SYS-CON

Disaster Recovery: Planning for the Worst

information STORAGE + SECURITY journal

www.ISSJournal.com

In This Issue:

VOLUME: 2 ISSUE: 6 2005

SECURING THE ENDPOINT

The Battle for the Desktop

〈26

PRESORTED
STANDARD
US POSTAGE
PAID
ST. CROIX PRESS

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

What's in Store for Q4

BY PATRICK HYNDS AND BRUCE BACKA

Okay, summer's over. Let's get back to work...

But first, let's look at what's new. Microsoft has taken WinFS, its new file system, out of the first release of its next operating system. The story is that WinFS will follow soon after the OS releases. For most of us, this is something we don't need to worry about for a couple of years, at least.

EMC has reduced prices on its Celera line of NASes. Network Appliance still dominates the high-end NAS business and its alliance with NTP Software for storage management gives it a significant advantage in the marketplace.

Speaking of NTP Software, they have released their Deep Scan technology for content filtering. Deep Scan uses sophisticated heuristics to determine the real content of a file without creating a huge impact on storage platform performance.

What's on tap for Q4? Based on the conversations we've been having, the number one issue for most people is extending the life of their existing hardware platform. While many people planned for rapid growth, compliance, which has come to mean you can't delete anything, has compounded the problem. The hardware and software facilities to handle compliance are just emerging. Most of us will be stuck with just muddling through for another two quarters at least.

Security, of course, is also on everyone's mind. (Hey, this is ***Information Storage & Security Journal***, right?) Both the privacy laws and compliance have put security at the forefront. But looking at what's going on, we are reminded of the day when we used to do network audits... If you wanted to find fault with someone's network, there were a few spots to check that were guaranteed to be trouble. Today's storage security is like that. We've yet to meet anyone who can score 100 in an audit of storage security. A lot of installations can't even get close. (Wanna try to be the first to score 100? Call us.)

Storage security is an area where there is a lot of work to be done. (Remember the recent story about a bank's off-site tapes that were lost with millions of unencrypted records?) The current state of the art is appalling. We'll all be working on this for years to come

All in all, it should be an interesting Q4. Stay tuned...

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

Wanted: Faster Backups

GUNNING BACKUP EFFICIENCY

BY ROBERT FARKALY

IN THE OLD West, the cowboy with the fastest gun survived. Some cowboys modified or bought cut-down "fast draw holsters;" modified the weight, balance, and length of their sidearms, and practiced drawing their weapons for hours – all to improve their speed.

Like them, with a few simple modifications to your backup environment, removing backup bottlenecks, and a little practice, you can successfully cut your backup and recovery time, making you a true hero in your organization.

Let's look at four areas that can be modified without much effort and help remove backup bottlenecks in the backup process.

Choose the Right Weapon: Tape, Disk and Virtualization

The cowboy would pick his sidearm carefully – and you have to choose your backup method and tools with the same scrutiny to have the fastest backup around. Traditional backup methods using tape alone can no longer be used to create fast backup. However, because tape is still the most efficient media for storing high volumes of data over long periods of time, and because of new government regulations that mandate specific data retention periods for many organizations, few companies can dispense with it entirely.

You can begin transforming to a faster backup and recovery model by combining the advantages of disk with those of tape. Disk becomes the first stage of data protection, while tape is for longer-term archiving. A common way to quickly convert from a traditional tape-based backup approach to disk-to-disk-to-tape (D2D2T) is by putting disk between the primary storage and tape and using it as a caching device, moving data from the disk to the

appropriate tape media thereafter. This first step immediately speeds backup and recovery.

Next, you can use and maintain data on a disk-based system using virtual tape, which not only makes the data easily accessible but maximizes the space used on the disk, and lets you use your physical tapes more efficiently. Make sure that the virtualization solution you use is the one best suited to your needs. Various forms include Virtual Tape (VT), Dynamic Virtual Tape (DVT), or Virtual Tape Library (VTL) on the primary device of the backup process.

Virtualization has another benefit – when file systems on disks are used for backup over long periods of time, system fragmentation occurs, causing performance to deteriorate. This is due to the constant writing, expiring, and rewriting. A virtual tape solution solves the problems caused by system fragmentation because virtual tapes don't use a file system; they're created directly from the logical volume management system.

The Right Holster: Networks and Connections

Many gunfights were lost because the gun never had a chance to fire – it was stuck in a slow holster. The right backup tools may never have a chance to perform if the network and the backup appliance's connections don't give it a chance. It's no surprise that the faster the data moves from the backup server to the backup storage device, the faster the backup will be. By making sure your network is appropriate for the backup speed required, you can save yourself a lot of headaches. iSCSI (using GigE) and existing Fibre Channel are ideal connections between backup servers and their storage. Simply adding an inexpensive network card to the storage device combined with a little system configuration will speed things quite a bit – especially on most high-speed corporate networks.

Another way to remove backup bottlenecks on the network or LAN is to use a different network for storage than you do for e-mail, file, print, and other business-critical software. We find that many companies still use the same network for backups and other business applications, and backing up over the corporate network impacts everyone's performance – including users' e-mail and Internet access. Unfortunately, the day of the nine-to-five job is gone, and most of people tend to work around the clock. If the network they use for business processes is the same network the IT department uses for backup, network congestion will most certainly occur. This causes nightly backups to take hours longer than they should – often causing them to fail and creating cranky employees come morning because they waited up all night for a multimedia presentation to download.

Installing a separate network for backup is relatively inexpensive – and worth the cost. You can easily set up a dedicated gigabit backup network and install a new network card in your application servers, buying yourself up to 10 times as much performance. Not only do you minimize complaints from night-owl executives, but you improve efficiency and save time on the backup process.

Modify Your Weapon: Fine-Tune the Software

Not only did gunfighters modify their holsters, they often modified their sidearms, changing the weight, the length and the balance of their guns. If the gun wasn't right, it wasn't used – knowing that the right tool for the job is the only tool to use. IT managers should use the right backup software designed for the system, environment, and process. In fact, it isn't enough to just select the right software, but you must follow the manufacturers' recommendations on tuning it for your specific application. We've found that working with a value-added reseller (VAR) can help you fine-tune the software. Most out-of-the-box default settings aren't recommended for the kind of speed you need. Not all software is equal – you'll find that tuning the buffers, caches, and block sizes can dramatically affect a system's speed and performance.

If one or more of your application servers takes too long to backup, you may want to look at installing a media server version of your backup software (backup media servers write directly to tape or disk) on those slow-to-backup application servers. By installing media server backup software on all critical servers, you can leverage the power of the software in multiple locations, making the process faster. With this kind of software, backup data is most frequently pulled from application servers over the LAN by a backup sever and written to a backup target disk or tape.

The backup data can be written directly from application servers to its disk or tape target. This way the software doesn't have to pull data across the LAN and saves time and money – all for the nominal cost of an additional software license.

Remember, if software sits on both an application server and backup server, you'll have to work with your VAR to fine-tune both places.

Using software to remove overhead is a time-saver as well. A large number of small files overwhelm a system if backed up file-by-file. Smaller packets of data tend to have just as much overhead and metadata as larger BLOBs (binary large objects) of data. Backing up small application files, graphics, and e-mail messages that have the same overhead as large files is a waste of valuable backup time. Small files moving across a network are also likely to interfere with larger BLOBs of data. Using image backup software to consolidate small files into BLOBs will make the data transfer more efficient, improving backup performance and reducing backup time.

A Trusty Friend: Servers and Systems

Horsepower isn't just for cowboys. If you're using a Pentium II or other older system, you should consider buying a newer system for backup. Backup is an intensive process, and older systems just can't pull the data from application servers fast enough to be efficient. By spending less than a few thousand dollars, you can significantly speed up your backup with a more up-to-date, efficient and faster server.

Using fast disk storage for application servers can help deliver the performance you'd like to see when backing up. External storage, including RAID, can also be a backup bottleneck, but to improve backup performance, use fast (usually external) disk storage on application servers. This will improve application response as well. Remember, backup speed can only be as fast as the speed at which the application server's disk delivers the data.

Besides fast external storage, you can improve backup speed by reconfiguring your backup jobs to run in parallel rather than serially. Many IT managers sequentially back up their servers (backup server one, followed by backup server two, then three and four, etc.). Sequential backup quickly fills up an overnight or weekend backup window – but using your system to absorb all the data stream simultaneously is a simple solution.

By using a multi-stream-capable backup-to-disk target, you can back up all of the servers at the same time by simply modifying your backup scripts. By doing this, you can start all the backups at the same time – reducing your overall backup time. With most sequential backups that overlap, each job will take a little extra time when processing multiple data streams at the same time, but the total time spent backing up will be cut.

Many general RAID disk systems aren't multi-stream-capable. However, specialized yet inexpensive backup and recovery appliances can process multiple backup data streams simultaneously. These specialized appliances, such as the Overland REO SERIES, can back up often problematic application servers that are newly equipped with media server software licenses.

By taking advantage of these best practices for backups, you'll quickly reap the benefits, allowing you to focus resources on new storage initiatives. Not only will you have the fastest backup around, you'll have time to ride off into the sunset.

"If you're using a Pentium II or other older system, **you should consider buying a newer system for backup"**

About the Author

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

Backup for Continuance of Operations

THE BACKUP DILEMMA

BY SCOTT GOLIGHTLY

EVERY PERSON WHO has ever been responsible for backing up data has had to ask themselves the same basic questions. They need to know what data has to be backed up, how frequently it changes, where to store the backups, and how quickly the data will have to be restored in case of disaster. The answers to these questions in a large way determine the media used to back up the data and the ultimate storage location for the backup. It seems that every day we read about a hurricane, fire, flood, or other disaster. Couple the natural disasters with the need for 24x7 availability and increasing government regulation and it's easy to understand why the disaster recovery plans of an organization are coming under scrutiny. For years the general pattern was to back up data to tape and store those tapes onsite in a vault or offsite in a secure location. But several recent developments have started people looking at other alternatives.

Many organizations still prefer tape for long-term data storage. It's especially attractive for data that's unlikely to change and lack of instant access isn't a problem. On the other hand, with the increasing capacity of disk drives and the reduced cost per megabyte for storage backing up to disk makes sense for data that has to be available relatively quickly. This article will look at mirroring some of your most important data to mitigate the effects of a disaster.

Continuance of Operations (COOP)

When disaster strikes the priority is always to get the company "up and running" as quickly as possible. Hopefully you've done your homework beforehand and know what data has to be available for work to continue. The data to be

restored can be broken into three tiers: tier-one data is the most important, tier-two data supports tier-one data and makes your operations run smoother, tier-three data is the data that can be restored later. (See Figure 1.) While each business has to decide what data must be available for work to continue, tier-one data will almost always include customer-facing data. This might be data used in a product catalog, user login information, or data used by your order-receiving systems. Tier-two data has to be restored quickly but won't stop your business from making money. Tier-two data might include an inventory control system, Web site personalization data, customer order histories, or accounting information. Finally, tier-three data is data that can be restored later. Tier-three data might include data warehouses, historical data, or reporting data.

For years databases have used mirroring and clustering schemas to reduce or eliminate the time it takes to get data from a more durable media like tape. By applying some of the same concepts you can ensure that operations stay up and minimize the time needed to restore from backup by mirroring tier-one data. In the event of a disaster you have to apply any completed updates to the mirror location and then redirect applications to retrieve their data from the backup location to continue working immediately.

What to Back Up

People say that "disk space is cheap," but it's not free. As you look at the data that you want to mirror you have to prioritize. You want to ensure that the data that's mirrored is the data that's essential for your business to keep on making money. It should be the first data that's restored in your disaster recovery plan. Other supporting data can be restored later. This is your tier-one data and, as mentioned, will probably include data that customers need to order products as well as data that's essential to the company's core competency.

To ensure that the proper data is mirrored and ready for a disaster, a company's IT and business people have to identify what processes are most critical. They should then identify the data that's used in those processes. The applications should be looked at carefully to see if they could work with only partial data. Most applications aren't designed to work with only partial data, but given the increasing number of apps that support occasional network connections and services that retrieve data you might be able to begin using an application before all of the data stores that it relies on are fully restored. Once you know all of the business factors and the data dependencies of your most critical applications you can prioritize the data and come up with a plan for switching over to a mirrored data source or restoring from backup to kick start operations quickly. The plan should

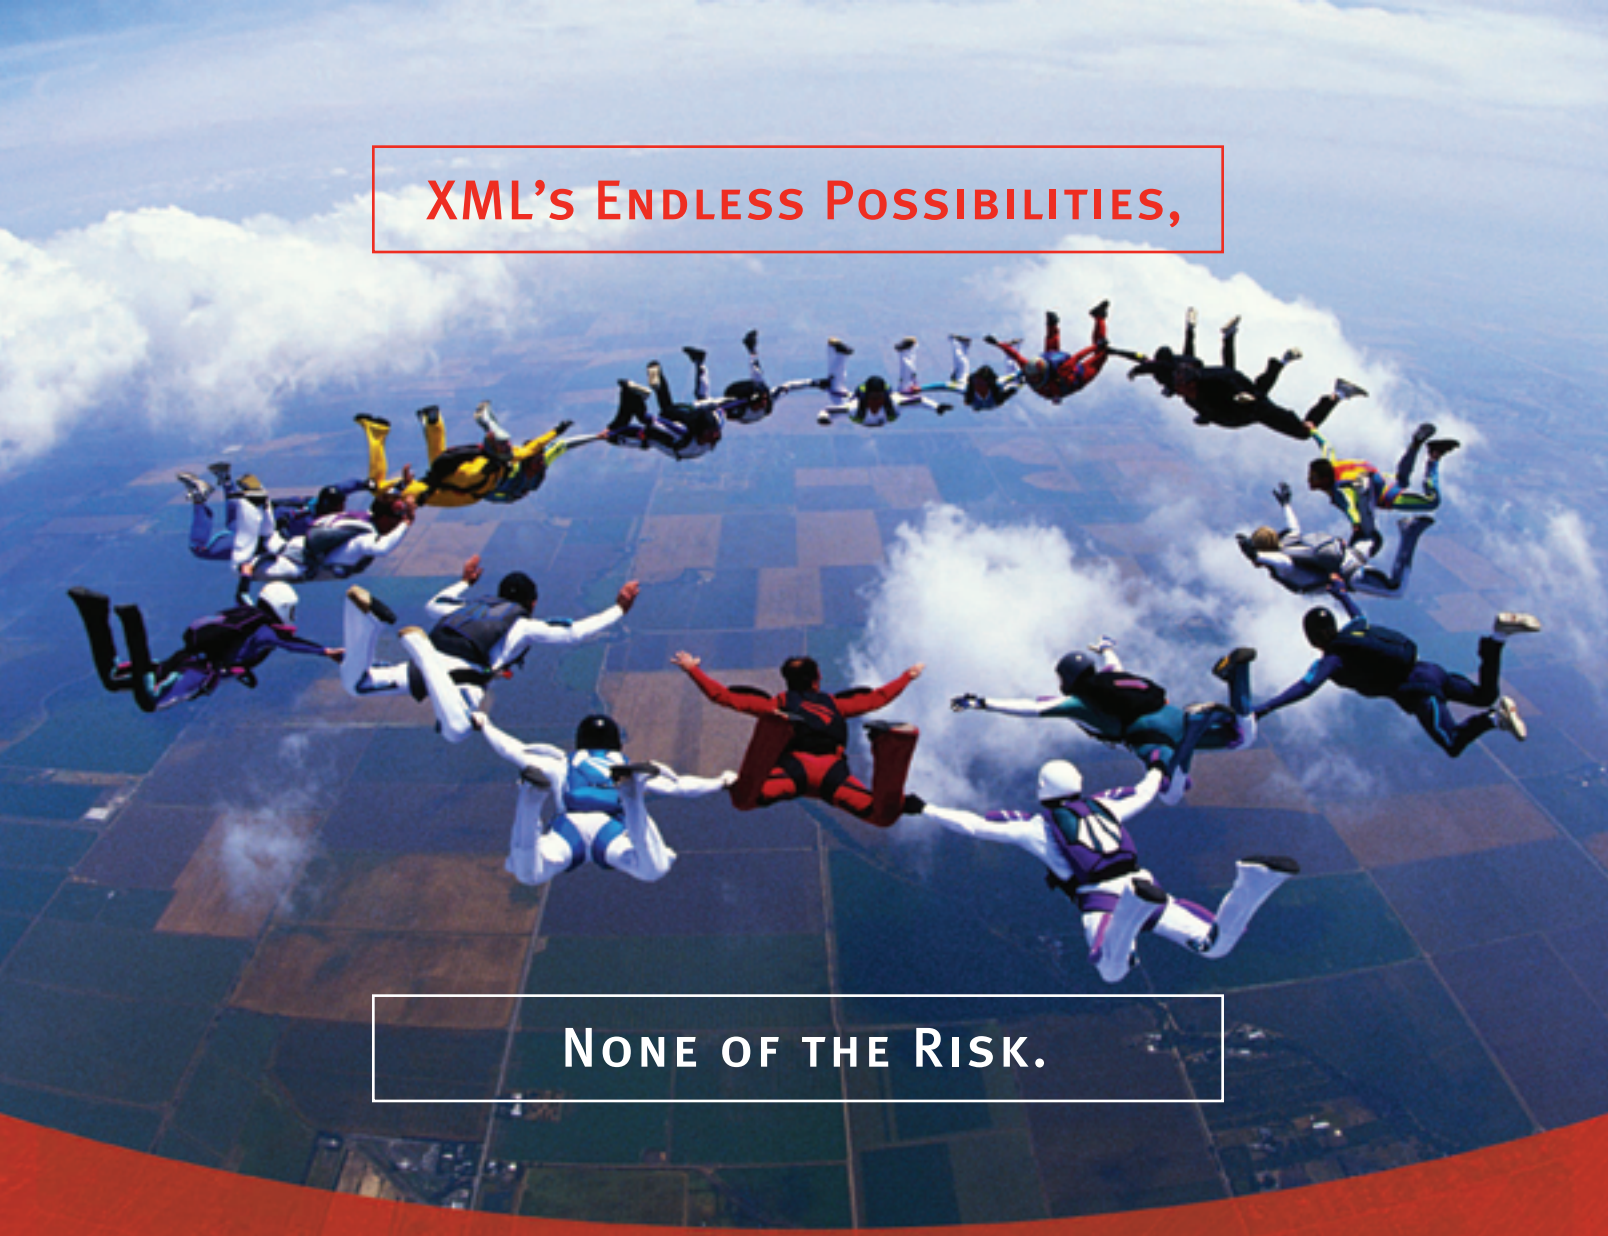

Forum XWall™ Web Services Firewall - Reinventing Security

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

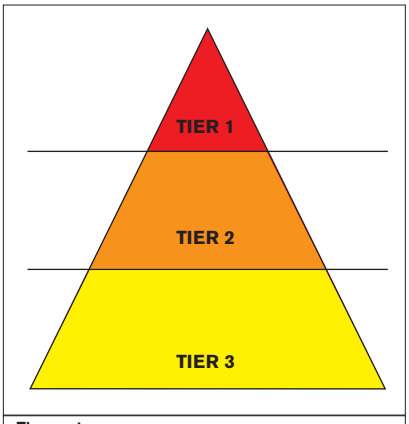

Figure 1

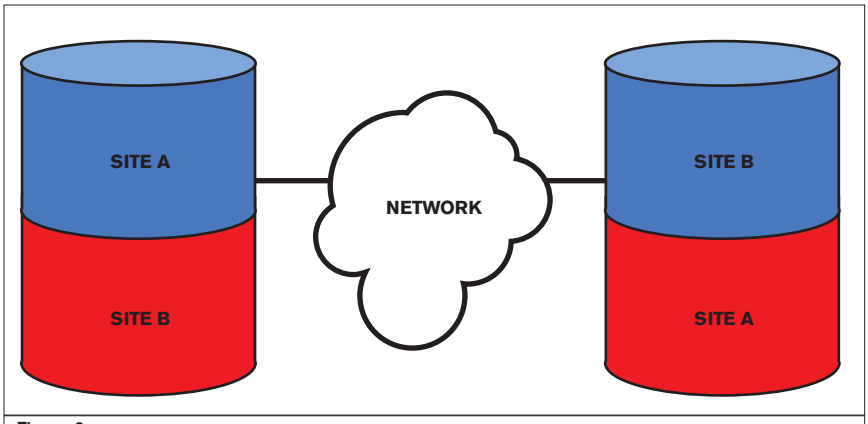

Figure 2

be reviewed periodically and any time there's a major change in the business, such as launching a new product line, a merger, or acquisition, or when the business experiences significant growth. You can also use the insight that you've gained into the interdependency of your data to drive development so new applications work well in your recovery plan.

Where to Back Up

Once you've decided what should be mirrored to another site you'll need to decide where that site should be. Obviously the primary and mirror sites should be far enough away from each other that they wouldn't be affected by a single disaster. By separating the sites geographically, you'll also ensure that they're on different power grids. And you need to consider staffing requirements. You'll want to be sure that someone will be there to make sure the mirror site is up and running when you need it. The sites can't logically linked either. The networks and other infrastructure should be able to operate independent of each other.

One other consideration is that the infrastructure be able to handle the expected volume of data. Since very few businesses will have the luxury of maintaining a complete data center strictly for disaster recovery it's likely that your backup site will be the primary site for some data. In the event of disaster at one of the sites, the other one will be the primary site for both sets of data. And as the data center gets used for more and more activities, you'll need to review it periodically and see if there's enough network capacity, storage space, and sufficient cooling to handle both the normal workload and any disaster-created spike in activity at the other site.

It might be easy for large organizations with multiple data centers to pick a pair of data centers that mirror each other's data. Smaller organizations geographically limited to one location or that have only one data center have to be creative in finding a partner to provide an off-site location for their data. By partnering with a company in another location you could store your backup data at their place and they could store their data with you. This would give each company a backup location without the cost of building a new data center. Figure 2 shows a typical arrangement with two sites, each mirroring some of the other's data.

Securing Your Backup Data

Regardless of whether you use a company-owned data center or a trusted partner, you should think about the security of your data. Having that much critical data in one place could be a tempting target for hackers or disgruntled employees. It's also easy to forget about monitoring or securing mirrored data. One way to safeguard the data is to encrypt it on a disk. You'll have to have a PKI infrastructure in place and add the recovery of encryption keys to your disaster plan. Another consideration when deciding whether to encrypt the data on disk is your applications. Most applications can't read encrypted data from disk and use it. You'll need a way to encrypt and decrypt data at the disk or network level so the application doesn't know the data was encrypted.

Other Considerations

As you look at what data you'd like to mirror to a remote location you should consider other technologies that would allow it to be available in the event of a disaster. For some of your data, caching it for a long time might let you keep it available in the event of a disaster and improve the performance of your application. The data that's a good candidate for caching changes infrequently and aren't specific to the person or application accessing it. A product catalog or list of services might be apt for caching for relatively long periods of time.

Throughout this article I've assumed that your tier-one data won't change while you're reading it. If the data could change you'll have to have a plan in place to ensure the changes are reflected in the primary site when you switch back to it.

Conclusion

With some planning you can set up a site for continuance of operations that will store the data most critical for you to function. In a catastrophe you can fall back on a mirrored data and ensure that your business appears to be functioning to the outside world while you buy yourself some time to restore data from your regular backup.

About the Author

Scott Golightly is Microsoft regional director and senior principal consultant at Keane, Inc., in Salt Lake City. He has over 12 years of experience helping his clients design and build systems that meet their business needs. When Scott isn't working he enjoys fishing, camping, hiking, and spending time with his family.

scott_j_golightly@keane.com

There's an easier way to protect your data.

Every 12 seconds, a PC is lost or stolen—most with confidential or sensitive information. In today's stringent environments with laws and compliance issues all around us, protecting your data is a priority that no one can afford to live without.

Fortunately, keeping data safely locked away is easy with SafeGuard® from Utimaco. Unauthorized users can't access, read or decipher protected data, and authorized users won't be inconvenienced or inhibited in any way.

SafeGuard keeps them out with industry-leading authentication, including pre-boot authentication for eTokens, SmartCard and Biometric identification systems. That's just the first line of defense. The entire hard disk's contents are protected via any of 10 worldwide industry standard encryption algorithms. And because SafeGuard encryption works in the background, users will never know it's working.

Utimaco also has security solutions to protect PDAs, smartphones, files and folders, and emails, as well as LANs, servers and storage, email gateways, embedded systems, digital signatures and stamping.

Protect your crucial data with Utimaco. Risk management has never been so easy.

For more information on how we can help you protect your mobile data visit our website at **www.utimaco.us** or call us at 1-877-UTIMACO.

Try a free demo of our latest SGE4.20 with FIPS mode and Computrace compatibility.

utimaco®
safeware
Security made simple.

Designing and Implementing a Security Architecture

ENSURE YOUR ASSETS ARE AVAILABLE, RELIABLE, AND SAFE

BY RICHARD WILLIAMS

INFORMATION SECURITY IS a top priority for many companies. Protecting information from external threats such as hackers, viruses, and spam, as well as governmental regulation requirements (SOX, HIPAA, NISPOM, etc.), are driving IT purchases beyond ROI as C-level executives seek to assure shareholders (and themselves) that assets are secure within the company complex. Viewed as today's growth market, many software/hardware/service companies are creating offerings to mitigate perceived risk or actual liability.

The security environment within some organizations may be somewhat lax – "safe" behind the routers, IDS, and firewalls. In this article, I'll discuss how to create a security architecture, including analysis, planning and prioritizing security needs, and I'll examine the following topics:

> Understanding security architecture
> Balancing threats, costs, and the value of secured assets
> Creating an architecture that fits the business framework
> A layered examination of security, including network access, application access, external access, and physical access

In addition, references are provided at the end of the article with links to useful information.

Understanding Security

Security architecture differs from other kinds of security in that it addresses requirements from a high-level perspective as opposed to a tactical perspective. When possible, you should understand your company's security requirements before specific security issues are implemented. It's as important to know your own assets, where they are deployed, and what they're worth to your company, as it is to know what threats they are facing.

Security architecture can become very complex. By looking at security from multiple perspectives, including external access and physical security, network security, application and computer-specific security, you'll be looking from the outside in as well as from the inside out. These perspectives must also be balanced against other business requirements, financial and otherwise. Whatever model or security architecture you use, you are trying to ensure that your assets are available, reliable, and safe. Consider Figure 1.

The confidentiality perspective prevents your competition from siphoning off the cream of your company's products. The integrity perspective protects your information from unauthorized modification with verifiable, auditable access records. The availability perspective ensures that information within your business is accessible at all times. Your security architecture should focus on delivering these three attributes. Securing your information while keeping the click-and-mortar business open and vibrant is a very challenging task.

Dollars and Sense

When planning your security architecture, you are governed by overriding factors including time and money. In some cases, spending one makes more sense than the other. For example, a small company pushing their first product into the marketplace may require a security architecture overridden by cost above all (if the product doesn't make it to market, having a safe infrastructure doesn't matter). They may have to phase in security measures over time. At the other end, non-compliance with a regulatory security standard could cost a company a large account, or even threaten its ability to remain open for business.

It's also important to understand that the strategic view of your enterprise security architecture is a view of where you want to be. Few companies can afford to start from scratch with regard to implementing security. For example, your company may currently address physical access with Intrusion Detection Systems, gateways, and firewalls. These are integral elements of a good architecture, but alone they may not adequately address the risk to your company.

To create the appropriate architecture for your business, you need to strike a balance between the value of assets being protected and the cost of the protection. As a general guideline, protect the highest valued assets most stringently. This may be your source code and the servers it resides in, or perhaps the marketing info including the initial public offering data. Tape backup into an offsite location may provide adequate protection for some businesses (based on the cost/value analysis), while others may require biometric access to the clean rooms where prototyping is occurring. Secure higher-priority assets first, and keep moving forward with planned steps to reach a secure destination.

Create a Security Architecture That Fits the Business Framework

As we have seen, there are multiple perspectives in a security architecture. Many models exist that may match one, some, or all of the important perspectives. There are many framework examples, including the Lattice, the Federal Enterprise Architecture Framework, the Clark-Wilson or Biba models, and many other reference models

(see the hyperlink section for reference links to these and other frameworks). In each case, the common goal is to create a balance between the business needs and the information systems that support them. Understanding what is important in relation to other things in your business helps you value both the assets and the corresponding protection you will afford them.

For example, Figure 3 is an X-Y graph that shows assets increasing in value (up the vertical axis), facing increasing risk over time (on the horizontal axis extending to the right).

This simplistic representation shows the most highly valued assets facing the least exposure to risk over time, descending in value to assets that can withstand increased exposure to risk over time. Whatever method you use, the value of assets in your enterprise needs to be determined. Revisit these models when you acquire additional assets so that their value is properly established and defended. In this way, there is an ongoing evaluation of what assets are present and their security needs within the business framework.

Network Security Architecture – It's Not Just Firewalls Anymore

As your customer community grows more sophisticated and begins to expect more protection from your products or services, the potential for accidental or intentional misuse or attack within your company grows as well. A majority of data loss in companies today occurs via credentialed accounts. Similarly, reliable and correct delivery of information on your LAN or WAN is no longer guaranteed via TCP/IP, with address spoofing and snooping available to anyone on your network, unless network security is active from the inside-out as well. Evaluate this short list of network security mechanisms as potential additions to your security plan:

> ***Data integrity checks and data encryption:*** Stored before and compared after critical data transmission, integrity checks can include encrypted totals, which can identify data transmission errors. Network transmissions using encrypted totals need to use the same encryption at each end of the transmission, either via the network or via the application after delivery. Using different encryption methods for different types of transmissions or different data streams makes data transmission even more secure. SSH, SSL, and Secure Telnet are examples of network applications that encrypt their data in transmission. If you have multiple locations (i.e., R&D in one office, finance in another), data transmission between the offices should be encrypted and the contents verifiable.

> ***Transmission logging:*** Storing an audit trail for transmissions or applications that transmit data can include the transmission date, time, transmission type, source, and destination.

> ***Transmission loss:*** In some cases, data loss on an otherwise reliable network can indicate port-scanning activity (i.e., someone viewing transmission samples looking for vulnerabilities). With 65,535 TCP ports on a system (using TCP/IP as the lion's share of network traffic), active data transmission to well-known ports such as http (port 80) or telnet (port 23) are the tip of the iceberg, but are often the first point of attack. To defend against this activity, keep port-scanning tools off of the network with a published mandate in security policies known to all employees, backed up by a periodic review of hardware and software inventory on computers. Keeping unused ports closed and current network patches on systems also enhances network security.

> ***Change control review:*** While many system and network administrators view change control as an impediment, reviewing network devices or software before they are introduced allows a larger perspective, including the security and business framework. The extra time spent here is inexpensive insurance over the system's life cycle. This is particularly true when your company is starting up operations. Documenting and following best practices produces repeatable, reliable results.

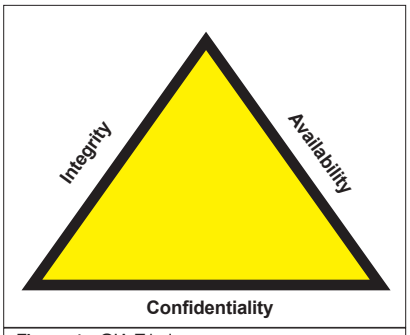

Figure 1: CIA Triad

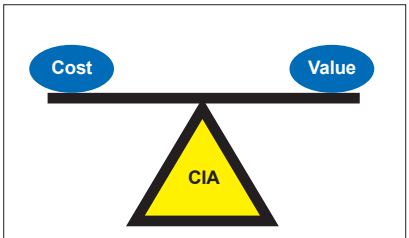

Figure 2: Balancing CIA

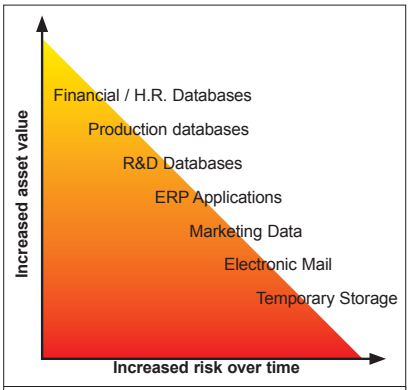

Figure 3: Assets vs risk

Application Security

Within an enterprise there are many applications used for data input or reporting, communications, database access and management, and Web services. Your actual matrix may be simple, or very complex, but each application should comply with your basic security architecture and business framework. It is important to provide the highest level of application security without impairing the business capability.

The Five Ws

Who, what, when, where, and why? These questions should have clear, documented, auditable answers before the installation of any application's software. Who is the application's primary user community? What is their business function? When do they require access to the application? Where is the application installed, and from where is it accessed? Why is the application important? How does it meet business needs? In addition, the answers should be periodically reviewed within the security architecture to make sure they remain relevant and adequately addressed throughout the life cycle of the application.

As each question is answered, security architecture issues will fall out. For example, a communications application is used by sales staff via remote access from anywhere in the world at any time. The access allows sales to enter orders, query inventory and/order status, query ERM application modules, and modify personal account information within specific sales parameters. Again, a visualization tool aides in this evaluation – consider Table 1.

Users	System	Who	What	When	Where	Why
Bill	Sellit01	smusa01	Creeate, query inventory, query ERM, modify personal account information			
Cindy	prweb	mkt01	Web design for Sellit01	7x24	office	Marketing, P.R
John	Sellit01	sales01	Create orders, query inventory/order status	7x24	VPN02, office	Sales rep region 1
Sandy	Sellit01	sales02	Create orders, query inventory/order status	7x24	VPN03, office	Sales rep region 2
Steve	Sellit01	sales03	Create orders, query inventory/order status	7x24	VPN03, office	Sales rep region 3

Table 1

You can see many communications and application security issues emerge from this simple case description. These issues may include remote access via VPN or IPSEC tunneling, http or https access, middleware application security, boundary testing, address checking, and security testing to ensure that credentialing to the queried applications is appropriate and at the level required to do business (but no higher).

Match each of the assets valued in your enterprise security plan against this simple set of questions and be prepared to address security concerns that emerge. Keep in mind that the goal is to enable business processing while safeguarding assets at the highest level possible. Often this is accomplished by providing the lowest level of access required for a specific business task as well as testing the application for security. You will decide if the protection is worth the risk of leaving your operations open, or at what level you can afford to provide protection.

External Access

Your company security architecture should also allow external access at the least privilege-required level. In the previous example, sales staff access may happen from anywhere in the world. Your security architecture should allow this access with a secure application, providing the highest level of security for accessing only the application(s) they require for their business function.

An example of this might be a company providing remote access to their development staff for a variety of services, including at-home development at all hours for principal programmers, file upload/download capability to outsourced marketing/public relations firms, or potential customers accessing the corporate Web site. In these scenarios, the "who, what, when, where, and why" may resolve to thousands of annual visitors accessing applications to get product, to pay for services, or to ask a general question. Access could occur from anywhere in the world, based on the specific application access.

The corresponding network security requirements to fit the business framework might include http and https access passed from public networks to the private corporate LAN or WAN, thus allowing middleware applications to query customer record databases and payment processing applications. These systems could be in separate data centers, requiring data transmission on the corporate network to pass from the internal Web/middleware systems to the database systems, to the financial systems, and return the requested information to the viewer while completing internal processing – all within stringent requirements for data security.

In a complex transaction model, having a security architecture and business framework provides guidelines and limits, helping to ensure that business is done efficiently while maintaining the highest level of security possible. It's no longer enough to determine that the data is secure in transmission. Denial of service attacks on the corporate Web server can be catastrophic when each second of real time represents hundreds or thousands of transactions. To keep this from happening, to detect it, or to analyze it, companies need to actively protect the business from these type of actions.

Physical Access

With today's phones, PDAs, handheld computers, and wireless laptops, the limits of physical access security have never faced stronger challenges, while the requirements continue to skyrocket. You should evaluate the kind of physical access required with the potential threat. For example, are your company's assets located in an area subject to natural or environmental threats, such as earthquakes, hurricanes, tornadoes or floods? Are your global resources in areas subject to terrorism or civil unrest? What about the likelihood of corporate data theft or destruction by disgruntled employees or ex-employees?

It is likely that your organization faces some of these risks. Does your staff walk away from systems with active logins, leave the server room door open, or leave keys in the server racks in machine rooms? The scope, detail, and expense of your physical access security plan should also be compared to the value of assets and secured to the highest degree possible without adversely affecting normal business functions. Installing screen locks that become active after 15 seconds of idle time may cause considerable productivity loss, as well as increase employee irritation. Requiring all documents to be shredded before disposal may only be required where vital data can be compromised.

The Sum of the Parts

Ongoing scrutiny, review, and modification of each of the areas presented provide a basic groundwork for security architecture. The key word is "ongoing" – security architecture is not a static process. You can't "set it and forget it." Implementing the maximum level of security required by each asset in your business is a task measured in man-years, not man-hours. But when compared to the value of your company's information systems, isn't it worth it?

Reference Section

> http://www.cisecurity.org/
> http://www.sans.org/score/
> http://www.itsecurity.com/dictionary/biba.htm
> http://e-government.cabinetoffice.gov.uk/Resources/FrameworksAndPolicy/fs/en
> http://www.attackprevention.com/ap/library/securitymodels.htm
> http://www.crime-research.org/news/07.06.2004/320/
> http://www.itsecurity.com/dictionary/cw.htm

About the Author

Richard Williams is a senior product marketing specialist for Symark Software in Agoura Hills, California, with over 20 years in systems administration, architecture, and design.

SNIA

STORAGE NETWORKING WORLD

COMPUTERWORLD

Learn How to Achieve Storage Networking Success

October 24-27, 2005 • JW Marriott Grande Lakes Resort • Orlando, Florida

Featured Speakers Include:

JOSEPH AMADO
Vice President, Information Services
Philip Morris, USA

YURI AGUIAR
Senior Partner, Chief Technology Officer
Ogilvy & Mather Worldwide

KEN BLACK
Global Storage Achitect
Yahoo!

KARLTON JOHNSON
Lieutenant Colonel, U.S. Air Force
US Army War College, Class AY06

JOSEPH TUCCI
President & CEO
EMC Corporation

The Leading Conference for:

- IT Management
- Storage Architects
- IT Infrastructure Professionals
- Business Continuity Planning Experts
- Data Management Specialists
- Network Professionals

To register or for more information, visit **www.snwusa.com**

Attendees at Storage Networking World Fall 2005 will see solutions from companies including:

as of 8/30/05

PLATINUM SPONSORS

adic Intelligent Storage™ · Agilent Technologies · BROCADE · Cisco Systems · CommVault Unified Data Management · Computer Associates · DILIGENT · dot HILL · EMC² where information lives · EMULEX · hp invent · HITACHI DATA SYSTEMS · IBM · intel · LSI LOGIC STORAGE ADAPTERS · McDATA · Microsoft · QLOGIC · Quantum · Seagate · Sun microsystems · symantec. · Xiotech

GOLD SPONSORS

adaptec · AppIQ · BakBone SOFTWARE · BLUE ARC · ciena · engenio · FCIA · Finisar · FUJITSU · Infortrend · ISILON SYSTEMS · ONARO · Overland STORAGE · pillar DATA SYSTEMS · Plasmon · sgi · SNIA · SPECTRA LOGIC · syncsort · verizon · WilTel · xyratex

CONTRIBUTING SPONSORS

3PAR · acopia · ARCHIVAS digital archiving · CENTREPATH · compellent · Data Domain · DECRU · EQUALLOGIC · Evaluator Group · EXANET · FalconStor · Hifn · intransa · FORSYTHE · KASHYA · LeftHand networks · Maxtor · neartek · NEC · ONStor · PMC-SIERRA · PROMISE TECHNOLOGY, INC. · Rainfinity · revivio · RIVERBED · sbe · Silverback · SOFTEK · TACIT NETWORKS · TMS · Topio · wasabi

MEDIA SPONSORS

Bio-IT World · Business Solutions · COMPUTERWORLD · DiGstar · Extension MEDIA · IDC · INFOSTOR · InfoWorld · Integrated Solutions · jobstor.com · NETWORKWORLD · STORAGE NETWORKING WORLD ONLINE · Sys Admin

PRE-CONFERENCE GOLF OUTING SPONSOR

Quantum

PLATINUM PARTNER PAVILLION SPONSOR

Microsoft

For sponsorship opportunities, call Ann Harris at 508-820-8667

Co-Owned and Endorsed by SNIA

Co-Owned and Produced by COMPUTERWORLD

Building a Secure Corporate Environment

VIABLE BUSINESS UNITS CAN MAKE POSITIVE CONTRIBUTIONS TO THE BUSINESS

BY STEPHEN W. FOSTER

This article is an excerpt from Larstan's The Black Book on Corporate Security. *This new book is available in bookstores and the first chapter is available for free at www.theblackbooks.com. Printed with permission from Larstan Publishing, Inc. All rights reserved. Copyright 2005.*

YOUR COMPANY NEEDS a secure data infrastructure, but how, exactly, do you set one up from scratch? Here, a former FBI agent who now serves as an information security officer reveals the best methods for creating a system that takes control of your information.

I'm a battle-hardened veteran of DMZ skirmishes. No, I'm not talking about the "demilitarized zone" imposed between North and South Korea following the Korean War in the early 1950s. Among information security officers such as myself, a DMZ is the euphemism for a computer host or small network inserted as a neutral buffer that separates a company's private network and the outside public network. It stops outside users from obtaining direct access to a server that contains company data. As you attempt to tailor a secure network to a company's overall business strategy, crucial and sometimes controversial issues such as DMZs emerge and they must be dealt with in a forthright manner. That's why building a secure corporate environment starts with communication.

Building a new information security team is no easy task and will be fraught with many obstacles. The building effort begins during the CISO's interview process, which will provide him or her with a window into senior management's philosophy on information security. The support they provide is essential to your success.

The first order of business in building any new program is the discovery phase. The CISO must get out of his office and meet other business managers face to face. Reaching out and developing a personal relationship is vital to your success. Today, too many managers rely exclusively on conference calls and e-mail. The information security team should also educate key managers within the company as to how security can partner with them to help enable their business solutions. CISOs should continually demonstrate to the business that the information security team is an integral part of the business process.

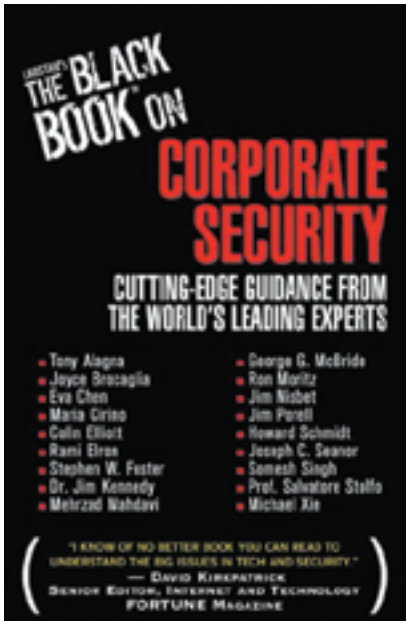

For example: Business unit XYZ requests that a risk assessment be conducted for a new DMZ they want to build. This DMZ will be used for outsourcing services to their external customers. The initial security assessment reveals numerous high-risk exposures. The business unit becomes very defensive, insisting that the security team is creating obstacles that will prevent them from being successful and meeting their deadlines. At this point some important hand-holding is definitely required. This should include detailed discussions explaining what the security team is trying to accomplish and how it will eventually enable their business goals. It should be made clear that the DMZ is going to be certified for operation and the security team is going to help them overcome any imposed security requirements. Once they understand that the security team is a full partner in the solution, attitudes will quickly change and compromises will become realities. A success story in the making.

It is imperative for anyone creating a security program to understand the needs of their internal and external customers. The CISO must understand the background and history of the company as well as each viable business unit. What are the company's products and services? What are the business environments they compete in? Who are their competitors? What are the company's strategic plans? How can information security be a value added and a market differentiator?

CISOs must also understand that the information security team does not own the computer systems, but are internal security consultants to the businesses who provide an important but supportive role. CISOs should also understand the industry their company is competing in as well as the company's proprietary products and processes. How does the company work with its customers and contractors in this industry? Many of your information systems may be dependent on these proprietary processes and the level of protection that is required. Understanding the critical assets of the company is another key goal and will drive the allocation of limited funding. Finally, you should identify industry peers that you can call on to leverage experience and ideas.

Insider Notes: *It is imperative for anyone creating a security program to understand the needs of their internal and external customers. They must learn the background and history of the organization as well as each business unit.*

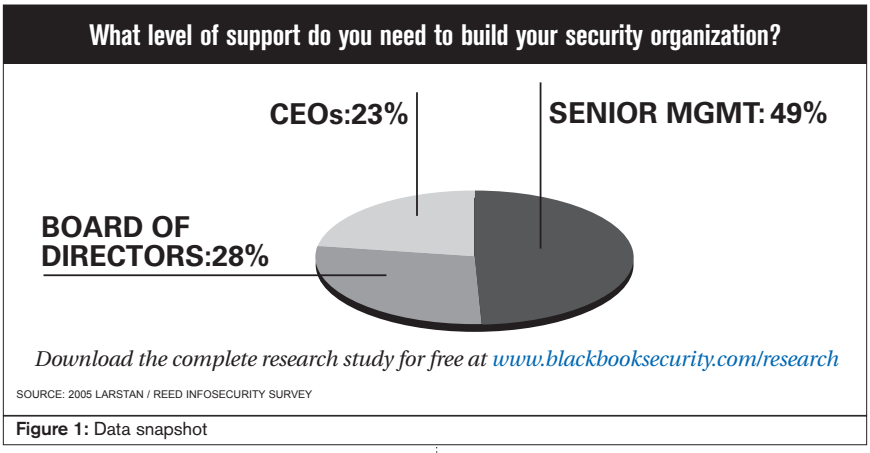

Figure 1: Data snapshot

Independent Assessment

A good way to obtain an independent view of your organization's information security posture is to conduct a full-scale security review (such as ISO- 17799) by a third-party security consultant as soon as possible. Also review prior assessments, internal IT audits, and SAS 70s for a comprehensive understanding of your company's IT security environment. By identifying the company's risk exposures and deficiencies, you can begin to develop your new information security "road map" for success.

Service Level Agreements

Another important step is to conduct a full-scale review of existing Service Level Agreements (SLA) and contracts for internal and external customers, as well as your security vendors. Ensure that they make good business sense and are in alignment with business strategies. Do your vendors provide a timely response? Are they giving your company the support it requires? Are your internal customers pleased with support from your antivirus team? When viruses impact the network, are they detected and cleaned within agreed SLA time lines?

Setting Expectations

Ensure that your organization issues a corporate-wide communication announcing the new CISOs arrival, your role, reporting structure, and support by senior management. This communication is vital to your future success.

You should also define the organizational structure for your department. This will include:

- Develop your vision statement
- Develop your mission statement
- Develop your organizational chart
- Develop function work streams to meet you goals
- Define expected roles, responsibilities and functions
- Define information security processes

During the CISOs interview process he or she should have negotiated the reporting structure for this new and critical role. Current industry trends support the CISO reporting to the chief legal council, chief financial officer, and/or the chief auditor. This is an important step toward maintaining independence (eGovernance) between senior IT managers, who often have different project priorities and funding requirements.

Another area for discussion is whether physical security should report to the CISO. The decision to incorporate all security into one reporting line may be simply based on the company's culture. There are many pros and cons on this subject and as such should be discussed on its own.

In essence, these are the basic building blocks required to build an information security team. Remember that your organization exists to support the business and therefore your information security team should reflect all strategic and tactical goals of the business.

Building the Security Roadmap

Once you have compiled and digested all of this information, organizational planning can begin. Again it is imperative for this plan to be in sync with the long-term business strategy of the company as well as its short-term tactical needs. Use nimbus maps, or flow-charting, but somewhere you need to get it all on paper. You should also consider hiring a project manager to coordinate and plan these activities.

Develop a program that will allow your team to demonstrate immediate progress to senior management. This can be accomplished by developing a project plan that incorporates incremental steps to achieve your goals. Hit some home runs quickly. Your "road map" should also drive the information security budget plan, ensuring that all designated priorities are properly identified and funded.

Insider Notes: No other security program will hit a home run quicker then the Information Security Awareness program. By communicating to the global user community, this program will also help you brand your new organization.

Establish Achievable Goals

Remember that your security "road map" should never advance unrealistic goals and objectives. Do not promise something you cannot deliver in order to impress your new boss. Once you lose your credibility, it will be hard to recover. Credibility is central to your continued success, especially with senior management.

If you are going to effectively challenge senior management, especially if their demands are unrealistic, you must always maintain confidence in yourself and your program. Information security is not considered a revenue producer and the information security program will only be important when there are serious risk exposures and/or compliance issues at hand. What happens when your information security team delivers a secure environment? Will the company continue to fund your team? Will it outsource or downsize? Start thinking about your second and third year strategic plans if you want to keep your program alive.

Implementation Steps – Setting Milestones for Success

Behind any good initiative there should be a well-thought-out strategy and detailed project plan. This includes tasking and ownership for deliverables. Tracking is essential, thereby providing regular status reports and metrics for all levels of management. All good project plans are going to have milestones and by creating incremental change, you not only demonstrate early successes, but also reduce risk exposures. An early success story that can provide company-wide exposure is a comprehensive security awareness program.

Security Awareness Program

The biggest bang for your planning efforts will be the creation of a security awareness program. The security awareness program will help you to brand your new organization and communicate throughout the global enterprise. Part of this branding effort involves educating the user community about the difference between IT security and the information security team. It sounds like a small issue, but you don't want the IT organization taking credit for all your hard work. Below are some easy winners:

- Create a security Web page on your company intranet with links from the company's intranet home page.
- Display pictures of your team and define their security functions, roles, and responsibilities.
- Begin weekly articles on your security Web site concerning information security tips and security topics that you want to socialize with your organization.
- Buy information security posters and display them at strategic locations around the company.
- Periodically deploy telephone message announcements concerning important security issues.

That's a success! Senior executives will walk out of their offices and see your posters. They will also surf the intranet and see information about your organization as well as read weekly security articles.

Insider Notes: The first step in transitioning from planning to implementation of an information security practice requires an understanding that this is a change process for the company. It is a new way of conducting business and it will precipitate a culture shift.

Once employees understand that there's a new security team in place, they will want to work with that team. At that point you should begin to develop and publish your security policies to the greater user community. After a strong start, just keep the momentum going.

The Steps to Implementation

This section is devoted to constructing your security department. Beginning with the planning phase discussed above, this section details how those plans are best put into action.

Communications Are the Key to Implementation

Transitioning from planning to implementation involves a major culture shift for the company. Intra-company communications and your awareness program will become essential factors when trying to change company habits, especially in a global environment. It is important to make sure that everyone is aware of the paradigm shift and becomes part of the solution and not the problem. There will be many partners in this process, especially since information security touches all areas of the company. In order to ensure success, management, company employees, business partners and consultants must be educated and trained.

The optimal place to start your security "road show" is to educate management on the goals and projects of your security program (including the vice president and director levels). Make every effort to meet them personally because you will need them to champion your security work efforts.

Be sure that your presentation is concise and to the point – providing scope, objectives, and an executive summary. Support from this level of management will ensure that the entire user community will cooperate with your strategic security goals and projects.

Draw on the expertise of the corporate communication and public relations teams. They will be critical to the success of your communications campaign. Your communications campaign should detail information security expectations for all users in terms of both compliance and cooperation. Your objective is to make them feel part of the team and the solution process.

The Importance of Goal Setting and Creating Good Public Relations

Precise goals and realistic milestones need to be established. It is important to build in some early successes. Some quick wins will demonstrate to the company that a return on their security investment is in progress.

As previously indicated, your security awareness program can be a quick success story. Weekly security messages and articles not only educate your users, but also give your team broad exposure. Employees will begin to observe information security posters on the wall or company bulletin boards. Information security screensavers should be installed on desktops and laptops. These quick security successes provide good public relations for security, get senior management's attention, and also move you to a higher level of security.

Employees will begin to recognize that there are new security processes and this will help change the culture in a cost-effective, painless, and transparent mode.

Defining Your Security Maturity Model

Everyone knows there is no such thing as 100% security, and as such there will always be risk exposures. So what is a "best-in-class" security model? This will depend on your industry, business practices, and management culture. You should generally drive toward developing an enhanced security model that has multiple layers of good security practices. Don't be confused with trying to define the perfect security model known as "best-in-class."

"Information security is not a revenue producer; as such, the security program is always on thin ice with the client. **Never provide unrealistic expectations.**"

Explore
Innovate
Share

Attend

Returning for an unprecedented 15th year – RSA® Conference 2006 features more than 275 exhibitors, 200 classes and over 14,000 attendees in search of the latest techniques and tools from the best and brightest information security professionals. If your job involves any facet of information security, RSA Conference 2006 is the best place on the planet for education, empowerment and enlightenment.

Register by November 18 and save $800 off the standard registration rate.
www.rsaconference.com

RSA® CONFERENCE 2006

FEBRUARY 13-17 | McENERY CONVENTION CENTER | SAN JOSE, CA

RSA CONFERENCE 2006 — 15 YEARS — EXPLORE • INNOVATE • SHARE

The World's Leading Information Security Conference and Expo

Cisco Systems · CA Computer Associates · Internet Security Systems · Microsoft · RSA Security · Sun microsystems · symantec. · TippingPoint · VeriSign

Your new security model will help you balance managing risk exposures with good cost-effective practices.

As referenced in the planning stage, you should engage an independent third-party security consultant to review and assess your IT environment from top to bottom. This will highlight your primary risk exposures and what steps require attention in order to develop an enhanced security model. Your information security program should also be in alignment with current and future business strategies. Third-party assessments should evaluate security processes, network processes, application processes, and business processes as an integrated solution.

Insider Notes: *To ensure buy-in, the new organization needs to build both a road show and an awareness and communication program to get the word out. The optimal place to start the road show is at the next tier down from the CEO staff, the line vice presidents and director.*

An ISO 17799 assessment will evaluate your overall information security environment and will also guide you to ISO compliance for your industry. With both of these assessments complete, your company will have its security framework and roadmap for the next couple of years. Now you can begin to map required security projects to the necessary funding required.

This will form the basis of your security business case, and should be presented to senior management for approval.

Pitfalls to Avoid

The two big pitfalls to avoid are not challenging unrealistic expectations and not surfacing jeopardies in real time. It is key to determine expectations, timelines, resources, and funding for your security projects. There is also nothing worse than letting a problem linger for days or weeks.

Recognize mistakes, take ownership, and come up with a plan to remediate. Overcoming your setbacks is about taking ownership (leadership), developing a new plan, and executing. On the other hand, don't sidestep difficult issues – accept the challenge, that's what leadership is all about.

- ***What are some of the most important characteristics of a security product vendor?***

There is no absolute set of questions to ask vendors, but below are recommended questions to consider when evaluating a new technology purchase:

1. Do you have a dedicated team to assess and respond to security vulnerability reports concerning your product?
2. What is your vulnerability response process and track record?
3. What process improvements have you made as a result of past vulnerabilities reported in your software?
4. What is your release strategy (are they grouped or individual releases)? In other words, how long do we have to wait for fixes to known software problems?
5. What training does your development and test organization receive on application information security matters?
6. What percent of your team is focused just on security issues?
7. Does your company monitor the latest attack trends in the underground (Cracker) community and consider how those trends affect your software?
8. Do you patch all currently supported vulnerable versions of your application/ platforms at the same time (or are they released as needed)?
9. Has a third party conducted an independent security review (code review) and what are the results?
10. Can you provide independent references that are using this product?
11. Application RFPs should contain the following:
    - *The terms and period of your security support agreements*
    - *Proof of security testing and vulnerability assessments during deployment*
    - *Review vendors' Common Criteria certifications or any other software certifications*
    - *Review application patch records (quality and quantity)*
    - *Future upgrades should be dependent on vendors' security records*

When evaluating any new technology, all company stakeholders should be included in the process. Remember that there will be many cross dependencies and you need all parties involved if you want the project to be successful.

Insider Notes: *The two biggest pitfalls to avoid are providing unrealistic expectations and not revealing dangers in real time. Set expectations for both time and funding, with time being the most important.*

Building a Security Organization

This section discusses in detail the measuring process employed when constructing a secured environment.

Understanding Metrics

A successful security organization is constructed on a foundation of four pillars: policy, awareness, risk, and metrics. The policy development, awareness, and risk assessment programs required have already been discussed earlier. This section will deal specifically with your security scorecard.

Your metrics should be designed to support business objectives, security operations, security projects, and to measure overall progress. These metrics will be necessary to support four primary areas: on-going security operations, new security projects, supporting internal users, and risk assessments.

Information Collection

The key to good metrics is good information, and the key to good information is a good method for collecting and evaluating that information.

The information collection process starts by identifying all information security work streams, functions, and processes. Policy, systems, users, and other resources are the drivers that will most impact your metrics. The ultimate goal is to develop the capability to automatically collect and track information that will help your team tell the information security story to senior management. Automation of the information collection process is essential if you want accurate and timely information.

Information security metrics are driven by two primary factors: the number of systems in the IT environment and the number of people who use those systems. A good Security Information Management (SIM) tool is a useful technology that allows the collection of logs from server and network devices to monitor, track, and review compliance matters. It can also be leveraged to provide your information security dashboard, which will allow your team to build and compile better metrics, thereby giving management more flexibility to make cost-effective decisions.

Collecting and developing metrics is meaningless without a correlation of events and an action plan. Cleaning viruses as they impact your network is a necessary work effort, but understanding how viruses get through your gateway and why they were able to infect so many systems is even more important. Enterprise correlation of information from many sources is the key to effective security management. By reviewing events from various systems and devices such as intrusion detection alerts, antivirus gateway logs, firewall logs, system logs, etc., you will develop a clear picture of how the virus entered your environment and was able to propagate itself. By making better use of reporting capabilities and correlating security events, you will be more effective at deploying your limited resources to mitigate risk exposures.

The Importance of Being Proactive

Get out in front of your primary risk exposures. To be proactive, an organization must not only be able to collect systems and device logs real time, but collate them into meaningful reports and take action on them. By collecting these logs and sorting them into meaningful reports, you may discover 25 failed logon attempts to your active directory. This should set off alarms and an investigation should be initiated to find out what took place. Maybe the twenty-sixth time the user was successful. Security departments need to be proactively looking at security metrics and making recommendations to senior management so corrective action can be taken.

A process to collect metrics can be developed, but it does little good if it isn't reviewed and acted upon.

Insider Notes: When researching these technologies, all stakeholders should be given a "say," from the network architects and engineers to the risk people. If it involves a security management product, a decision should not be made in a silo.

For example, one indication of a problem might be having a high number of password resets. You can usually predict that your call center, depending on whether it is in-house or outsourced, is probably receiving a lot of calls.

This costs money. You should be asking why are we having so many password resets? Is it that passwords are too complex and people can't remember them? Is it because employees are just forgetful, lazy, or not paying attention to their work? What is driving the password resets? You may have to employ a security awareness program to educate and train users on password usage, which is much more cost-effective than burdening the helpdesk with a deluge of calls.

Another example: correlation of internal maintenance scans reveals that a number of servers have high vulnerabilities. Your team has worked hard to develop good hardening standards, but you know there is always going to be the human factor to consider. IT organizations are always making changes to production, and new security patches are always being announced. Investigation of these vulnerabilities may indicate a poor patch management process or maybe a breakdown in your change control process. Regardless, metrics are not just for gathering information and generating reports, they are tools to solve security problems and reduce costs.

When Are Metrics Successful?

Developing a security dashboard and measuring your success through good quality metrics is another step toward achieving your information security maturity model. Achieving a full level of maturity is probably not possible since conditions will constantly change, and the same is true about metrics.

Probably the best measure of success is your ability to solve security problems and empower business processes through good metrics.

In order for senior management to understand the continuing value that your team adds to the organization and their return on investment, the CISO must continue to effectively communicate a strong information security posture, and that is accomplished by providing good strong metrics.

Presenting for Success

Develop your security "roadmap" early and provide good metrics to support that story. While the credibility and demeanor of the CISO is important, the presentation of your metrics to senior management will require simplicity and meaningful information that can be translated into cost-effective solutions. Using too many bar and pie charts may not communicate your story effectively, especially to non-technology managers. It has to be simple, but effective. Senior management will undoubtedly ask hard and challenging questions, so be prepared to support your metrics.

Periodically, your metrics are going to tell a negative story, and while that will be understood, also be prepared to provide a corrective plan of action.

Building a Security Organization

Third-Party Networks

Customers, business partners, and outsourcers are beginning to require companies to provide information concerning their security posture (security questionnaires) as a prerequisite before conducting business with them.

Much of this activity is driven by Sarbanes-Oxley legislation and SEC disclosures. As you build your security organization, you will be required to factor customers, business partners, and outsourcers into your security equation.

A company is only as secure as its weakest link, and when you extend your network to a third party, you have effectively increased your risk exposures exponentially. Companies planning to conduct business in the next five years will be required to certify to third parties that they are meeting all legal and regulatory requirements, such as Gramm-Leach-Bliley Act (GLBA) and Sarbanes-Oxley. If they are conducting business in Europe, they will have to meet stricter compliance with the European Data Protection Act (EDPA). Information security programs are at the forefront of these work efforts. Design your security model to reach for the highest bar so your organization will be in compliance with all critical laws and regulations.

Insider Notes: *It is also important to note that metrics are meaningless without action. Identifying the number of viruses hitting the network, successfully deploying data files, and fighting the virus is fine, but the question is why hasn't that virus been caught at the gate?*

Information Security as a Business Enabler

Your security program should find ways to leverage existing or new security solutions with other business units. It is essential that the information security

How critical are project managers in building an information security organization?

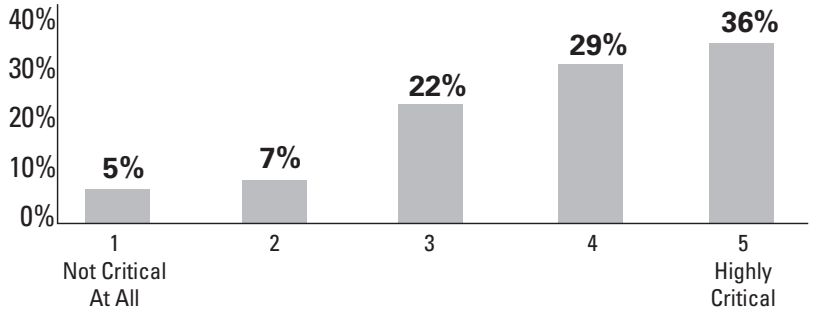

Based on the metrics, it confirms my belief that as a CISO, you should establish your own security project management office within your organization to ensure success.

Download the complete research study for free at www.blackbooksecurity.com/research

SOURCE: 2005 LARSTAN / REED INFOSECURITY SURVEY

Figure 2: Data snapshot

team develops close working relationships with the business units to understand their goals and business problems.

For example, your security team is trying to develop an encryption standard across the global environment. While developing this encryption standard, the security team must take into consideration all business units if they are going to have an effective enterprise-wide solution. At first it doesn't seem like there is a satisfactory enterprise e-mail encryption solution available. Eventually a member of the team recommends a Public Key Infrastructure (PKI) solution.

During this process the information security team learns that HR is concerned about their staff sending and receiving sensitive employee information through e-mail. Additional research by the security team learns that the merger and acquisition group is also concerned about sending sensitive information through e-mail.

After surveying other business units the team also discovers that this issue is isolated to just HR and the M&A groups. Working with the business units to understand their business needs and making them feel part of the solution becomes essential. The last thing you want to do is purchase an enterprise-wide PKI solution to find out that it is not being used and your limited funding is wasted.

After a little work on your part, you find out that you have a small group of users who will require e-mail encryption. Your team can accomplish their mission by deploying PGP to the HR and M&A teams. They will require some education and training to use PGP, but your team has established credibility, because you listened to their requirements and found a viable solution. It's all about being part of the team and offering the business affordable solutions.

Another example: the information security team would like to purchase and implement a SIM solution. After meeting with various IT and business units to discuss this project, the security team finds out that they can leverage this technology to other groups and enhance their ROI. The infrastructure team would like to use the network utilities tools and dashboard that SIM offers. The services division would like to leverage the central logging capabilities that SIM offers.

When a VP of sales or business development offers to partner with your information security team on new initiatives or customer offerings, you will know that information security has been successful. Differentiating your organization and security environment from other competitors will be critical, especially in a global economy.

Keep in mind that any time the information security team creates a security challenge for the business, it is essential that your team works with them until a solution is found, balancing good security practices with good business requirements. Ultimately, information security should not be viewed as the cybercop, but rather as a positive business enabler.

Insider Notes: *When senior management sees the value you add to the organization (a return on investment), and that you are producing appropriate metrics, you become a viable business unit, not just a drain on the bottom line.*

Security as a Market Differentiator

For many years before 9/11, senior management viewed security as a cost center that did not add to the bottom line. When companies had tight budgets, security was often one of the first line items to be cut. That thinking needs to change, and the CISO is on the front line of that effort.

Security departments need to be viewed as viable business units that make positive contributions to the business. As discussed, metrics are one of the key tools that will help you sell that story.

Information security is beginning to transform itself into a viable business unit and market differentiator. As third parties continue to require good security practices before doing business, CEOs will have to view information security as a central component of its core business plan. A company without good security practices that meets all legal and regulatory requirements will be a big loser in the emerging global marketplace. Sarbanes-Oxley has set a minimum security standard in the U.S., but there are other security benchmarks, such as the EDPA. When considering outsourced partners, a company should review their SAS 70 (financial audit for IT controls) for information security controls and compliance matters.

Companies that publicly advertise or certify that they have met generally accepted security standards will be able to use this as a market differentiator. In the long run, it will drive new revenue channels. Companies that have not developed or invested in good security practices will trail the pack.

About the Author

Stephen W. Foster was Chief Information Security Officer at Avaya Inc. He joined Avaya after a distinguished 20-year career with the Federal Bureau of Investigation.

"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™

The Vulnerability Management Lifecycle Approach to Application-Layer Security

AN EFFECTIVE AND HOLISTIC APPROACH

BY AARON NEWMAN

THE YEAR 2005 was distinguished by 50 security incidents compromising approximately 50 million pieces of sensitive information. Already it is the worst year in history for database hacks. High-profile data theft incidents, such as those experienced by ChoicePoint and CardSystems, exemplify what industry veterans already know: traditional security measures, such as firewalls, do not provide in-depth security at the application level, leaving database applications vulnerable to intruders. If the stakes weren't high enough, regulatory compliance requirements have further upped the ante. After all, whether it's significant transactions (relative to Sarbanes-Oxley) or personal information (relative to California Senate Bill 1386, the Health Insurance Portability and Accountability Act, and so on) – most data spends 99% of its life in the database. Protecting databases – an organization's "crown jewels" – requires a more comprehensive, layered approach.

At its core, security is about risk reduction. One of the most effective database security practices, defense-in-depth, involves multiple layers of protection to reduce the risk of intrusion. This layered approach is analogous to the defensive layers surrounding a medieval castle: drawbridge, moat, the outer wall, the inner keep, archers manning the wall, soldiers stationed outside the wall, etc. No single level of defense is infallible, and all these layers cannot ensure the castle will be 100% impenetrable. However, together, these combined layers of protection can make the castle – and its crown jewels – significantly less vulnerable to invaders.

Database security is similar. Protecting the database encompasses more than setting permissions. There are many layers of protection to consider when safeguarding your databases.

Most large organizations have antivirus software, firewalls, and sometimes even intrusion detection systems (IDSs) to protect their networks and host operating systems. Though these defense tools serve a purpose in protecting servers and networks, they are not designed to detect application-level attacks, nor are they capable of stopping such threats before damage is done.

Firewalls provide protection only at the network level – examining packets and determining whether an incoming request should be granted access to a given port. Firewalls do not understand database vulnerabilities or protocols (such as SQL) that may be used by attackers. Firewalls are also typically located on the edge of the network, where they are ideally situated to watch for attacks from outside the enterprise, but not threats from insiders.

In addition, by their nature, firewalls simply admit too much traffic to provide foolproof application protection. In today's world of virtual organizations and electronic commerce, an enterprise cannot afford to lock out customers, suppliers, distributors, remote employees, or contractors.

Similarly, though many enterprises have deployed IDSs to improve network security, these tools do little to protect core databases and applications. These systems scan the network and compare traffic and usage patterns to either historic trends or against the "signatures" of known network attacks. However, most IDSs are passive, scanning for suspicious traffic and alerting the network administrator, but not taking any action to stop the attack. They are also designed as forensic tools, gathering evidence to analyze an attack after the fact, as opposed to thwarting it in real time.

Firewalls and IDSs certainly have a place in a multi-layered security system. But they are not enough to protect organizations from internal and external threats, while still allowing appropriate access to applications and databases. A modern enterprise needs application-intelligent equivalents of its existing network- and host-based security platforms, which can discover, assess, and dynamically protect applications and databases against rapidly-changing security threats.

Many organizations have already employed an effective life-cycle management methodology at the network and host operating system levels. Enterprises should apply these same principles to the application layer as well.

The vulnerability life-cycle management process includes four main components as shown in Figure 1.

1. ***Establish a baseline:*** Through intelligent and complete discovery. It's 9:00 p.m., do you know where all your database applications are? Given today's easily deployed databases and the pace of change in an organization, simply knowing where the database applications are is non-trivial, yet crucial to implementing security. Only with a detailed understanding of what database applications are deployed, where they are, what releases they are running, etc., can true application security begin.

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™

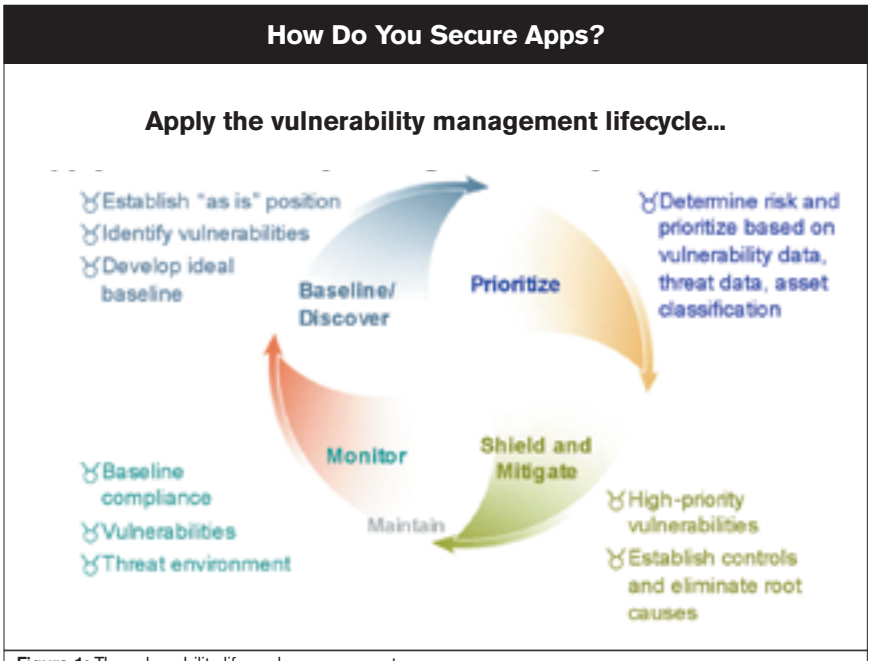

Figure 1: The vulnerability life-cycle management process

2. ***Prioritize:*** Assess risk based on asset classification and vulnerability. Proactively assessing the vulnerability of application components helps organizations minimize risk and evaluate compliance with their security policies. With this prioritization in hand, firms can inform the implementation of the next steps in the life cycle, or direct the roll-out of this process across their infrastructure, based on the importance of the systems at hand to their business and the severity of the present vulnerabilities.
3. ***Shield and Mitigate:*** Via vulnerability assessment and encryption. Having identified which systems are most critical to the business and quantified the number and severity of vulnerabilities on these systems, a firm should then conduct vulnerability assessment tests and proactively "harden" these components by removing the present. The system should ensure the installation of all current patches, that passwords have been changed from their default settings, and that required security configurations are set. Also, the system should produce meaningful security audit reports – prior to and after application deployment; they ensure new components get deployed securely and stay that way. With this step complete, organizations take a major step toward improved security. By regularly eliminating the known weaknesses an attacker might try to leverage, organizations significantly reduce the risk of a breach.

   Encrypting the most sensitive data further bolsters protective efforts. This "last line of defense" ensures that even if your database is breached, its most critical information remains protected . This step is crucial not only to thwart an attacker that manages to gain access to the database despite other protections, but also to defend against unauthorized access to data by legitimate users.
4. ***Monitor:*** Utilize your established baseline to monitor for vulnerabilities and your threat environment. The ability to detect and block attacks as they happen is essential. After all, zero-day threats (new vulnerabilities) and insider abuse are a simple fact of life for which threat signatures are either unavailable or do not apply (respectively). Thus, to complement proactive hardening efforts, organizations require real-time protection from both rapidly spreading new security threats and rogue employees who don't need to break-in to the databases – they already have access. Real-time protection also helps guard unpatched systems during the sensitive gap between when a new vulnerability is published and when patches are universally applied.

Clearly, database intrusion detection and security auditing comes with complexities. Monitoring your databases is useful, but monitoring is only effective when done in conjunction with a well-conceived and balanced security plan. Database monitoring is a layer of defense augmenting your overall database security strategy. When database monitoring is performed in parallel with vulnerability assessment and encryption, you can develop an effective life-cycle approach to database security.

When considering the use of database monitoring and security auditing solutions, make sure you select a tool compatible with other database security products. You can achieve an effective and holistic approach to security when you incorporate and integrate the solution at the different layers. As a result, you can fortify your castle (database) and crown jewels (sensitive data) from modern-day barbarians.

"Firewalls do not understand database vulnerabilities or protocols (such as SQL) **that may be used by attackers**"

About the Author

Aaron Newman is co-founder and the chief technology officer of Application Security, Inc. (www.appsecinc.com). In his current role, Aaron is responsible for defining the overall AppSecInc product vision. Widely regarded as one of the world's foremost database security experts, Aaron is the co-author of the Oracle Security Handbook, printed by Oracle Press.

ENGAGE AND EXPLORE...

The Technologies, Solutions and Applications that are Driving Today's Initiatives and Strategies...

CALL FOR PAPERS NOW OPEN!

SOA WebServices Edge 10th International conference+expo 06

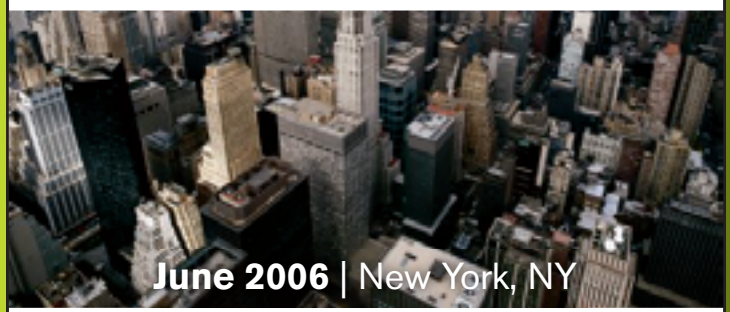

The Sixth Annual SOA Web Services Edge 2006 East - International Web Services Conference & Expo, to be held June 2006, announces that its Call for Papers is now open. Topics include all aspects of Web services and Service-Oriented Architecture

Suggested topics...

- *Transitioning Successfully to SOA*
- *Federated Web services*
- *ebXML*
- *Orchestration*
- *Discovery*
- *The Business Case for SOA*
- *Interop & Standards*
- *Web Services Management*
- *Messaging Buses and SOA*
- *Enterprise Service Buses*
- *SOBAs (Service-Oriented Business Apps)*
- *Delivering ROI with SOA*
- *Java Web Services*
- *XML Web Services*
- *Security*
- *Professional Open Source*
- *Systems Integration*
- *Sarbanes-Oxley*
- *Grid Computing*
- *Business Process Management*
- *Web Services Choreography*

CALL FOR PAPERS NOW OPEN!

2006 ENTERPRISE> OPENSOURCE CONFERENCE+EXPO

The first annual Enterprise Open Source Conference & Expo announces that its Call for Papers is now open. Topics include all aspects of Open Source technology. The Enterprise Open Source Conference & Expo is a software development and management conference addressing the emerging technologies, tools and strategies surrounding the development of open source software. We invite you to submit a proposal to present in the following topics. Case studies, tools, best practices, development, security, deployment, performance, challenges, application management, strategies and integration.

Suggested topics...

- *Open Source Licenses*
- *Open Source & E-Mail*
- *Databases*
- *ROI Case Studies*
- *Open Source ERP & CRM*
- *Open-Source SIP*
- *Testing*
- *LAMP Technologies*
- *Open Source on the Desktop*
- *Open Source & Sarbanes-Oxley*
- *IP Management*

Submit Your Topic Today! www2.sys-con.com/events

Sponsored by

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

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

Attention Exhibitors:
An Exhibit-Forum will display leading Web services and OpenSource products, services, and solutions

For Exhibit and Sponsorship Information - Call 201 802-3066

Produced by SYS-CON EVENTS

The Battle for the Desktop

SECURING THE ENDPOINT

BY TODD BRENNAN

AS NETWORK PERIMETERS become more porous, and endpoint security becomes even more critical, companies are struggling with the problem of unwanted software – whether it's new, unknown, and potentially malicious software, or simply known non-business applications.

Now, a new approach to endpoint security called Automatic Graylists™ is enabling IT professionals to regain control over spyware, malware, and other unwanted/unapproved applications with real-time, network-wide visibility and control – for maximum security with minimal effort.

The Problem of Unwanted Software

Unwanted software represents a root cause of a wide variety of IT problems, not just time, but money. For example, Microsoft estimates that spyware causes one-third of system crashes and unlicensed software can result in settlements in excess of $100,000.

Unwanted software includes worms, viruses, spyware, vulnerable applications, unlicensed software, and unsanctioned applications, and can arrive on corporate networks in many ways. The scope of the problem ranges from mobile users connecting to infected networks to end users unknowingly infesting their office desktops with spyware. Users can also knowingly download non-sanctioned applications such as peer-to-peer file-sharing, instant messaging, and mp3.

"Despite increasing public and corporate awareness about cybersecurity, the number of computer vulnerabilities in the second quarter of 2005 increased 10.8% compared with the first quarter," according to a new survey from the SANS Institute, which develops data and research on information security. In all, SANS discovered 422 new vulnerabilities, up from 381 in the first quarter.

To complicate matters further, spyware producers continue to make the legal case that their software is neither malicious nor illegal. While unwanted software encompasses a variety of applications that have different implications and characteristics, today's solutions typically focus on narrow aspects of the problem. So different tools have evolved to solve a subset of the problem. For example, most current security products attempt to detect only malicious applications of certain types and new malicious software is discovered every day. In fact, in September of 2005 alone, Sophos identified 1,233 new viruses. Unfortunately this approach has caused a proliferation of agents on systems that increase complexity and administrative effort.

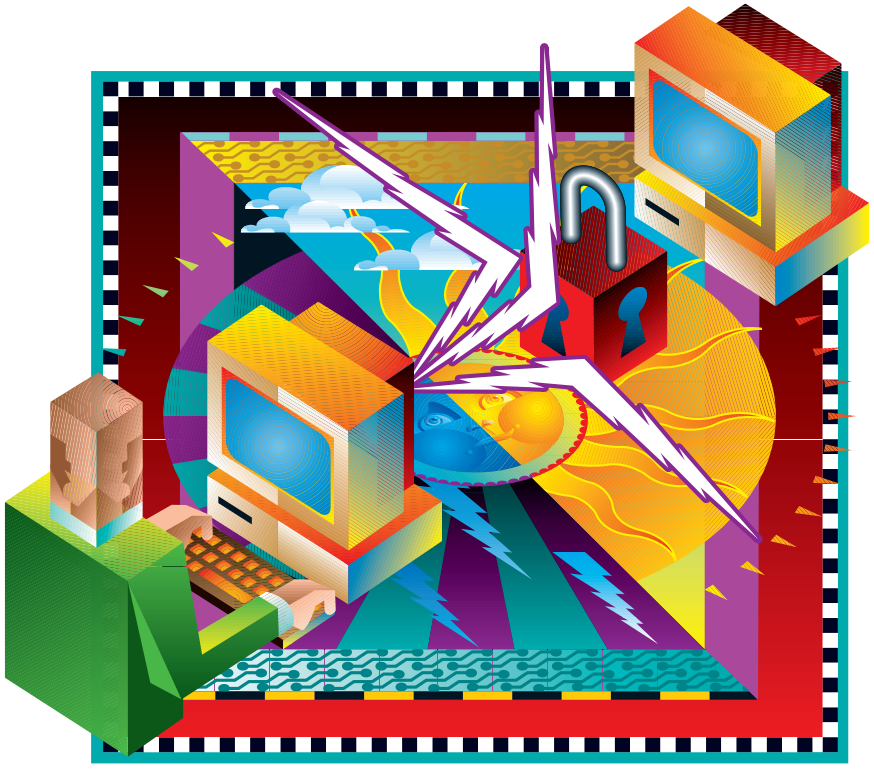

According to Gartner's Security Software Forecast, "Revenue for the security software sector will grow at a compound annual growth rate of 16.2% through 2009, with new license revenue reaching $11.4 billion." This double-digit growth indicates that current solutions to today's problems are simply not working.

Today's Solutions

Current solutions to the problem of unwanted software take a narrow view and typically target one problem, such as spyware. This requires IT administrators to install and manage an increasing number of agents on their systems. Furthermore, most products depend on perpetually out-of-date "blacklist" technology that compares executable files against a "known bad" list of bit patterns or behavior signatures. This approach leaves systems exposed to new, unknown, and potentially unwanted, software.

Other products use a "whitelist" approach that limits systems to a "known good" list of executable files. While this

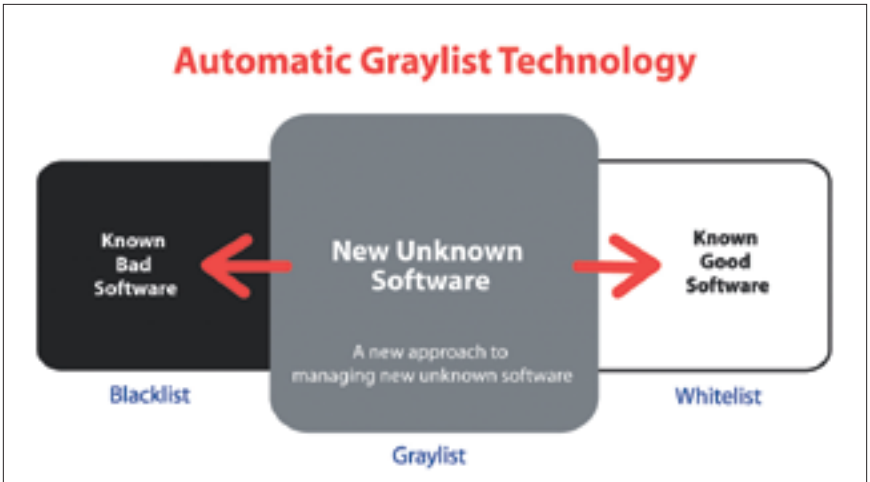

Figure 1: Automatic Graylist technology tracks and controls new and unknown software not previously approved by IT

approach in theory provides the best protection, these solutions are often overly restrictive and whitelist administration is daunting.

When these approaches fail, the result is late detection, false alarms, and complexity. Furthermore, none of these approaches gives IT visibility and control over all the executables in their infrastructure. Clearly, what is needed is a solution that lets administrators see and monitor what has arrived, when it arrived, who has executed it, and where it is now.

Introducing Automatic Graylists

So far what we've learned is that blacklists are eternally incomplete and whitelists are difficult to maintain in dynamic environments. In desktop environments, new applications, plug-ins, and Active-X components – wanted or unwanted – appear frequently and don't yet exist in a black or whitelist. Both approaches require additional information to increase security coverage and reduce administrative burden.

Now there's a powerful new mechanism to control unwanted software that integrates with existing IT processes called an Automatic Graylist. Unlike other alternatives that attempt to detect what's wrong or bad, Automatic Graylists track new unknown files – before they've been classified – and help automate their approval. When combined with whitelists and blacklists, Automatic Graylists overcome the drawbacks of today's other endpoint security solutions while providing the best possible protection.

Automatic Graylists don't rely on malware signatures or behavioral patterns and let the customer, not the vendor, decide what software is appropriate and approved to run on his infrastructure.

In this new model, new application files are detected in real-time as soon they appear on systems and are automatically added to the Automatic Graylist. They can be easily approved or banned, based on current security policy. This supports a dynamic environment, letting the administrator define the appropriate security policy for groups of desktops, laptops, and servers. With this information, IT professionals can create policies that automatically track and control software that hasn't been centrally rolled out or pre-approved by IT. When applied on a network-wide basis, this Automatic Graylist approach provides the best protection against unwanted software without a complex or costly administrative burden.

For host groups where graylisted software is blocked, this approach provides zero-day protection as new unknown files are effectively put on probation. For example, enterprise customers could block new files for 30 days after which the files are dropped from the Automatic Graylist. This buys customers the time to react while anti-virus signatures catch up with new zero-day attacks.

Be Prepared

IT managers, security managers, and executives are watching with mounting concern as the unwanted software on enterprise endpoints increases. Furthermore, current security products offer limited protection, at best, against these applications and the costly damage associated with them. Current technologies don't provide enough early detection and lockdown capabilities to deal with new unknown and potentially malicious code, which is typically the most disruptive. And enterprises can't respond quickly enough to avoid harm.

Now, with the introduction of Automatic Graylists, security administrators, for the first time ever, have gained visibility and control over unwanted files that install and execute on user systems. For more information about Automatic Graylists, please visit www.bit9.com.

References

- *Information Week, January 10, 2005*
- *Business Software Alliance, news releases, http://www.bsa.org/usa/press/*
- *http://www.sans.org/press/q2-2005update_release.php*
- *Summary of legal actions by spyware vendors: http://www.benedelman.org/spyware/threats/*
- *http://www.sophos.com/pressoffice/pressrel/us/20050930topten.html*
- *Gartner Forecast: Security Software, Worldwide, 2005-2009, Executive Summary. March 30, 2005*

For additional information about endpoint security, malware, and the problem of unwanted software in general, visit www.bit9.com, www.gocsi.com, www.gartner.com, and http://www.virusbtn.com/index

About the Author

Dr. Todd Brennan brings an entrepreneurial spirit and deep technical background to his role as chief technology officer of Bit9. Previously, Dr. Brennan founded Okena (acquired by Cisco in 2003), where he devised new techniques to defend against emerging computing threats. There he was also responsible for creating the company's technical and market vision, assembling the management team, and raising capital. Prior to Okena, he was a research staff member of the Satellite Communications Division at MIT Lincoln Laboratory. Earlier, he was a software engineer in the VLSI Advanced Methodology Group at DEC. Dr. Brennan got his PhD and MS degrees in electrical and computer engineering from the University of Wisconsin, where he won numerous teaching awards, and holds a BS in electrical engineering from Cornell University.

kim@bit9.com

Securing the Enterprise Beyond the Perimeter

DEPERIMETERIZING SECURITY ARCHITECTURE

BY RICHARD MOULDS

RECENT HIGH-PROFILE SECURITY breaches have taught us a clear lesson: organizations that rely primarily on a secure perimeter to protect sensitive data are fooling themselves. This year, hardly a week has passed without headlines about a security breach involving sensitive data.

However criminals get the data, whether through a traditional perimeter breach, use of insider credentials or outright theft of physical storage media, the lesson is the same. Organizations can no longer regard everything inside the traditional perimeter (people, machines, and networks) as "trusted," requiring only a "soft" approach to security that consists primarily of procedural controls and weakly enforced permissions.

It's an approach to IT security that's like a candy M&M: once criminals penetrate the hard shell that protects the network from the wholly untrustworthy public Internet, they can easily devour the data at the soft center. Actually they often don't have to penetrate the perimeter at all. They can simply go around it by stealing unencrypted backup tapes, for instance, out of the back of a cargo van.

Not only are attackers constantly blowing open security cracks in perimeter security, but enterprises themselves are also willingly, and often unwittingly, contributing to the perimeter's disintegration.

For example, virtual private networks frequently tunnel through the perimeter, which often provides all-or-nothing access to network resources. Web Services, which are starting to finally fulfill the early hype, are meant to interconnect business processes and often reach into the core of an enterprise network. Factor in the mass of mobile devices, wireless networks, portable media storage and off-site data archival, and it's not outlandish to suggest that

there really isn't a perimeter at all. Instead, enterprises need a "jawbreaker" model in which the network is "hard" all the way through to the center.

Drivers Behind the Jawbreaker

Unfortunately the traditional perimeter model doesn't just fail to provide adequate security. It's also far too expensive and inefficient to deploy, given today's far-flung workforce. Enterprises have to manage an exploding number of network connections for employees working at home, traveling and staffing remote offices, not to mention the connections they've built to the networks of partners, outsourcers, and customers.

Enterprises need a unified management approach to the identities of users, their rights and roles, and ultimately the enforcement of those rights. The search for a unified approach has led many security experts to believe that security will soon be deperimeterized.

In a deperimeterized world, every user is "remote," whether he's on the corporate campus or in a coffeehouse halfway around the world. Instead of building a perimeter around the network, in a deperimeterized architecture there's a virtual perimeter around every user or internal system that establishes "islands" of trust that securely exchange information.

The Jericho Forum (http://www.opengroup.org/jericho), a security organization recently founded by corporate CIOs, is taking a stab at defining the requirements for both the short-term and long-term transition to a deperimeterized world – a unified world with an inherently less expensive, more consistent approach to identification, authentication and authorization. By and large, its vision doesn't require the development of brand new, whiz-bang technologies, but rather strings together existing technologies into a unified whole.

The Jericho Forum's vision is no pipe dream. It's already underway. Computer manufacturers like Dell, Hewlett-Packard, IBM, and Fujitsu have all made trusted platform module (TPM) technology a standard feature in their enterprise-class laptops, enabling users to securely lock away in hardware the secret digital keys that are crucial to encrypted communications. These keys let users securely encrypt and decrypt information with their laptops, and give administrators the ability to verify not only that a user is safe, but also that the user's machine is safe.

Dell, for one, has gone a step farther and has put smart-card technology in its laptops so network administrators can assign a digital identity to each user instead of relying on the notoriously insecure usernames and passwords.

Pervasive Encryption

So a world in which every user is a secure "island" raises important questions like how one know who's actually "on" each island?

The foundation of a deperimeterized security architecture is knowing whether users and their machines are who and what they should be. Enterprises will have to use strong methods of authentication such as smart cards, USB tokens and ulti-

mately biometrics to validate users and embedded digital identities to recognize devices such as laptops, phones and even peripherals.

It also begs the question: How will these islands communicate securely with one another?

At the end of the day, the only sure way to enforce confidentiality is though encryption. No enterprise in its right mind would ever send sensitive data across the Internet without encrypting it first. That mindset is now starting to be applied to all networks. There are well-established means for securing data as it travels "outside" the traditional perimeter, means that can be re-applied in a deperimeterized world. SSL, virtual private networks, and Web Services will all be used to link up the islands protecting data "inside" as it moves between cubicles or campuses.

You also have to ask: How will enterprises protect sensitive data and the processes that use them once they've arrived on the islands?

The reality is that pockets of stored data are virtually everywhere and that much of this data is sensitive in nature. In a deperimeterized world, the situation is probably going to get worse. There is a "data at rest" problem that goes well beyond backup tapes. There will be need to be the islands responsible for protecting the data on the island – whether the data is stored in a database, file system, tape drive, or the laptop's hard drive. In some cases, tightly integrated access controls may suffice but, once again, encryption will often be used to provide a last line of defense. If all else fails, a thief's efforts will be in vain – he may have access to data, but because it's encrypted, he won't see anything except gobbledygook.

Clearly, encryption plays a pivotal role in a deperimeterized security environment. But as encryption penetrates deeper into enterprise operations, enterprises will need to deploy new systems to manage – cost-effectively – the exploding number of private keys on which pervasive cryptographic security will depend. There will have to be a mechanism for recovering lost data and separating duties.

It's a big challenge, but once deperimeterization becomes a reality, the payoff will be enormous. Not only will the headlines about security breaches recede but enterprises will be able to expand their networks efficiently and securely to include remote employees, new branches, partners, customers, and outsourcers.

It's only a matter of time before the walls fall down. The question is whether there will be systems and policies available that can raise the security bar sufficiently to cope. Life in a deperimeterized world might be a liberating experience and should certainly be less costly in the long run.

The security industry still has plenty of work to do. What seems clear is that the using cryptography will become more widespread, often under the covers, but nonetheless a fundamental component behind strong authentication and enterprise-wide data protection.

About the Author

Richard Moulds is nCipher's vice-president of marketing. He has a bachelor's degree in electrical engineering from Birmingham University and an MBA from Warwick University in the U.K.

rmoulds@ncipher.com

Disaster Recovery Plans: Be Prepared

LEARNING FROM RECENT EVENTS

BY MIKE ROZLOG

IT WOULD SEEM only logical that after 9/11, one of the most horrific days in American history, corporations large and small would be ready for unforeseen catastrophic events. However, by one recent estimate, less than 38% have put a complete disaster recovery plan in place – the policies, processes, procedures, and architecture to deal with unforeseen events. In the wake of Hurricanes Katrina and Rita, IT managers are again forced to reassess how well prepared they and their organizations are to manage through and recover from natural or man-made disasters.

Understanding the strategic goals and requirements for surviving a catastrophic event is one thing, but actually having a set of guidelines in place for handling the tactical issues involved is quite another. Ultimately, the goal is to recover and restart business operations quickly and efficiently. But successfully arriving at that goal involves doing a hundred little things before, during, and after the event. This can present a challenge for IT organizations just trying to keep up with day-to-day operations, but when one thinks of what's at stake, it's imperative. Advance planning, preparation, awareness, and testing are the keys to the success or failure of a disaster recovery plan.

Planning for the Worst

CIOs and IT managers should ask themselves, "How do I deal with 21st century threats?" Imagine what you would do if you were in faced with a disaster. Would you be in good shape, or would you find yourself scurrying to implement a recovery plan? If the latter, now ask yourself: How long can my company's IT operations afford to be offline?

Preparing data and IT systems for potential disaster requires a combination of well-planned procedures and thoughtful policies. All companies should ask themselves this "What if..." question at every stage of their

operation: "What if the power was down for over 10 days? Do I have a plan to deal with it?"

The problem today is that most software development organizations view disaster recovery at best as an after thought, although prevention is the key. The inability to resume everyday operations quickly and protect resources can be detrimental to a business and its community, and the companies most prepared for the unexpected are the ones that will reduce the risk of operational downtime, protect their valuable intellectual property and get back up on their feet quickly.

Approach Planning Incrementally

Organizations have to determine what it will take to protect them during a disaster and how long they can manage before full restoration is required. By asking the "What if..." question, organizations can be better prepared for an emergency by defining an initial plan and then refining it with more and more detail. It's best to create such a plan incrementally, starting out by asking the simple questions first, and then moving to more complex queries as appropriate for the business. Keep in mind that this project is an on-going one. As your business changes and new initiatives are defined, your plan has to change too. So reviewing your disaster recovery plan regularly should be a key part of the project.

When looking at disaster recovery from an application lifecycle management perspective, questions such as these should be addressed:

1. From a Requirements Stage: What are the business, system, and data requirements needed in the event of a disaster?
   a. Are these items isolated so they can be moved in an emergency?
2. From a Design Stage: Is there adequate design for failure, is the architecture defined to enable recovery?
   a. Where are the hardware, software, and other assets located and where will they be if something happens? What are the steps to restarting or continuing these assets?
3. From a Development Stage: So the systems that are being built include concepts of failover, redundancy, and co-location?
   a. Do the architects and developers understand the importance of disaster recovery as it relates to the systems being created, and has management participated in these requirements?
4. From a Testing Stage: Could we implement a copy of our environment and test to ensure that it is stable and live in the event of an emergency?
   a. Where are the assets going to be located in an emergency?
5. From a Production Stage: Could we move all operations to another location seamlessly and immediately if we had to and is the location and infrastructure needed to do this available and ready?
   a. What are the steps needed to move the company assets?
6. From Configuration Management Stage: Is there adequate backup and redundancy built-into our large software investment?
   a. Do you have a recent copy of your entire business off-site? And is it stored in a safe place?

Disaster Recovery Planning Checklist

Besides asking these questions, there are other measures IT executives can take to protect their people, information, infrastructure, and assets. The following are some steps companies can start to do today to ensure that their systems have the best chance of survival:

- Create a disaster recovery plan. Once you have a plan in place, communicate it, rehearse it and keep it updated.
- Review software and hardware contracts to ensure that proper licensing and contingencies are in place to help in the case of an emergency.
- Enhance the company's software development methodology to include the guidelines needed for disaster recovery in every phase of the development process.
- Test the plan to see if it will work. Contemplate the worst possible case scenarios. If your company already has a disaster plan in place, great. You're in far better shape than most. Once you've made a plan, test it. As we all know, plans look great on paper, but when it comes time to execute things don't always go as you expect.
- Backup. Many of the plans for data recovery are way too limited. Here are a few questions companies can ask themselves to ensure their backup plans are extensive enough:
  I. Do I have a current backup? Is there data missing from real-time to what was backed up, and does that matter?
  II. Do I have an alternate location from which to conduct business? If you have a good backup, and the data is current, is there a place to run the software or actually failover to?
  III. If the power outage is widespread and all businesses are down locally, will the company supplying the recovery location be overwhelmed? If so, what are our alternate plans?
  IV. Does the disaster recovery location have the staff to actually run the business? If not, what is the plan to get competent people there?

Many of us today are thinking about our own disaster preparedness. What lessons can be learned from recent events? What can be done differently next time? Is it possible to be prepared for every contingency? Coming up with better plans for the future – and executing on them – may take some time. However, adapting existing technology to new and unforeseen circumstances is something that we can start doing today. The unpredictable nature of events that can cause IT disruption continues to be a threat and one that businesses can't afford to dismiss. ■

Resources

- http://www.borland.com/us/products/alm/index.html
- http://www.borland.com/us/products/core_sdp/index.html

About the Author

Mike Rozlog is Borland Software Corporation's chief technical architect. He spends a great deal of time discussing and explaining all the technical and business aspects of Borland products and services to audiences and analysts worldwide. If you get a chance, check out his blog at http://blogs.borland.com/MichaelRozlog/ which is just getting started. Mike has been published many times, and his latest collaboration is Mastering JBuilder *from John Wiley & Sons, Inc.*

ISSJ | Advertiser Index

Advertiser	URL	Contact	Page
E7Software	www.e7software.com/risk	800-824-4717	21
Forum Systems	www.forumsys.com	866-333-0210	7
InfoSecurity	www.infosecurityevent.com/storage		Cover III
Imation	www.imation.com	www.imation.com/usedtape	9
ISSJ	www.issjournal.com	888-303-5282	33
NTP Software	www.ntpsoftware.com/learn	800-226-2755	23
RSA Conference 2006	www.rcaconference.com		17
SafeNet	www.safenet-inc.com/hse/15	800-697-1316	Cover IV
Storage Networking World	www.snwusa.com		13
SurfControl	www.surfcontrol.com/go/threatshield	800-368-3366	Cover II
The Security Awareness Company	www.thesecurityawarenesscompany.com	727-393-6600	31

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

Reducing TCO Through Mainframe Resource Optimization

AND MEET THE DEMANDS OF THE CUSTOMERS AND BUSINESS

BY JOHN ALBEE

ORDERING ADDITIONAL MAINFRAME hardware was once a regular, accepted part of the budget cycle. This process made capacity planning a far less challenging task than it is today. Performance problems, regardless of the cause, were easily addressed by adding more hardware. Performance analysts and capacity planners were able to deal with performance issues with little concern about the cost.

Economic uncertainty and recent world events have changed this paradigm. Now that every dollar in the information systems budget must yield a maximum return on investment, hardware upgrades are delayed as long as possible. In today's world, simply adding new hardware is not the most efficient or cost-effective way to manage performance problems. At the same time, reductions in staff from downsizing and the retirement of experienced mainframe technicians are causing mainframe technical expertise to diminish. These issues make performance management that much more of a challenge.

The demand for continuous systems availability and reliability is increasing exponentially. What was once a reasonably controlled user population has expanded to everyone with an Internet connection. Web-enabled legacy applications are causing transaction volumes to explode, putting a greater strain on IT resources.

"Do more with less," is the mantra, but what is the best way to accomplish this while providing the required service and performance? While hardware costs are dropping, software and people costs are increasing. As the total cost of ownership (TCO) rises, each business transaction becomes more costly. One of your many challenges is to control costs while meeting service level objectives. In today's world, simply adding new hardware is not the most efficient or cost-effective way to manage performance problems. Contrastingly, TCO can be best reduced by optimizing your existing resources, improving application performance, and deferring costly CPU upgrades.

The Old Ways Aren't Enough

The traditional methods of dealing with performance issues are seldom adequate in today's environment. Many system programmers and performance specialists have learned to work around performance issues.

Not too long ago, well-defined batch and online processing windows made it possible to change processing times in order to take advantage of well-known periods of low activity (valleys), where resources were more plentiful. Today, while batch is still a key workload, online processing occurs 24/7, turning the picture of yesterday's peaks and valleys into a plateau of near-constant demand. Online applications are the priority workloads day and night. Deferring work is not an option, and moving it can be a risky proposition without a way to test the impact.

Because of this, many companies looked to migrate work to distributed systems (DS) such a UNIX and Windows, but the costs of rewriting applications often proved to be prohibitory. In addition, three-tier environments were heralded as the "next new thing," but many companies became aware of the lack of cross-system expertise to manage the enterprise.

Adding to the challenge, hardware upgrades and tweaking system parameters often resulted in smaller performance improvements than expected, considering the outlay of time and money. In many shops, more than half of performance problems originated from inefficient application design and, with pressing business deadlines, programmers are forced to make it work, rather than make it work well, allowing for errors.

If optimization and tuning opportunities are ignored during the development cycle, you will pay for it later – in time, people, dollars, or an application's inability to scale. No matter how much CPU or system timing is done, inefficient applications place additional demand on the system. Industry analysts have demonstrated that it is 10 times more costly to resolve a performance problem in production than during development and testing. Time and again, these performance problems translate into lost business opportunities.

The New Ways Exist

The mainframe environment is dynamic, with daily changes for maintenance and new development. The ability to tune a complex batch window or to manage a high-demand CICS system is rapidly becoming a lost art. The right tools are essential to manage these dynamic and complex environments. The old manual tuning and optimization processes that worked so well in the past are simply not adequate to meet the demand and data volumes that exist on today's systems.

To address the changing environment,

companies must leverage performance and capacity management solutions that enable you to get the best results from existing resources. These automated solutions should:

- Model performance and plan for growth
- Manage application quality
- Optimize batch processing
- Optimize CICS processing

Model Performance and Plan for Growth

To reduce costs and process data efficiently, identify targets – workloads that use a large amount of costly resources – without putting excessive artificial loads on the system. To do so, companies should implement a performance management solution that allows IT managers to track work down to the individual address space and drill down to find candidates for resource optimization. In addition, enabling such a solution will allow users to analyze CICS, IMS, DB2, and MQ transactions.

After you have identified candidates for resource optimization, it is important to test tuning options and moving workloads (or parts of workloads), and changes in the transaction mix, to ensure that production response time and turnaround remains within agreed limits.

With these tools, it is easy to test a myriad of solutions and select the best price/performer. If the hardware costs of disaster recovery (DR) are of concern, DR strategies can be tested to ensure that acceptable performance can be achieved in a variety of situations. Though the CPU impact can generally be assessed with a spreadsheet, the impact on throughput and response times requires an understanding of queuing theory (which is the core of analytic modeling).

A common question that is posed to capacity planners is: "How much will this new application cost when everyone is using it?" Users can answer this question by modeling volume changes down to the individual address space. This process demonstrates the cost of maintaining acceptable performance at the new volumes. Knowing exactly how much hardware is needed, and when it's needed, simplifies the budget process.

Application Tuning

Applications and systems are more complex now than ever before. Specialization has become the norm. Application developers and systems programmers can have an impact on performance, but they do not focus explicitly on application performance. Without an automated way to manage application quality, a shop might never recognize a poorly designed application that consumes excessive resources yet meets acceptable service levels as a performance improvement opportunity.

The demand for quick time-to-market coding forces developers to push code into production too quickly. Time is limited for adequate design analysis and testing, and other factors, like high-level languages, further complicate the environment in which they run. Due to the lack of attention to application tuning, it is often difficult for an application programmer to know which code structures will result in less efficient processing.

Until recently, the goal of improving performance through application tuning was not considered critical – to reiterate, new features and increased functionality were the goals. In reality, many application-tuning opportunities relate to problems that were introduced when the application was coded. Therefore,

application tuning is a significant opportunity for large cost savings.

Application quality management (AQM), a methodology for proactively optimizing mainframe application performance throughout the application life cycle, automatically targets candidates for performance analysis while prioritizing opportunities for analysis. The AQM process provides automated application measurement, automated and targeted performance diagnosis, and prioritizing performance analysis, which results in significant IT savings through deferred upgrades and resource optimization.

Manual tuning procedures are time consuming and inefficient, and few organizations have the luxury to operate this way anymore. By using this process in the development cycle and automating the application tuning process, you can avoid performance disasters.

Batch Optimization

Even with a gifted team of performance analysts, optimally managing the complex batch environment and a myriad of CICS regions is a daunting task. In the 24/7 world, the static settings that govern batch jobs and the static tables that manage CICS usually fail to deliver optimal performance. Yet it is impossible for anyone to continually tune the settings around the clock. There exist solutions that optimize batch and CICS processing to ensure good performance at low resource cost.

Batch processing continues to grow, and the batch windows continue to shrink. Continuous applications processing and changing business models place more demand on existing jobs. Optimizing batch performance is no longer an option – it is now a critical necessity. Batch optimizing automates the complex task of batch performance tuning by increasing parallelism, optimizing data access, and keeping as much data in memory as possible during processing.

Three areas offer opportunity for batch optimization.

Data Optimization

A common myth that is associated with file buffering is that more is better. Data optimizing helps evaluate the file type and determines the access method of each file to dynamically optimize buffer resource allocations. Optimizing I/O access can translate to huge reductions in the elapsed time for batch processing because data stores on DASD is still significantly slower than data in memory.

Job Step and Job Parallelization

Combining data and job optimization strengthens the impact of an optimized batch workload. The batch cycle comprises job steps and jobs that are both dependent and independent of a previous step's processing. When independent jobs and steps are run in parallel, the resources complete sooner, making resources available for other sources. As more jobs complete sooner, less batch processing time is required by each workload. This process shortens the batch cycle for jobs within the critical path.

Job Piping

Batch optimizing improves parallelism by allowing jobs to pipe data into each other, causing a job stream to complete faster. The technology also extends to the piping capability across LPARs within a sysplex. By piping data across two or more LPARs, jobs can run in parallel on alternate partitions or processors to further increase parallelism, offering significant reductions in batch processing times and helping widen the ever shrinking batch window.

CICS Optimization

CICS has always been a core component of any large business, and now this subsystem has expanded its role by becoming a cornerstone of business strategy. Web-related traffic transactions often differ from legacy application traffic. The dawn of the Internet age, increased distributed processing, and increased enterprise database access have increased the importance of CICS. The number of regions continues to grow, and the use of a single application owning region (AOR) is no longer limited to a single application. CICS environments are becoming increasingly complex, forcing IT personnel to find new and innovative ways of proactively managing performance. The service level agreements that companies have in place must be met, or client relations and business suffer. Simply put, IT managers need a way to correct problems with minimal impact to users and other crucial workloads.

Automation is crucial in the complex CICS environment. IT staff should look to implement solutions that dynamically optimize CICS performance in response to workload peaks and valleys, as well as dynamically optimize resources, often before a performance monitor would have found it.

In addition, since the primary CPU cost of CICS relates to the large number of MVS waits that CICS issues as it processes the transaction workload, managing these calls to ensure that critical work completes faster with an optimized environment is essential to business flow. The workloads use fewer resources, so more are available to handle other workloads.

The goal of every IT organization is to reduce the cost of running mainframe applications while meeting the demands of the customers and its business objectives. The traditional approach of moving workloads or simply adding more hardware is no longer the best option. The goal of a business is to keep the applications running more responsively while consuming fewer resources, allowing them to defer expensive upgrades while providing customers with the service that they require. ■

"The goal of every IT organization is to reduce the cost of running mainframe applications while meeting the **demands of the customers and its business objectives.**"

About the Author

John Albee is director of mainframe solutions, BMC Software.

FREE
Exhibit Hall Admission!
www.infosecurityevent.com/storage

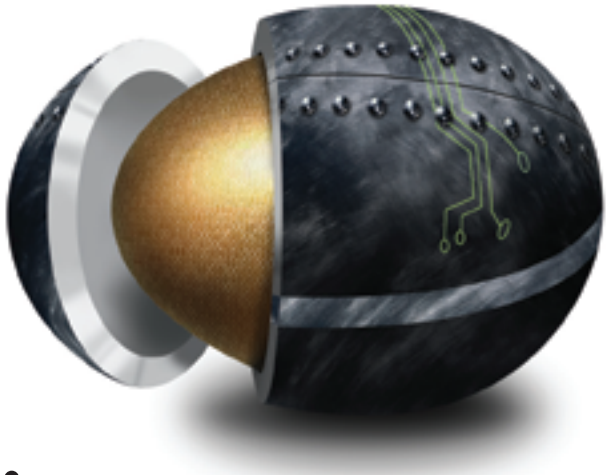

Protecting your assets is more challenging than ever.

Attend infosecurity NEW YORK

December 6-8, 2005
Jacob K. Javits Convention Center

FREE KEYNOTES

Tom Ridge, Opening Keynote
Former Secretary, US Department of Homeland Security
Wednesday
December 7th • 11:30 am

Bruce Schneier, Luncheon Keynote
Founder & CTO, Counterpane
Thursday
December 8th • 11:30 am

Honorary Conference Chairman

Howard A. Schmidt
Chief Security Strategist, US CERT
President & CEO, R & H Security Consulting LLC and Former White House Cyber Security Advisor

Infosecurity delivers information, education, and networking for a more secure and compliant infrastructure

The stakes have been raised. Your job of protecting your business and its information assets continues to increase in difficulty. Protecting against outside threats like hackers, spyware and targeted Trojans is now coupled with the need to be current on the latest regulations and compliance issues, guarding against data leakage, and the ever present insider threat. Infosecurity is geared to assist you in maintaining the confidentiality, availability, integrity compliance and risk mitigation of your organization.

Conference Tracks

Security Leadership Conference Series:
- **T1:** Security Management
- **T2:** Technical
- **T3:** Emerging Threats
- **T4:** Defense-in-Depth

T5: Wireless & Mobility Security – NYMISSA
T6: Compliance & Governance – Concordant, Inc.
T7: Privacy: Issues for Clients, Customers and Your Organization – IAPP
T8: Best Practices in Achieving Financial Justification – Larstan's Black Book Series

CISSPs/SSCPs Earn Up To 18* CPEs

Only (ISC)²®, the leader dedicated to educating, qualifying and certifying information security professionals and Infosecurity, the global leader in information security events, can offer such a high caliber education program. The CISSP® and SSCP® credentials identify you as an individual capable of developing and implementing solid information security practices.

**Combination of full conference and pre-conference workshops*

Real Risks – Real Solutions

Information security technology is complex and ever-changing to meet continued threats. Visit with over 150 leading vendors at Infosecurity to evaluate competing and complementary products and services such as Access Control/Authorization, Assessment & Audit, Authentication, Content Filtering, Encryption, Network Security, Perimeter Security, Security Management Products, Storage and more.

REGISTER TODAY AND SAVE!

For early-bird conference discounts and free exhibition admission, register online before October 18 at **www.infosecurityevent.com/storage**

Premier Education Sponsor: (ISC)²
Exclusive Global Media Sponsor: SC Magazine
Premier Media Sponsor: CSO
Silver Sponsors: PatchLink, Citadel Security Software, Websense, Juniper, AppSense, watchfire, Network Intelligence, nCircle, Secured eMail, TriGeo, AirMagnet, Internet Security Systems

Education Sponsors: ISSA, Concordant, Black Book Series, iapp
Media Sponsors: CIO Insight, IT Hub, Network Computing, Baseline, IT Architect, eWeek, Secure Enterprise
Featured Channel Media Partner: Business Solutions
Technical Track Sponsor: ICSA Labs
Produced and Managed by: Reed Exhibitions

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