COMMUNICATING WITH SAM

Network Security: Where Are the Real Threats? Part II

Sam Johnston

Question:

Over the past year, it seems that high profile "hackers" have been in the news on a regular basis. These security breaches, many of which have impacted some of the world's largest companies, have created a real awareness within our senior management of the need to properly protect our information. We have taken measures to protect the high-risk points. Our Internet connection is protected by a firewall, and we have intrusion detection software to monitor our AS/400 for host security breaches. We also regularly force our users to change their various network and system access passwords, and we randomly do physical security checks to ensure that users do not remain logged onto the network or systems while away from their desk. Despite these actions, senior management is concerned that we may not be taking all the security measures necessary to protect our valuable data. Based on your experiences, what points in the network represent risk, and what can be done to reduce security risks?

Answer:

Security is perhaps the most discussed, and to a certain extent the most misunderstood, information technology discipline that we currently deal with. It is not uncommon for information technology trends to reflect the views commonly expressed in the mainstream business press.

However, as an IT professional you need to be aware that by the time a topic is a regular item in the popular press, you better have a sound strategy for dealing with issue, because playing catch-up when you are under the senior management microscope can be a painful experience.

Your question, due to the sensitivity of the issue, is extremely complex, but believe it or not, the answer you need is based upon a simple principle. The key to unlocking the answer, and demystifying the solution, is rooted in the fact that in this connected world of interlocking dependencies there is a need to protect all valuable data and network resources. More simply put, if you can abandon the temptation of protecting individual devices or hosts, and start accepting that security is a network-wide need that demands a system embedded in the utility of the network, then you will be successfully prepared to fight the battle.

In Part I of this series, we published Cisco's list of the top fourteen security vulnerabilities, and they are worth review as they emphasize the importance of focusing on a network security model. They are:

1. Inadequate router access control: misconfigured router ACLs can allow information leakage through ICMP, IP, NetBIOS and lead to unauthorized access to services on your DMZ servers.
2. Unsecured or unmonitored remote access points are one of the easiest means of access to your corporate network.
3. Information leakage can provide the attacker with operating system and application versions, user groups, shares, DNS information zone via zone transfers, and running services like SNMP, finger, SNMP telnet rusers, sunrp, NetBIOS.
4. Hosts running unnecessary services (such as sunpc, FTP, DNS, SMTP) leave ways in.
5. Weak, easily guessed and reused passwords at the workstation level can doom your servers to compromise.
6. User or test accounts with excessive privileges.
7. Misconfigured Internet servers, especially CGI scripts on Web servers and anonymous FTP.
8. Misconfigured firewalls or router ACL can allow access to internal systems directly or once a DMZ server is compromised.
9. Software that is unpatched, outdated, vulnerable, or left in the default configurations.
10. Excessive file and directory access controls (NT/95 shares, UNIX NFS exports).
11. Excessive trust relationships such as NT Domain Trusts and UNIX .rhosts and hosts .equiv files can provide hackers with unauthorized access to sensitive systems.
12. Unauthenticated services like X Windows.

13. Inadequate logging, monitoring, and detection capabilities at the network and host level.
14. Lack of accepted and well promulgated security policies, procedures, guidelines and minimum baseline standards.

Source: Cisco Systems, "Make Your Network Safe for E-Business"

The interesting link among the top fourteen security vulnerabilities is the fact that most can be resolved through simple adherence to good IT practices. Despite the fact that the popular press paints a picture of a network of sophisticated hackers that are capable of cracking the Pentagon using black magic techniques to attack your data, the reality is that digital security is no different than physical security. More often than not, the criminal is able to make the victim vulnerable because the victim makes it easy for them!

So now we know where and how we will be threatened, how do we architect a security infrastructure to protect our valuable assets?

As we evolve to a network security model, the world of physical and digital security, as we have alluded to earlier, start to draw many parallels. In both cases, the appropriate security architecture is driven by security policies that reflect the unique nature of your business. A security policy, whether it is designed to protect physical or digital assets, has a few key things in common.

The policy starts with an assessment of the risk profile of the business you are in, and the inherent risks of what you do. This is weighed against your risk tolerance as an organization, and the cost of reducing your risk profile to zero or nearly zero. This may or may not alter what services you plan to offer, if the cost of protecting the assets at an acceptable risk level generates a cost structure that delivers on unacceptable ROI. The answers to these questions will lead to a policy framework that simplifies the process of creating the suitable network security architecture.

If we draw on principles from physical security, we will come to understand the value of a network versus system or host-based security architecture. When we only protect the host or source of value, it is like a bank relying entirely on the vault for security. There is a high probability that the vault will resist the intruder for a long enough duration to permit the police to ward off the robbery, but think of the damage that can be caused while the robber has control over the bank. As we know, banks do not rely entirely on the vault, but rather have layers of security that ensure that intruders can be detected and isolated before gaining access to the vault. The goal of network security is no different – detect and isolate the intruders on the perimeter where you can limit and mitigate the potential damage.

If we look at the physical security that most of us have implemented within our offices and premises, we see some common layers to the architecture:

- ***Identification:*** the need for employees to carry identification cards for access
- ***Perimeter Security:*** use of secure entry points to control who can enter the premise
- ***Secure Transport:*** the use of armoured cars and secure escorts to protect valuables in transport
- ***Security Monitoring:*** security guards and cameras to watch secure entry points for intrusion
- ***Security Management:*** the security command post to review event logs and ensure policy enforcement

Digital network security is no different, in that there is a need to address each one of these layers, with options at each layer that vary in investment, complexity and protection levels. Obviously this is where security policy provides a guiding hand in determining which element is the appropriate choice for your business. The scope of this article does not permit a detailed review of the various technologies, so I have tried to provide a general blue print of how we each layer could be addressed.

Identification:

Good digital security starts with identification, or the triple-A server, which will provide Authentication of who you are, Authorization to access resources and Accounting information on transactions for audit purposes. Using technologies such as RADIUS or TACACS+ (Terminal Access Controller Access System Plus), the AAA server typically controls user access from the outside to internal resources.

Other identification techniques include digital certificates whereby an algorithm dynamically assigns a password to the user's digital certificate to authenticate the user's public key.

These systems also provide the benefit of in effect creating a single password for all access, and as the system automatically and dynamically assigns the password based on the algorithm, it eliminates password administration and the risk of password sharing.

Perimeter Security:

At the perimeter, protection is achieved by firewalls that ensure that only acceptable traffic passes through the front door, and by using network address translation they conserve and hide addresses. Firewalls also protect the network from DoS attacks.

Although we think of firewalls at Internet entry points, they can also be crucial to creating intermediate perimeters within the Intranet, remembering again the scope of internal threats. Thus, the firewall plays the role of protection from both external and internal intrusions.

Secure Transport:

In addition to enabling Web access, firewalls also permit VPN site-to-site and client connectivity over the Internet to back office resources previously only reached via private WAN connections. This is done through technologies such as IP SEC, and encryption technologies, which use algorithms to cipher data into unintelligible formats for secure transport over the Internet.

Security Monitoring:

We have already briefly noted the need to use intrusion detection systems (IDSs) to monitor hosts and network resources for attacks. There are two forms of IDS. The one noted in your question is referred to as host-based IDS, or HIDS, which focuses on application or host specific attacks and can intercept OS and application calls on a specific host to actually prevent attacks.

The other form is network IDS, or NIDS, which looks over the entire network, but is more passive than HIDS in that NIDS issues an alert upon discovery of an attack. Think of HIDS as the bank vault, and the NIDS as the security check points in the bank that lead to the vault. In other words, you need both, and like the bank analogy, there are various checkpoints where NIDS should be deployed. →

The obvious placement points for NIDS is at the external perimeter, such as the Internet connection point, and in the DMZ where semi-trusted traffic is permitted.

However, recalling that internal threats, where security enforcement is lowest, represent the largest financial risk, it is important to not to forget about NIDS on the LAN segments of your network.

The key in deploying NIDS is tuning the configuration to ensure that it is mitigating legitimate attacks, and not creating a flood of either "false-positives" (alarms generated by legitimate traffic) or "false-negatives" (attacks that are missed). For NIDS mitigation, there are two options to consider.

First, there is the use of "shunning" traffic, whereby NIDS can block a host from coming into the network for a predetermined time if it detects an attack from that particular host over a particular protocol. The problem with "shunning" is that address spoofing can allow a hacker to create a DoS for a legitimate host.

The second NIDS option is TCP resets, which terminate an attack by sending a TCP reset message to the attacking and attacked host. This is just a simple overview of NIDS, and before implementing the technology it is crucial to perform a detailed review of your environment to ensure that you understand how to tune the implementation to optimize security, while minimizing the flood of false alerts.

Security Management:

The security control centre is crucial to ensuring that your investment is not wasted. All the devices mentioned above create logs and information that needs to be captured centrally. The challenge you will face is the volume of data that will be generated. You will need resources dedicated to security management, and you will need to invest in tools and applications that will allow you to sort out which logs are important. The security control centre also needs to be the control point for secure access to devices for configuration, as configuration settings will be important to ensuring that your security needs are met.

Security management plays the crucial role of evaluating results versus the policy guidelines so that you can determine whether your objectives have been met. Remember that implementing network security is the simple part – ongoing management of the systems is where the real investment lies, and where the real return will be derived.

Although I have tried to give you a sense of the underlying technology that drives network security, don't be too worried if it is seems overly complex. The key message that you need to understand is that the technology is meaningless unless you have a security policy that is a reflection of your business needs.

We can draw on physical security practices to create the architectural blue print necessary from a network security implementation perspective, and this will simplify the task of selecting the right technology. However, without sound policies, the likelihood of selecting the right technology will be extremely low. TUG

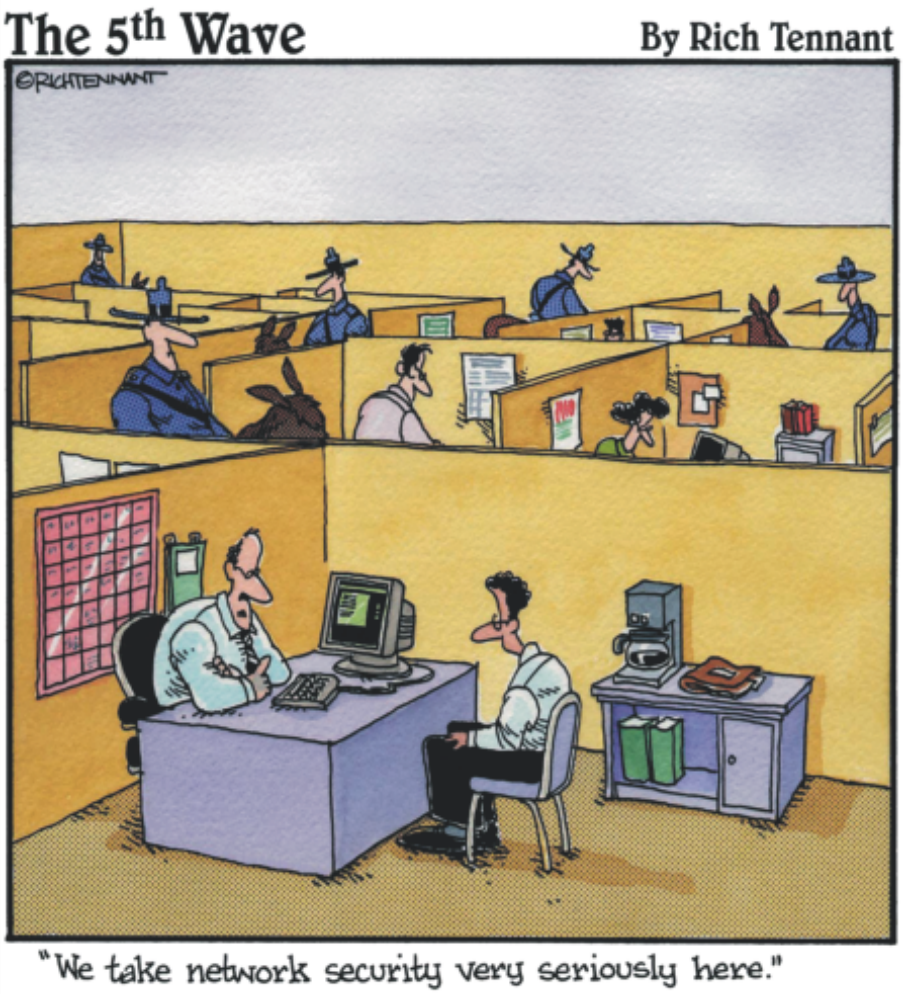

Sam Johnston is a partner and Chief Technology Officer of Intesys Network Communications Ltd., providing value-added networking and e-commerce solutions to theAS/400 community. He can be reached at (416) 438-0002 or via e-mail at sjohnston@intesys-ncl.com.

Any TUG member wishing to submit a question to Sam can forward their typewritten material to the TUG office, or to Intesys. The deadline for our next issue is Friday June 8th, 2001.

The most cost effective management solution for your AS/400!

SNAPSHOT/400

The ultimate real-time AS/400 graphical interface

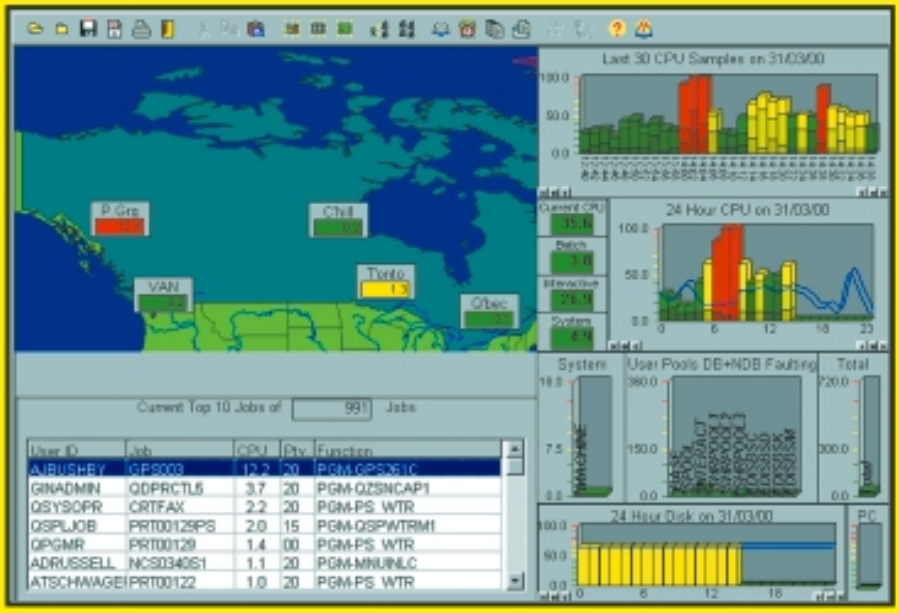

Imagine what you can see at-a-glance, without entering a single command...

- what is happening on the system - right now
- how the system has been running the past hour
- how it is running compared to how it normally runs
- how many jobs are running on the system
- what the Top-10 jobs are
- how the memory is being utilized
- what is the current capacity of the available disk
- identifies & manages hold ups & bottle-necks
- what is the network-related component of the end-user response time – ***very powerful!***

The answers are right in front of you ....updated every 60 seconds

...and many Canadian companies agree with us...

"Not only has Snapshot/400 automated system monitoring & reporting, it has saved us from purchasing other tools because it is so versatile & flexible"

Monica Harper
Parmalat Canada

parmalat

"Snapshot/400 has helped us in maintaining our operations at a high level in a complex system management environment. Snapshot/400 has been the perfect tool."

Andrei Centea
Canadelle

Sara Lee Branded Apparel Canada

"Snapshot is the only tool we use to monitor the performance of our AS/400. We have found nothing that provides the same reliability & flexibility."

Ken Lloyd
Highland Valley Copper

Highland Valley Copper

"The problem with most Performance Monitoring today is they are 'after the event' solutions; but not Snapshot/400."

Darrin Casey- IBM Global Services Australia

IBM

"Your product is great. I would recommend Snapshot/400 to anyone looking for a real time system monitoring and management tool"

Bill Shaw –
Amkor Technologies - USA

Amkor Technology

"With Snapshot/400 projected on the wall, management can see we have total control. The Capacity Planner was so easy to use compared to IBM and Best/1"

Nick Cummins
Parker Hannifin - England

Parker

See our booth (#F10) at the TEC Showcase, April 24

SNAPSHOT/400

Call Toll Free 1-800-767-5495
www.snapshot400.com