Investing Wisely in Windows Security

by John Q. Walker, Ph.D., *NetIQ Corporation*

johnq@netiq.com

Investing in "information security" for the Microsoft Windows computers in your network is a lot like investing in "securities." As with any investment, you need to be proactive and vigilant even after you make the initial budget outlays to secure the systems. This paper looks at ways IT budgets are typically allocated for Windows security and suggests treating these limited budgets as you'd treat your personal investments. Six poor security investments are introduced and discussed, followed by six investments likely to produce excellent returns.

Contents

Treat Windows Security like Your Other Investments

Securities and security – we know we must invest wisely in both. That's why many of us devour the monthly financial magazines. They have front-cover headlines with titles like "Where to Put $1000 Now," "Safest Investments for the New Year." These magazines have a mixed bag of readers: some who are complete novices and need introductory articles, and others who are experts looking for the latest tips.

Reading these magazines over the long term, you see that they have common, repeated recommendations, like investing in high-quality growth stocks, lowering your credit-card debt, and avoiding get-rich-quick schemes that are too good to be true.

Windows security can be viewed similarly. Certain data-security topics appear over-and-over again. Discussions of Windows security usually include the active buzz surrounding the latest threats and what to do about them. In addition, they draw from a body of accumulated wisdom on ineffective ways to spend your security budget.

This paper describes security spending with the same approach used by monthly financial magazines. We offer our opinion on six poor investments, places where your budget dollar probably won't get the best return. We then offer six good investments. Consider these as you do your budgeting. We conclude with some overall advice that applies remarkably well to both the financial world and IT security world: insurance and diversification!

Six Poor Investments

Poor security investments often result from lack of planning. It takes focused effort to anticipate the problems that can occur and build a good plan for reacting to them. Most money-wasting IT investments involve frittering away resources on things that don't matter much in the larger scheme—or fighting disasters, where everything has to be dropped and lots of things get broken.

The SANS Institute maintains a Web page listing "Mistakes People Make that Lead to Security Breaches" [1]. It's worth browsing their site for some reinforcement of the foibles described below.

1. Securing Low Risk Areas

Your security budget is finite, and you can't protect everything.

Know what's important to protect. Prioritize your time and resources – don't waste them on stuff that's not so important (no matter how much fun it may be!).

Consider this candidate list of what's important.

- Human lives
- Customer records and information
- Employee records and information
- Financial data
- Trade secrets, business processes
- Employee productivity
- Reputation of the organization

Given the above list, the following don't seem quite so important.

- Hardware and software
- Things that are insured
- Things easily replaced
- Productivity of the security team

Start by securing the servers – they're the place where the most valuable information is usually congregated. Give each database server, LAN file server, and Web

server close attention. The security of Microsoft's Internet Information Server (IIS), its Web server, got lots of press coverage recently. IIS is widely deployed, often by those with little experience with its vulnerabilities, so is a frequent target of hackers. Use Microsoft's *IIS Lockdown Tool* [2] to make each IIS installation remarkably stronger.

Securing desktops and laptops is a bit more of a quagmire (as we'll discuss below). Secure servers are your responsibility as an IT manager.

2. Relying on End Users to Do Security Checking

Some basic security steps should be taken to "harden" desktops and laptops.

- Conduct regular vulnerability assessments of their hardware and software, and close the vulnerabilities that are identified.
- Deploy personal firewalls on network-connected computers.
- Perform frequent virus checking and destroy or quarantine files found to be corrupted.

Vulnerability assessment (VA) software examines a system's hardware and software, and identifies the holes exploited in active invasions. It doesn't necessarily stop invasions; it guides you in patching the security holes used by invaders. VA software asks questions like: What ports are open? What files should be encrypted? What exploitable applications are running?

Personal firewalls patrol the network traffic entering and leaving a computer.

Anti-virus (AV) detection tools watch for the byte sequences that indicate a computer virus has attached itself to a file. It takes just a few infected computers to keep a virus alive [3]. Modern computer viruses survive because they can continue to evolve.

The problem is that VA and AV software have rules and definitions that are updated frequently by their manufacturers. When new (initially vulnerable) software is shipped or when new vulnerabilities are discovered in existing software and hardware, preventative measures are soon developed to close the holes. To be effective, VA and AV software needs to run frequently (even constantly).

End users cannot be trusted to update their software or even to run it. They have other things to focus on, and they usually don't understand the risks. They'll almost never have the most recent updates. Relying on end users is a poor follow-up to a sound investment in these modern defensive software products.

You need to control and enforce security processes from a central location. See that the VA, firewall, and AV systems can be deployed, updated, and monitored centrally.

New invasions are being invented as you read this. Computers that are not checked frequently with VA and AV tools become more vulnerable to outside attack as time passes.

3. Performing Event Correlation Manually

Many computers and network devices keep logs of the events they encounter. These events are usually very specific to a certain operating system, application, or network component. When intruders attack any part of a computer network, it's likely that telltale signs are left behind, written as events to the logs of the computers on the front line of the attack.

You'd like to see these events as they occur, rather than going back to each system weeks later and dumping its system logs. Clues to what's happening in a computer network are widely available – they're just spread out all over the place. You'd like

to have the events from different systems correlated to detect broad attacks. You'd like the log data to be consolidated and synchronized so you can see what is happening where. Most important, you'd like to generate automated responses to intrusions, corresponding to the security policies you've established in your organization.

Security event correlation [4, 5] in any large organization is a huge data processing task, however, well beyond the capabilities of humans. Copious amounts of event data need to be consolidated, redundancy in the data eliminated, patterns discovered in the events, and then actions initiated to respond to what's discovered. Yet we see people continue to perform event correlation manually.

The security systems in a large organization can accumulate more than a terabyte of event data over a seven-day period. In addition, that data must be kept online for some period if the intent is to perform any forensics after an intrusion has been detected.

When you can centralize events recording and handling, you also get the ability to correlate events across an organization. For example, suppose someone attempting to crack a password moves from workstation to workstation to avoid detection. Under normal circumstances, this method wouldn't raise any alarms, because the only way to notice the moving intruder would be to look on a computer-by-computer basis. Event correlation systems can see "the big picture" by receiving events from all these locations. They can correlate these actions and detect a pattern that raises an alert, as well as initiating an automated response, such as disabling user ID for some period of time.

Even after extensive data reduction, the task of correlation and pattern matching requires a strong analysis engine. In particular, you want the pattern analysis to identify points of failure across the networks, systems, and applications. NetIQ's *Security Manager* event correlation engine is aimed at the tasks described here.

4. Having Too Many Alarms

The simplest reaction to finding something wrong is to raise an alarm – a notification that an intrusion or other anomaly has been detected. The rules defined in your detection systems obviously need to be good enough to:

a) Reduce the total number of alarms, and
b) Reduce the number of false-positive alarms.

Having too many alarms can make you as vulnerable has having none. Aesop's fable about the boy who cried "Wolf" taught us that false alarms make you distrust the systems altogether [6].

Pete Cafarchio, technology program manager at the International Computer Security Association [7], recommends that the security team not "turn on alarming features until the intrusion detection system is up and running for 30 days." By that time, they should have a better understanding of how the system works.

5. Giving Users Clean Machines

A computer's exposure to outside threats begins the moment it is booted up and attached to a network. It's tempting (and cheap) to hand your users computers that arrive fresh from the manufacturer, but these computers are extremely vulnerable. It's the job of the IT team to see that the computers are as invulnerable and as manageable as possible from the first day.

Invest in creating a secure build image, to be installed on all user computers. Never install a virgin operating system and

application suite on a computer and then hand it to a user.

Start with a clean computer, with a hard disk that has been wiped clean. Prepare the disk with the strongest file system available. Use the Encrypting File System, in conjunction with NTFS for Windows 2000 and XP, for state-of-the-art file system security. Install the latest version of the operating system that you're comfortable supporting.

Then, apply the latest patches to the operating system. Microsoft keeps the latest updates for user operating systems at windowsupdate.microsoft.com. These may change daily or weekly – consider using a service to keep up-to-date. Install your organization's applications and their latest patches.

Next, run a vulnerability assessment against that system, set at the highest level of "pickiness." It may find hundreds of potential holes in the system, usually ranked in importance from critical down to low. Spend the time to research and close the vulnerabilities it finds. This may take several days, but it's a good investment – you'll be replicating these fixes across many computers, with the intention of avoiding identifiable intrusions. Leave the VA agent on the computer so it can be updated and run remotely. NetIQ's *Security Analyzer* is well suited for this job.

Install a personal firewall on the user computer. Windows XP comes with an integrated personal firewall; personal firewalls are available from various manufacturers for most desktop and server operating systems. Ensure that the firewall is initially set in its most secure position; open any ports or change any other defaults you know will be needed.

Get the latest anti-virus software for the operating system, and make sure it's up-to-date. Run it to assure no viruses have been introduced so far, and again, leave an AV agent on the computer so it can be updated and run remotely.

This is now the hardened "gold standard," the computer to clone. Replicate this base image with each new installation.

6. Fighting Uncontrolled Fires

When a severe security breach occurs, you need to drop everything. Right then is when you must reduce the depth and breadth of the damage. Doing firefighting like this frequently is a poor way to spend your security budget.

Firefighting is costly. Everything that's productive and proactive stops, maybe throughout the entire organization. In lieu of forward progress, you try to reduce the amount that you fall back. Schedules slip; people get stressed; they lose sleep; accidents become more likely; morale declines; rotten conditions prevail. There's lots of collateral damage, including, potentially, your reputation or the reputation of the whole organization.

Develop a firefighting plan. Establish a set of processes to be followed when an intrusion does occur. Plan ahead for fighting fires, to reduce the chaos when they arrive. Hold "fire drills." Everyone on the team should have clear assignments and should be able to tell when one goal is complete, so they can move to the next step. Let every incident become a lesson on how to prevent or reduce the next fire. How can you change your security policies so it doesn't happen again?

Six Good Investments

Most investment opportunities with high leverage involve planning ahead, so you are well prepared for certain likely situations. These six investment ideas are

intended to position you for the broad range of possible breaches and lapses.

The SANS Institute maintains another excellent Web page, listing "How To Eliminate The Ten Most Critical Internet Security Threats" [8]. Microsoft offers "The Ten Immutable Laws of Security" [9]. Pass these top-ten lists to your security team.

1. Create a Well-Educated Security Team

Good computer and network security is rarely taught in schools. The reason for this omission is evident: students armed with state-of-the-art knowledge about the ins-and-outs of security would probably find the time to experiment on the tempting system in front of them: the school's. Also, both the state-of-the art and the state-of-what's-easy-to-do advance quickly in information security. So, in an area of extreme importance, every student, at every level of education, is seriously undereducated. What they've learned, they've usually learned on their own, on the street (well, probably "on the net").

The stakes are different for someone whose job now focuses on their knowledge of security. As an employer, it's important to have thorough, in-depth, ongoing education for your security team.

Those who know system security well often learned their trade from the other side of the fence, as an intruder. That's a valuable skill, but doesn't necessarily translate to the classic roles of the defender in managing computer security: prevention, detection, and reaction [10]. Members of the team need to master the processes in each of these areas.

It's important to understand user psychology. Most breaches are enabled by lapses in human engineering. Use insight into intruders to think ahead of them. What knowledge does your team need to prevent new attacks, to defend against their damage, and to react as they occur?

Intruders take advantage of vulnerabilities. Educate your team on solid processes that avoid vulnerabilities. Send them to classes, buy them books, and give them a time and place to experiment. You should insist on courses like "Windows 2000 Designing Security" and "Windows 2000 Network Security Design" For those seeking or renewing their MCSE certification.

The education shouldn't stop. Look at the weeklies; there's constant focus on the security problems with today's hottest new technologies, such as encryption, viruses, VPNs, and wireless networking.

2. Know Your Users Well

Insiders cause most security problems. Despite a huge focus on external intrusions and attacks, studies show that "about 70% of the security risk is internal - either by accident or maliciousness," says Marius Swart, GM of security solutions at Internet Solutions [11]. "However," Swart adds, "a thorough analysis of potential threats can substantially reduce [the] risk of a security breach."

Thus, minding the users at home is an essential part of analyzing threats and developing a security plan. Creating a thorough plan means establishing or revising your security policies, then putting processes in place to enforce those policies. Sophisticated software designed to coordinate and enforce security policies system-wide should be a part of any attempt to reduce internal security threats.

For example, your security policies should describe how new users are handled as they join an organization: how are they identified, and what can they access? When users change roles or jobs (or leave the area), are the access controls updated? How are contractors or temporary

workers handled when their work and contracts end?

Just as important are the policies applied to the members of the security team. For example, it doesn't make sense for a single team member to have full access to all the resources in an organization. Each team member should have an appropriate range of control, with all actions subject to ongoing auditing.

To secure your system is to know your users and administrators well, and to incorporate that knowledge into the way the system operates. Strong authentication of each user should be performed, to be certain that they really are who they say they are. Stringent access controls should be enforced. Access policies must be clearly and consistently applied: what objects is each user allowed to read, write, modify, create, or delete? Is the data they manipulate properly authenticated? Are the access controls consistent across all systems? Are changes to the validation, authentication, and access control audited? How is data integrity checked?

Carefully-monitored processes are needed to implement a thorough set of policies like these. NetIQ approaches these requirements with the products in its Security Suite. For example, NetIQ's *Directory Security Administrator* provides a central view of objects, such as files, folders, and printers, across an entire organization. It shows the objects in the directory structure and displays their associated permissions. Administrators can easily identify the assigned rights and privileges of each object, and make changes to their access control lists (ACLs) as necessary. Directory Security Administrator's "permissions search" capability lets administrators easily determine how permissions are set and manage the rights of users, groups or machines on objects throughout a domain. Directory Security Administrator helps IT personnel simplify administration by automating manual tasks and increases security by tracking and enforcing access rights.

"Microsoft's Active Directory ACL-based security model is powerful, but because it's so extensive, it can be difficult to manage," said Tom Kemp, senior vice president of products at NetIQ. Directory Security Administrator helps IT professionals get the most out of Active Directory's sophisticated security capabilities.

3. Perform Postmortems

Whenever a security incident occurs, take some time after the fact to go back and review it. How did the intrusion occur? How could it have been prevented? How could it have been detected earlier, before it did too much damage? How could the users and the security team react better, to reduce the damage? What roles weren't covered well?

Postmortems are a rich source for plans and actions in your organization. Write down the observations, and turn the lessons into something manageable, like a Top 5 or a Top 10 list. The key is that it's your list, not one from Microsoft or from a magazine. Factor the list back into budgets, job descriptions, and employee feedback reviews. Use it as a mechanism to make improvements and move your team forward – not to blame them.

4. Investigate Microsoft's .NET Framework

Early versions of Windows NT and OS/2 were touted as crash-proof, since they protected applications from stepping on each other's memory. Yet, a dozen years later, programs still crash. We've learned that memory protection was a simple first step to robust and secure computer environments.

Microsoft's .NET framework is designed to offer broader and more granular

protection. It aims to protect the wide set of objects that programs use, such as files, networks, printers, user identities, and privilege settings.

To best exploit the .NET Framework, programs should be coded to use its new application programming interfaces. For example, rather than developing and debugging their own encryption schemes, applications can use the DES or RC2 calls in the "CryptoAPI." They can authenticate users with all the richness of certificates, Kerberos, or Passport. Web applications can readily verify logon credentials or manage cookie generation.

Start your .NET education now. When applications are being upgraded, rewritten, or replaced, does it make sense to incorporate .NET features? When does it make sense to deploy servers implemented with the .NET framework? Awareness of .NET will probably soon become part of all your budget and hiring decisions; make the investment now to become fluent in its technology.

5. Use Virtual Private Networks

A virtual private network (VPN) lets you connect from one computer to another over any network. To the two computers, the network connection appears private and secure. VPN technology has matured rapidly during the past couple of years, particularly with the adoption of a group of standards known as IPSec.

IPSec technologies offer VPNs various levels of 1) authentication, verifying the identity of the two computers, 2) assurance that the data has not been modified, and 3) assurance that the data cannot be observed by anyone capturing network packets.

VPNs give you wonderful flexibility in your networking choices. For example, you might choose to run your entire set of private transactions over the public Internet, eliminating the need to maintain a private network.

Modern software tools, such as NetIQ's *Chariot*, can show you how VPN settings affect performance for different traffic types. Such tools can also identify network configuration errors or limitations of your service provider. Most importantly, they can demonstrate capacity limits: how many of your users can connect to your VPN and get reasonable performance.

The most vulnerable spots in a VPN are the two ends of the virtual connection. If you're using a VPN to connect to another organization for business purposes, consider requesting a security audit of the computers at the other end of your VPN. Take special security precautions for the computers your users use to connect from home or on the road.

If you're using Windows 2000 or XP on your laptop, desktop, or servers, you already have an excellent software-based IPSec stack shipped as part of these operating systems. Even though the price of hardware VPN servers has dropped appreciably and VPN client software is often free, you probably don't want to do this alone – there are still too many things that can go wrong. Consider using a VPN service to worry about the hard stuff (but be sure to track their performance and audit their security).

6. Exploit Microsoft's Extensive Resources

Microsoft views Windows security as key to its ongoing success. In January 2002, Bill Gates announced an intense focus by Microsoft on "trustworthy software." Their three-pronged approach is aimed at the right places: 1) making Microsoft's software less vulnerable to attacks, 2) making online transactions private and under the control of the users, and 3) making their software easier to use and more reliable.

Microsoft has an impressive area of their Web site online today; it's the place where the members of your security teams should start for information: www.microsoft.com/security/. There's another area with detailed technical information for IT professionals: www.microsoft.com/technet/security/.

Microsoft's *Security Toolkit CD* is free – the best kind of investment. It takes a few weeks to arrive, so see reference [12] and put your order in today.

Another free offering is the "Product Security Notification" [13]. This is an e-mail notification service that Microsoft uses to send information to subscribers about the security of Microsoft products. Anyone can subscribe to the service and unsubscribe at any time.

An Overall Investment Plan

You've seen some poor ways to spend your security budget and some good ways. Here are some thoughts that parallel the investment themes introduced at the beginning.

Diversify Your Investments

Spend money for the short term and for the long term. A short-term example: focus your team on dealing with the very latest virus; it's what's likely to bite you next. In the long term: educate, train, and treat your security team as if you know they'll be around for a long time.

Invest in all three stages of security management: prevention, detection, and reaction. Do some things in-house, but farm some things out to expert services teams who you can also draw upon when your resources get overextended.

It's More like Insurance

Putting money in computer security is more like buying insurance than buying market equities. You're trying to limit your downside risk. The money you put in is unlikely to show growth; rather, it's intended to avoid losses that could occur if you don't make the expenditures. Security investments rarely help revenues; however, they may affect profit by avoiding painful spikes in expenses.

Track Your Investments

Good security may be like good insurance, but that isn't to say you shouldn't see returns on the money you put into security. Good security practices can manifest themselves in the bottom line through increased availability. As security practices are implemented and refined, organizations should see the investment recouped in terms of reduced outages that are a result of security-related causes. Here are some examples:

- Fewer unauthorized accesses to sensitive files and directories.
- Fewer unauthorized remote administrative accesses to servers and workstations.
- Reduced exposure to administrators who can access all accounts.
- Less time spent restoring computers and files corrupted by viruses.

Measure and track what's going on in the security team, to help with future planning and to give the team feedback for improvement. Measure successes and failures. For example, how many security breaches occurred? In what areas? How many security-related help desk calls arrived?

Our Investment in Windows Security

Our focus at NetIQ is on the topics introduced here. We'd like to help you make wise security investments for your

organization. We're working intimately with Microsoft to build robust management software for its corporate software offerings. Microsoft licensed NetIQ technology as the core of its Microsoft Operations Manager (MOM), for managing Windows 2000 and XP operating systems and Microsoft .NET Enterprise Servers. Our companies collaborate to develop and market management solutions for servers, clients, applications, and devices.

Microsoft designated NetIQ as its premier independent software vendor (ISV) for solutions built on MOM. These solutions include management of non-Microsoft environments, management of Windows NT 4.0 operating system-based environments, and advanced security and network performance management.

To learn more about NetIQ's security products, see www.netiq.com/solutions/security/.

About the Author

John Q. Walker is a director at NetIQ Corporation. He is one of the founders of Ganymede Software, heading its software development team before it joined NetIQ in spring 2000. In earlier jobs, he managed teams responsible for designing and developing high-speed networking at IBM. Another job stint involved inventing ways to break security on a new computer system, reporting his methods to the development team. He co-authored a book on portable network programming for McGraw-Hill. Dr. Walker holds a Ph.D. in software engineering; his Master's degree focused on system testing. He can be reached at johnq@netiq.com.

Acknowledgments

Chris Cander and Scott Hollis suggested this paper. Gracious thanks to the readers who really helped improve it: Jeff Aldridge, James Coggins, Jeff Dozer, Chris Farrow, Joseph Kubik, Mike Mychalczuk, Susan Pearsall, Robby Rose, Chris Selvaggi, Carl Sommer, and Edith Sorenson.

References

1. "Mistakes People Make that Lead to Security Breaches," SANS Institute resources, www.sans.org/mistakes.htm.
2. Microsoft's IIS Lockdown Tool, www.microsoft.com/downloads/release.asp?ReleaseID=33961&area=search&ordinal=2
3. "Virus researchers: Internet needs immune system," Elizabeth Weise, *USA Today*, February 27, 2001, www.nettime.org/nettime.w3archive/200103/msg00003.html.
4. *Security Event Correlation – Where Are We Now?*, John Q. Walker, NetIQ Corporation, November 2001, www.netiq.com/Downloads/Library/white_papers/.
5. Dennis Drogseth, "Taking event correlation seriously," *Network World*, March 20, 2000, www.cnn.com/2000/TECH/computing/03/20/event.corr.idg/.
6. "The Shepherd's Boy and The Wolf," Aesop's fable number 43.
7. Ruttrell Yasin, "Security Mandate: Silence False Alarms," *Internet Week*, April 8, 1999, www.internetwk.com/story/INW19990408S0009.
8. "How to Eliminate the Ten Most Critical Internet Security Threats," SANS Institute resources, www.sans.org/topten.htm.

9. "The Ten Immutable Laws of Security," Microsoft Corporation, www.microsoft.com/technet/columns/security/10imlaws.asp.
10. *Secrets and Lies: Digital Security in a Networked World*, Bruce Schneier, John Wiley & Sons; ISBN: 0471253111, August 2000.
11. "The enemy within: Don't overlook internal security," Marius Swart, *ITWeb*, July 18, 2001, www.itweb.co.za/sections/techforum/2001/0107180819.asp.
12. Microsoft Security Toolkit CD availability: www.microsoft.com/security/mstpp.asp.
13. Microsoft Product Security Notification: www.microsoft.com/technet/security/bulletin/notify.asp.

Copyright Information