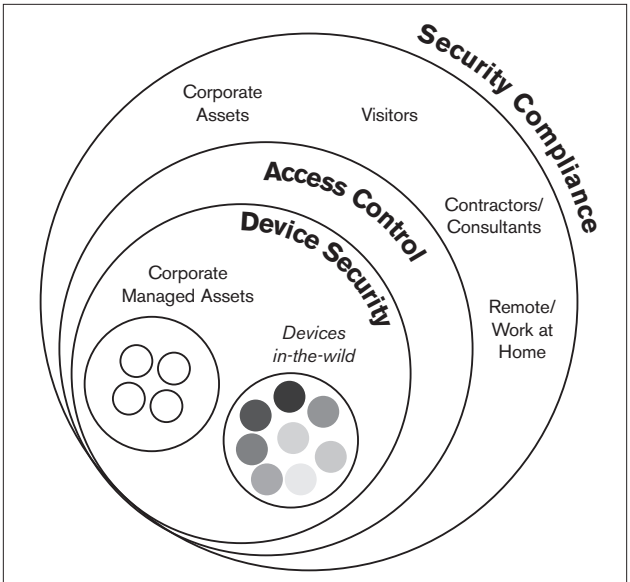

Figure 1: Compliance-based endpoint strategy

Upon successful authentication, the device's security posture is discerned. Implementing endpoint access controls through user authentication requires a greater degree of coordination and integration between infrastructure elements of the network.

> ***Local Agents*** - In some situations, as is the case with personal firewall or HIDS agents, the agent software on the endpoint device can act as the enforcement point for controlling access to the network. It relies on having agents installed on all devices.

Until the security posture of the endpoint device is determined, the device is "quarantined." This can be achieved at layer 2 with VLANs and port-level authentication, or at layer 3 through access control lists, IP address assignment and routing restrictions. Whichever method or methods are used, access control provides a mechanism for quarantining unknown devices and devices that don't meet an organization's security compliance requirements.

Security Compliance

A compliance-based strategy takes a different view of endpoint security. Rather than relying on a single limited set of technologies for securing endpoints, compliance implements a policy-based approach by matching the appropriate security policy to each endpoint device. This approach recognizes that some enterprise-managed assets can be required or even forced to use only a standard limited set of security technologies on managed endpoint devices. It also accommodates other security solutions that, while not the corporate standard, satisfy the security requirements through other security technologies on unmanaged endpoint devices.

In addition, a compliance-based approach can allow for variance in implementing specific security requirements across a broad range of managed and unmanaged endpoint devices. For example, Trend Micro's anti-virus software might be the corpo-

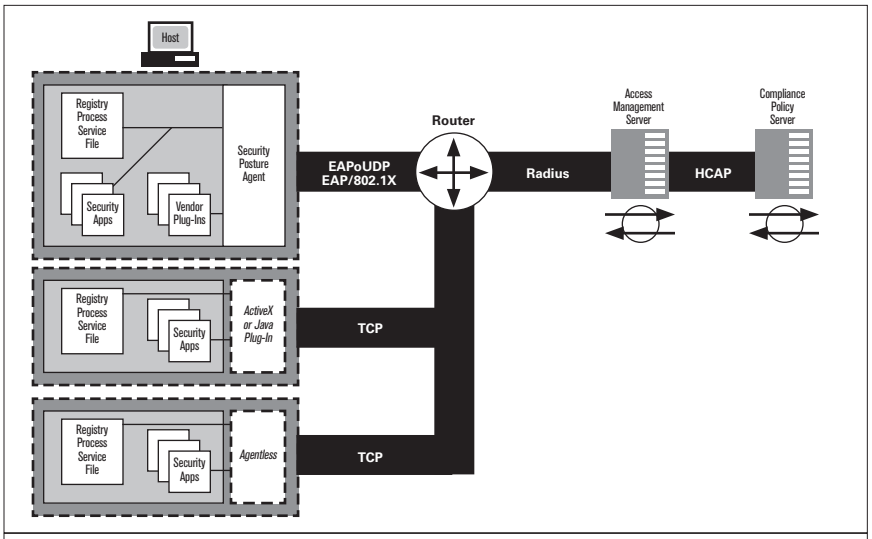

Figure 1: An example of a security compliance configuration

rate standard for all enterprise-managed devices, but a wide range of commercial and Open Source anti-virus solutions are acceptable on visitors', contractors', and employees' home computers. A personal firewall might be overkill for most corporate desktops but a number of different personal firewall products might satisfy the security requirement on road warriors' laptops. Additional endpoint security requirements or restrictions such as unauthorized peer-to-peer and messaging software, required patch management agents, and Web browser security setting requirements can be applied to each user's endpoint device as appropriate.

The compliance-based approach requires a significant amount of customizability. Merely testing for security patches and anti-virus software isn't enough. An enterprise's endpoint security requirements are typically more extensive. Policies for peer-to-peer, file sharing, instant messaging, application macros, and other security concerns are required. Other security requirements such as required security applications, patch delivery systems, Windows update settings, and Web browser security settings should all be enforceable. Connected hardware such as USB drives, flash drives, and iPods can also be of concern. Extensibility is also important. Since no vendor can anticipate every enterprise's unique requirements, the ability to create unique and custom compliance policies easily is also required.

A compliance-based endpoint security approach isn't built on reliance on a single endpoint security technology such as one vendor's personal firewall, patch management, or HIDS agent. End-user devices can have overlapping firewall, patching or HIDS software already installed, or they may not have the administrative privileges to install such a large software agent. Endpoint security solutions built around these single-purpose technologies are forced to rely on other alternatives when attempting to deal with unmanaged devices.

Many provide a so-called "agent-less" option in situations where pre-installing an agent isn't an option. In most cases the agent-less alternative means downloading and installing an ActiveX or Java browser software agent that executes as part of the end user's browser application to do an initial one-time test of the endpoint device. In most cases, agent-less options are a bit misleading because a software agent is still used, it's just delivered over the network and installed as a browser plug-in rather than a persistent agent.

There are exceptions to the faux agent-less approach though. Rather than a software browser plug-in, true agent-less endpoint security technologies use a direct network connection from an internal testing server to the endpoint device. Since no software is downloaded or installed, the typical problems with heavy software and browser plug-in agents are avoided. True agent-less solutions also offer the benefit of retesting the endpoint device during its network connect session without making the Web browser remain open and running on the endpoint device. Care should always be exercised when discerning whether any technology is truly agent-less or merely an ActiveX- or Java-based agent plug-in.

The Final Answer

It's clear by now that endpoint security isn't just about adding another security agent to every endpoint device connected to the network. There are network infrastructure considerations, access control options, and most importantly compliance policy needs that must be met.

While endpoint security might be the latest and greatest security technology craze, in the end it really isn't about technology. It's about endpoint security policy compliance for every network-connected device. A variety of technology approaches may be required to fully meet enterprise endpoint security requirements. Taking a security-compliance approach enables organizations to maximize the effectiveness and benefits that can be achieved.

"Endpoint security isn't just about adding another security agent to every endpoint device connected to the network"

About the Author

Mitchell Ashley is CTO and VP of customer experience at StillSecure where he's responsible for product strategy and the development of the StillSecure suite of network security products. Mitchell has more than 20 years of industry experience and has held leading positions in data networking, network security, and software product and services development.

mashley@stillsecure.com


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

SPEED
ESSENTIAL™ SECURITY
from SafeNet

Now you can have both speed and security.

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

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

For a free copy of the **Frost & Sullivan white paper**, "WAN Services and Encryption: Protecting Data Across Public and Private Networks," visit *www.safenet-inc.com/hse/0810*

Call 1-800-697-1316 to be SafeNet sure.
www.safenet-inc.com/hse/0810

SafeNet®
The Foundation of Information Security

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