



Security Whitepaper

The challenge.

Today's challenge is to create secure, well-behaved networks that offer services only to authorized users, and that provide the right services to those authorized users. Networks that are immune to hacks, attacks, and disruptions, both planned and unplanned (like network loops). These are necessary to support critical services like voice, and to keep private information private.

That's the goal, anyway. We can't think of a single network administrator we know that doesn't have this aspiration. Yet most of today's networks are not prepared for these challenges. Why is this so? One reason is that securing networks is somehow a larger task that can be bitten off by most administrators. It's something that has to be carved out of an already full day. Where does one start? Which of the hundreds of possible products to use? How to get whatever you do decide to buy to work together? How much will this cost, in money but more importantly, in extra technical time and management attention, when time and management attention is already in short supply?

These are very legitimate reasons for delaying the implementation of such a project, but whistling past the graveyard will probably not be an option for very much longer. For sure it'll be an immediate top project with the first big security breach or major network interruption. If you start before the

crisis you can be the master of the process. If you start after the crisis you'll be working in tune with someone else's schedule.

If you're still on your own schedule, you have the luxury of dividing whatever you do into small bites. Projects are like watermelons, best consumed in manageable bites. Or, to use a metaphor from your desk, piles.

Putting order to the process: Pile One.

You can divide the security pile into two smaller piles: the pile representing network self-correction and self-defense and the somewhat larger and more involved pile involving network access control, network protection, and reporting.

Let's start, as anyone would, with the smaller pile. Network self-defense is important for a number of reasons:

- First, you probably already have most of what you need here, and you merely have to energize it and integrate it to turn it on. So there's not a big PO in the cards.
- Second, you make this defense automatic, which means that you get a lot of protection for little to no ongoing management attention. Computers—actually, ASICs inside the switches—take care of the work and tell you when they have identified and defeated some problem.
- Third, when this is done your network's stable and reliable

enough to support critical services like emergency notification, voice over IP, paging and video instruction. It'll also be self-managed, which will give you more leisure time to contemplate covers for your additional security holes.

Most late-model network equipment has built-in protection against denial of service attacks, forged BPDU attacks, ARP broadcasts from unauthorized sources, and other hacks and attacks. This helps the switches themselves defend against attacks that would take them out of service. It also helps them to block attacks that pass through them. You should install these types of switches at critical network points, like MDF's, network cores (computer rooms), and at the connection to the Internet, partner networks, and any other foreign networks. You just install the switches and reap the benefits. No configuration or ongoing management is necessary.

Some late-model switches also have automatic loop protect, which senses network loops and automatically disconnects them from the network. Loops are a sort of unknowing attack (or, if the person knows something about network technology, a sneaky means to interrupt services). At any rate, you need to consider installing switches that support this feature in all new installations and upgrades. Some manufacturers' switches—like Hewlett-Packard—have retrofitted this capability into some of their existing switches. So check for this capability and if it's available through a software upgrade, by all means make that upgrade.

Another automatic setting is a proprietary feature of Hewlett-Packard switches, and available on their higher-end, layer 3 switches. This is Automatic Virus Throttling (AVT), which is

somewhat of a misnomer since the switch doesn't look for virus signatures. A switch with AVT enabled examines all its ports all the time and reacts to sudden and drastic changes in transmission rates. When it picks up one that exceeds the pre-set sensitivity level it throttles down the traffic or turns off the port. You get to set the sensitivity level: low, medium or high. This is an excellent thing to employ on the edges of the networks, since each port supports only one device, and aberrant behavior is easier to see. You can employ it on switches at the distribution and core levels, but it loses some of its effectiveness, since downstream ports serve many more users than one, and bad behavior on one end station may be masked in the aggregate transmissions of many. Still, this should be set on all switches that have the capability; just set the sensitivity to low on distribution and core switches to cope with massive attacks.

These last three are very low-level configurations that can run automatically without manager intervention—set and forget, if you even have to set them in the first place.

There's another facet of network self-defense, and that's reacting to threats on the network by packet inspection. Here there's a separate software application somewhere that collects and analyzes packets and sends alerts and/or makes disconnection decisions automatically, and independently of network management. This is still a set-(configure) and-forget feature, but it's a bit more complicated than the first three. Let's discuss NetFlow and its standards-based derivative, SFlow.

NetFlow was one of Cisco's many great ideas. How about designing a system whereby network equipment or a manager of that equipment self-analyzes its own transmissions and

takes automated action? It's a higher-level variant of AVT, because instead of just looking at relative traffic levels it examines the contents of packets to pick out problems like viruses, worms, hacks, and unauthorized activity. NetFlow has a couple drawbacks, though—it is very processor-intensive and requires Layer 3 to work. This limits it to high-end switches that have the horsepower to process their regular traffic and also process NetFlow. Other network equipment manufacturers, like Enterasys, also use NetFlow in their security offerings.

SFlow is the standards-based variant of NetFlow, developed by Hewlett-Packard. SFlow has the same purpose, but instead of capturing all packets it proscribes a random sampling of packets. The theory is that a random packet sample will over a short time give the same signature as a 100 percent packet capture. It surely is less taxing on switch processing power, which means that even edge switches can be fitted out with SFlow and capture packets at the very edge of the networks. SFlow can also be run on Layer 2 switches, so these edge switches can be inexpensive—inexpensive enough to install them everywhere.

What to do with those packets once they are produced and saved? You need some sort of engine to collect them, analyze them, and take action on them. SFlow collectors have been out there for a while, but they were the only products of companies that were set up to produce just collectors and consoles. This made them expensive. Network equipment manufacturers are now beginning to offer their own products, and since they're made to help sell more network equipment, they are more reasonably priced. These products collect and analyze SFlow packets, and

some can take automated action on SFlow alerts.

How might this work? Some switch somewhere on the network—let's say, at the edge of the network in some minor wiring closet—randomly collects packets from every conversation and forwards those packets to a collector and analyzer. The collector/analyzer application reads the packets from a certain port on the switch and determines that some dangerous activity, or even a greatly increased usage is happening on that port. The collector would then tell the network management system to take action, which would invariably be sending alerts to the human network managers, and perhaps taking action against the originator of the activity in question. The extent of the activity depends on the capability of that edge switch, but in the least it could involve turning the network port off, and keeping it off for some predetermined length of time. It might involve switching the port into a different VLAN. It might involve learning the MAC address of the originator and then freezing that MAC address off the network (MAC address lockout). It might involve restricting a particular type of traffic to the port to solve a network congestion problem (e.g., you can't download videos so prodigiously while other important work is going on).

The point here is that the network is continually assessing its traffic, and reacting to threshold violations with automated action that's applied consistently and promptly. This will increase network reliability and stability. And when you have the switches that support SFlow somewhere in your network, this SFlow collector and analyzer is a very inexpensive add-on.

These improvements can help take your networks to the utility stage—where they can be expected to run continuously,

and under self-control, even through unintended and malicious happenings. Once again, these capabilities generally come with your network equipment and can be used as is, or with the purchase of an inexpensive bit of software. If you have these capabilities embedded in some or all of your network equipment, you should turn them on.

Putting order to the process: Pile Two.

This larger pile is somewhat more difficult, and has several components: network access control at both the Internet connection and at the edge of the network, protection against hacks, virus and worms, traffic control, workstation scanning and remediation, and auditing and reporting.

We'll start simple and move up to as complicated as you will ever want to get.

Network access control (or prevention) for accesses from the Internet into your network is easy to do. This can literally be a one-appliance solution. What's important to understand that this one appliance is NOT just a firewall. Firewalls offered satisfactory protection five years ago, but since that time applications, hackers and thieves have perfected ways to bypass the relatively crude Layer 3 protections offered by firewalls. You need to move upscale with the bad guys and invest in an IPS. This can be put behind the firewall, or if you invest in a combination appliance, it can also replace the firewall. You can invest in one of these combination units—called UTM, or Unified Threat Management, units for less than \$10,000 in most cases. Some of these units also have built-in web content filters, virus and worm filters, and VPN termination engines. It's now possible to install a single appliance that can do all of your core network-to-Internet filtering and control, and consolidate all of those various different relationships,

different management screens, and different support agreements—for a pretty significant savings in time and money.

This allows you to protect the network from pretty sophisticated threats and exploitations. It also allows you to restrict outbound traffic from your network to the Internet to just those applications and accesses that you want. It also allows you to shape and prioritize outbound traffic. This generally has enough favorable impact to eliminate the need for your next Internet connection upgrade, so if you are planning an Internet connection upgrade one of those these units will pay for itself.

Cisco offers their ASA line of enhanced firewalls with this type of additional protection. Juniper Networks, Fortinet, Packeteer, SonicWall and other vendors also offer these devices, with widely varying ranges of costs and capabilities. Lightspeed Systems has a product called Total Traffic Control that also falls into this category, except that it also has a client agent that does anti-virus checking. In Lightspeed's case the appliance is a regular PC, which may have some throughput issues, particularly on larger Internet pipes.

The safest network is one that's never accessed. No chance of errant or malicious behavior! That's not practical, of course. People have to access the networks, and every access increases the chance of risk and security compromise. We've talked about checking bad behavior at the core of the network, where it leaks into the Internet or other foreign networks. Now we need to talk about control at the other end, where people attach to your wired and wireless networks. Two issues: making sure that only authorized people actually get on your network, and then making sure that when they access the network

they do so with “clean” machines that will not infect or disturb the rest of the network.

Network access control gets easier and easier. If you have wrestled with an early RADIUS system you’ll be gratified to know that it’s much changed since then, and in a good direction. Microsoft and Novell embed RADIUS into their server offerings now, and if you don’t have RADIUS there you can pick it up for free for or a negligible cost elsewhere. Our favorite: a low-cost, pre-configured RADIUS appliance. You have historically needed software at both ends—on the PC, too, to do an 802.1x challenge-authentication, with the network switch in the middle arbitrating network access. The newer lines of switches and wireless access points now support WEBAUTH, which is a proxy system—the network equipment holds the client supplicant and only passes on a web page to the client attempting to access the network. So you don’t have to maintain any software on the client to run a network access system. This is important for several reasons: first, it’s a lot less work than a conventional RADIUS security system (nice for you), second, you can authenticate PC’s to the network that you don’t control—like guest users, and third, you can authenticate any device with a browser, like PDA’s and those new ultra-mobile laptops. These latter devices may or may not have 802.1x supplicants. If you need more rigorous and ultra-secure network access, there’s always MAC authentication and standard 802.1x authentication with certificates to fall back on.

An effective access control system can make sure that all devices attaching to your networks go to only those areas that they are supposed to go, and see only the resources that they are supposed to see. These are normally kept apart by the flexible application of

VLAN’s (Layer 2-level control), although in some cases—those with higher-end, layer 3 equipment at the network edge—they are done by setting Access Control Lists (ACL’s) directly on the edge network port. These ACL’s are generally tied to some existing security database like Active Directory or EDirectory or some other LDAP-compliant database.

An intelligently-imposed network access control system can also divide your network into virtual parts that can be joined only at shared resources while running separately otherwise, even though it’s actually only the same wire (or fiber). This has the benefit of isolating segments so as to limit damages from a hack or breach on one of them. Keeping the blast zones small, as it were.

Such a network access control system can be established quickly and easily, if you have network equipment that supports WEBAUTH. Two or three days is generally enough time to implement a workable system, if you have established all of the VLANs that you will need and have that back-end security database up and running.

If you don’t have WEBAUTH in your network equipment and don’t want to implement MAC authentication or 802.1x you can invest in a third-party access control system that overlays practically any type of network equipment. Examples of these are Cisco Clean Access, Bradford Networks’ Campus Manager, and Still Secure’s Safe Access. These systems are designed to do more than just access control, however. They also do scanning of PC’s to make sure that they have allowable configurations and the presence of working anti-virus software. Most of them also have the ability to get their users to self-remediate, or fix their

own problems, if trouble is detected by the scanning tool.

These systems were much in demand a few years back, when newly-written worms and viruses took networks down for months on end. Since that time the evident problems of worms and viruses have largely gone away. Today's worms and viruses are not written to take down networks, but to harvest information off of them. Today's hacker goal is mostly not to disrupt networks, but use them for identity and information theft, so keeping the networks up and running well is also now a hacker goal! So, with the overt problems gone away there's been less pressure to do one of these systems (they are expensive and hard to run), and sales of this type of product have languished.

But if your goal is network access control, particularly over a mixed-vendor network, and also some sort of proactive (touch the workstation) utility, you have a number of options to choose from.

The problem with these systems—besides their management burden—is that unless they have permanent agents they don't enforce behavior—they merely periodically scan for the presence of installed and working anti-virus and anti-worm software, and some specified configuration, and then keep the PC off of the network if it's caught short during a scan. What the users do in between scans is not detected. Still, some regimen of periodic scanning is better than no control at all, and so these systems have had their places in the market. Most of these are found at colleges and universities, which have to accommodate a significant number of non-controlled personal computers and give them general access to the "heartbeat" applications. The benefits of some configuration control outweigh the management burden and the periodic delay in accessing networked

applications caused by scanning and remediation.

There's a new class of software that does client control and anti-virus and anti-worm checking. These client agents live on the PC at all times and check the PC to make sure that no virus or worm infests them, and that the configuration of the PC is not changed. They also can enforce behavior, keeping PC's from attaching to unauthorized sites, even if they are taken away to other places and use other (less secure) Internet connections. These agents can't be uninstalled or turned off. This is a much stronger way to lock down the PC than by doing periodic scanning, since it's an agent that's always on when the PC is on. In fact, the cost of a combination agent that does this level of absolute control and also does virus and worm scanning and deletion generally costs less than just the cost of a standard virus and worm agent from Trend Micro, Symantec, Panda or Norton. If you're running a one-to-one program in education you should be investigating this type of agent.

Pile Three: Additional protections.

This third pile involves making your communications and access to your resources even more secure. They are pertinent in many situations.

Passwords are great, but password theft is not. Until you hire only people who can memorize complex passwords instantly and completely, and then never write them down you'll run risks of password theft. Those passwords can generally be found taped to monitors and placed under keyboards, or (more secure here) installed inside desk drawers. Someone who steals a password can get into restricted resources, or onto restricted network

segments. You can do a few things to prevent password misuse.

The first one is to associate login with PC MAC address, so the login sequence has to be used at a particular PC. No longer can that user name and password be taken to a computer lab somewhere and used to hack into the system surreptitiously. This is frequently a capability of network equipment manufacturers' network access control applications, so if you're using one of these it pays to investigate it, and then if present implement it.

The second one is to put location and time parameters on system access: access can only be done from specified locations at specified times. For example, access to the financial system can't be done at any time from any dorm, and access to the same system can't be done from any academic building after, say 5 PM. This access control is normally enforced in the network access control (RADIUS) system, working with switches that are capable of doing this type of restrictive control.

A third—and more secure—way to protect user names and passwords is to implement some sort of two-factor authentication based on biometrics. This is not as convoluted as it once was, and our favorite one uses keystroke rhythm to make sure that only the password holder is getting into the password-protected system. You may be required to implement such a password control system in the future. Remember here that it's not necessarily onerous or complex.

Encryption of some types of traffic may be required, even if placed on a secured network. That's because you don't know if an authorized user is tapping into and recording conversations, passwords, and data output (ever hear

of Cain and Abel? Google it, download it from CNet, install and run it, and then fear it! and know that many of your users can do exactly what you just did). Some types of attempts to harvest information can be identified by the serving network switches and alarms sent to the network managers, but not in all cases. It makes sense, therefore, to protect the most sensitive data being sent even on internal networks with encryption. Some devices, like voice over IP equipment, can automatically encrypt their transmissions. You may find it useful to run encrypted VPN's inside your networks for some PC users. All that takes is that relatively inexpensive PC client that also does anti-virus and worm scanning, personal firewall and web access filtering.

It's said that any direction is OK when you don't know where you are going. And that's true in the world of network security; particularly if you don't know where you are. When you implement security systems you need information on how the network's doing. Most of the security products referenced in this document come with built-in data logging and reporting utilities, with varying degrees of completeness and polish. You must have some feedback and analysis loop always going, so that you can see where your remaining security gaps lie, identify problem users, and fine-tune the tradeoffs between security and access. All of this information should be fed back into the system as corrective inputs.

We have talked about a lot of stuff here. But remember that it all doesn't have to be implemented to improve what security you already have, or even to assure adequate security. One or two items may be sufficient. And don't forget that you may have different classes of users with different security requirements and different consequences. Targeting security to

where it's most needed may be a way of economizing on time, effort and money, while still getting what you need.

Bringing it all home.

So what would we implement if we had the luxury of a clean slate and enough money to do the job? We'd implement as much as possible that could be set to run and cure problems automatically. We'd try to use as much of what we already had to accomplish this purpose, and invest judiciously where we needed to. From bottom to top, here's what we'd do.

- Get network switches and wireless network equipment that supports SFlow, and perhaps AVT
- Get switches that support automatic loop protection
- Get network switches and wireless network equipment that supports WEBAUTH
- Get one UTM at the core of the network, and one smaller UTM at the edge of each local network that needed traffic grooming before hitting the WAN
- Implement a RADIUS-based authentication system and use WEBAUTH as the low-complexity front-end
- Turn off network segments when they are not supposed to be serving people
- Install permanent client agents that do anti-virus and anti-worm scanning as well as enforce behavior
- Install a recording and reporting device to monitor the performance of these security systems and report on needed changes

Or, we'd move along this hierarchy of improvements until the money on hand ran out, and prepare network

improvement requests for as much additional protection as we thought necessary. We'd also divide our user population into security groups with different requirements and lock down the high-risk groups first. That's another way of dividing your mass of security requirements into manageable segments. If your guests, students, or contract employees don't need high security and if you can wall them off from the other, more critical segments, you can cover them inexpensively.

In the long run we'd expect to save money on the consolidation of the various appliances at the Internet connection. We'd probably also spend a little less money on these new combination anti-virus PC control agents than we'd otherwise spend on regular anti-virus agents. We would probably also spend less money on Internet bandwidth upgrades. But money savings, while nice, is not the main thing. We'd expect to shift our work time away from fixing network problems and responding to virus and worm attacks to other, more proactive work. Remember that handling unforeseen and therefore unscheduled problems is the least effective type of work. And we'd expect to be more proofed against an embarrassing data breach or problems with accessing unwanted Internet content.

Remember that in the field of network security, any place is a good place to start, and if you can't do everything, you ought to do what you can. Any progress to the level of a secure, stable network is a good thing.

About MXN Corporation. *MXN is a company that specializes in the design and implementation of networks and networked services like video and voice. One of our competencies is the design and implementation of network security systems, and we have a growing practice in helping organizations cope with their current security challenges.*