Computers & Security, 20 (2001) 491-503

Managed Security Monitoring:Network Security for the 21st Century

Bruce Schneier

The Internet is critical to business. Companies have no choice but to connect their internal networks to the rest of the world—to link with customers, suppliers, partners, and their own employees. But with that connection comes new threats: malicious hackers, criminals, industrial spies. These network predators regularly steal corporate assets and intellectual property, cause service breaks and system failures, sully corporate brands, and frighten customers. Unless companies can successfully navigate around these, they will not be able to unlock the full business potential of the Internet.

Traditional approaches to computer security don't work. Despite decades of research, and hundreds of available security products, the Internet has steadily become more dangerous. The increased complexity of the Internet and its applications, the rush to put more services and people on the Internet, and the desire to interconnect everything all contribute to the increased insecurity of the digital world.

Security based on products is inherently fragile. Newly discovered attacks, the proliferation of attack tools, and flaws in the products themselves all result in a network becoming vulnerable at random (and increasingly frequent) intervals.

Real security is about people. On the day you're attacked, it doesn't matter how your network is configured, what kind of boxes you have, or how many security devices you've installed. What matters is who is defending you. The only way to stay ahead of new vulnerabilities and attacks is through detection and response. In the real world, this translates to alarm systems and guards. On the Internet, this means active network monitoring. You need the best possible people defending you at the moment you're attacked. You need instant detection and effective response. To get that, you need to be monitored by the best Managed Security Monitoring (MSM) company.

Security monitoring is a key component missing in most networks. Monitoring provides immediate feedback regarding the efficacy of a network's security—in real time, as it changes in the face of new attacks, new threats, software updates, and reconfigurations. Monitoring is the window into a network's security; without it, a security administrator is flying blind. Wherever a network is in the process of building security, monitoring is critical. It's the first thing you need to do.

0167-4048/01$20.00

The Importance of Security

When I began working in computer security, the only interest was from the military and a few scattered privacy advocates. The Internet has changed all that. The promise of the Internet is to be a mirror of society. Everything we do in the real world, we want to do on the Internet: conduct private conversations, keep personal papers, sign letters and contracts, speak anonymously, rely on the integrity of information, gamble, vote, publish digital documents. All of these things require security. Computer security is a fundamental enabling technology of the Internet; it's what transforms the Internet from an academic curiosity into a serious business tool. The limits of security are the limits of the Internet. And no business or person is without these security needs.

The risks are real. Everyone talks about the direct risks: theft of trade secrets, customer information, money. People also talk about the productivity losses due to computer security problems. What's the loss to a company if its e-mail goes down for two days? Or if ten people have to scramble to clean up after a particularly nasty intrusion? I've seen figures as high as $10 billion quoted for worldwide losses due to the ILOVEYOU virus; most of that is due to these productivity losses.

More important are the indirect risks: loss of customers, damage to brand, loss of goodwill. Regardless of how the million-credit-card-number theft at Egghead.com turned out, some percentage of customers decided to shop elsewhere. When CD Universe suffered a credit card theft in early 2000, it cost them dearly in their war for market share against Amazon.com and CDNow. In the aftermath of the Microsoft attack in October 2000, the company spent much more money and effort containing the public relations problem than fixing the security problem. The public perception that their source code was untainted was much more important than any effects of the actual attack.

And more indirect risks are coming. European countries have strict privacy laws; companies can be held liable if they do not take steps to protect the privacy of their customers. The U.S. has similar laws in particular industries — banking and healthcare — and there are bills in Congress to protect privacy more generally. We have not yet seen shareholder lawsuits against companies that failed to adequately secure their networks and suffered the consequences, but they're coming. Can company officers be held personally liable if they fail to provide for network security? The courts will be deciding this question in the next few years.

As risky as the Internet is, companies have no choice but to be there. The lures of new markets, new customers, new revenue sources, and new business models are just so great that companies will flock to the Internet regardless of the risks. There is no alternative. This, more than anything else, is why computer security is so important.

The Failure of Traditional Security

Five years ago, network security was relatively simple. No one had heard of denial-of-service attacks shutting down Web servers, common gateway interface scripting flaws, and the latest vulnerabilities in Microsoft Outlook Express. In recent years came intrusion detection systems, public-key infrastructure, smart cards, VPNs, and biometrics. New networking services, wireless devices, and the latest products regularly turn network security upside down. There are literally hundreds of network security products you can buy, and they all claim to provide you with security. They regularly fail, but still you hear CEOs say: "Of course I'm secure. I bought a firewall."

Network security is an arms race, and the attackers have all the advantages. First, network defenders occupy what military strategists call: 'the position of the interior'; the defender has to defend against every possible attack, while the attacker only has to find one weakness. Second, the immense complexity of modern networks makes them impossible to properly secure. And third, skilled attackers can encapsulate their attacks in software, allowing people with no skill to use them. It's no wonder CIOs can't keep up with the threat.

What's amazing is that no one else can either. Computer security is a 40-year-old discipline; every year there's new research, new technologies, new products, even new laws. Every year things get worse.

If there's anything computer security professionals have learned about the Internet, it's that security is relative. Nothing is foolproof. What's secure today may be insecure tomorrow. Even companies like Microsoft can get hacked, badly. There are no silver bullets. The way forward is not more products, but better processes. We have to stop looking for the magic preventive technology that will avoid the threats, and embrace processes that will help us manage the risks.

Security and Risk Management

Ask any network administrator what he needs security for and he will describe the threats: Web site defacements, corruption and loss of data due to network penetrations, denial-of-service attacks, viruses and Trojans. The list seems endless, and the endless slew of news stories prove that the threats are real.

Ask that same network administrator how security technologies help, and he'll discuss avoiding the threats. This is the traditional paradigm of computer security, born out of a computer science mentality: figure out what the threats are, and build technologies to avoid them. The conceit is that technologies can somehow "solve" computer security, and the end result is a security program that becomes an expense and a barrier to business. How many times has the security officer said: "You can't do that; it would be insecure?"

This paradigm is wrong. Security is a people problem, not a technology problem. There is no computer security product—or even a suite of products—that acts as magical security dust, imbuing a network with the property of "secure." It can't be done. It's not the way business works.

Businesses manage risks. They manage all sorts of risks; network security is just another one. And there are many different ways to manage risks. The ones you choose in a particular situation depend on the details of that situation. And failures happen regularly; many businesses manage their risks improperly, pay for their mistakes, and then soldier on. Businesses are remarkably resilient.

To take a concrete example, consider a physical store and the risk of shoplifting. Most grocery stores accept the risk as a cost of doing business. Clothing stores might put tags on all their garments and sensors at the doorways; they mitigate the risk with a technology. A jewellery store might mitigate the risk through procedures: all merchandise stays locked up, customers are not allowed to handle anything unattended and so on. That same jewellery store will carry theft insurance, another risk management tool.

More security isn't always better. You could improve the security of a bank by strip-searching everyone who walks through the front door. But if you did this, you would have no business. Studies show that most shoplifting at department stores occurs in dressing rooms. You could improve security by removing the dressing rooms, but the losses in sales would more than make up for the decrease in shoplifting. What all of these businesses are looking for is adequate security at a reasonable cost. This is what we need on the Internet as well—security that allows a company to offer new services, to expand into new markets, and to attract and retain new customers. And the particular computer security solutions they choose depend on who they are and what they are doing.

Protection, Detection, and Response

Most computer security is sold as a prophylactic: encryption prevents eavesdropping, firewalls prevent unauthorized network access, PKI prevents impersonation. To the world at large, this is a strange marketing strategy. A door lock is never sold with the slogan: "This lock prevents burglaries." No one ever asks to purchase: "a device that will prevent murder." But computer security products are sold that way all the time. Companies regularly try to buy: "a device that prevents hacking." This is no more possible than an anti-murder device.

When you buy a safe, it comes with a rating. 30TL—30 minutes, tools. 60TRTL—60 minutes, torch and tools. What this means is that a professional safecracker, with safecracking tools and an oxyacetylene torch, will break open the safe in an hour. If an alarm doesn't sound and guards don't come running within that hour, the safe is worthless. The safe buys you time; you have to spend it wisely.

Real-world security includes prevention, detection, and response. If the prevention mechanisms were perfect, you wouldn't need detection and response. But no prevention mechanism is perfect. This is especially true for computer networks. All software products have security bugs, most network devices are misconfigured, and users make all sorts of mistakes. Without detection and response, the prevention mechanisms only have limited value. They're fragile. And detection and response are not only more cost effective, but also more effective, than piling on more prevention.

On the Internet, this translates to monitoring. In October 2000, Microsoft discovered that an attacker had penetrated their corporate network weeks before, and might have viewed or even altered the source code for some of their products. Administrators discovered this breach when they noticed twenty new accounts being created on a server. Then they went back through their network's audit logs and pieced together how the attacker got in and what he did. If someone had been monitoring those audit logs—automatically generated by the firewalls, servers, routers, and so on, in real time, the attacker could have been detected and repelled at the point of entry.

That's real security. It doesn't matter how the attacker gets in, or what he is doing. If there are enough motion sensors, electric eyes, and pressure plates in your house, you'll catch the burglar regardless of how he got in. If you are monitoring your network carefully enough, you'll catch a hacker regardless of what vulnerability he exploited to gain access. And if you can respond quickly and effectively, you can repel the attacker before he does any damage. Good detection and response can make up for imperfect prevention.

Prevention systems are never perfect. No bank will ever say: "Our safe is so good, we don't need an alarm system." No museum will ever say: "Our door and window locks are so good, we don't need night watchmen." Detection and response are how we get security in the real world, and they're the only way we can possibly get security on the Internet. CIOs must invest in network monitoring services if they are to properly manage the risks associated with their network infrastructure.

Monitoring Network Security

Monitoring the security of a building implies several things. It implies a series of sensors in and around the building. It implies a central alarm that rings if the sensors are tripped. And it implies some kind of response to the alarm. (If there's no response, the alarm is useless. Think of a car alarm going off in the inner city; if no one notices or cares, it has no value).

Network monitoring implies several similar things. It implies a series of sensors in and around the network. Luckily, these are already in place. Every firewall produces a continuous stream of audit messages. So does every router and server. IDSs send messages when they notice something. Every other security product generates alarms in some way.

But these sensors by themselves do not offer security. You have to assume that the attacker is in full possession of the specifications for these sensors, is well aware of their deficiencies, and has tailored his attack accordingly. He may even have passwords that let him masquerade as a legitimate user. Only another human has a chance of detecting some anomalous behaviour that gives him away.

Building alarms detect simple things: a broken window, an intruder walking down a forbidden hallway. Building alarm companies are really nothing more than glorified paging services; they notice alarms and alert the right people. Unfortunately, network monitoring is much more complicated.

The first step is intelligent alert. Network attacks can be much more subtle than a broken window. Much

depends on context. Software can filter the tens of megabytes of audit information a medium-sized network can generate in a day, but software is too easy for an attacker to fool. Intelligent alert requires people. People to analyze what the software finds suspicious, and to delve deeper into suspicious events, determining what is really going on. People to separate false alarms from real attacks. People who understand context.

By itself, an alert is only marginally useful. More important is to know how to respond. This is the second step of good network monitoring. Every attack has a response. It could be as simple as shutting off a particular IP address to repel an attacker. It could be as drastic as taking a corporate network off the Internet. Again, people are the key. Software can only provide generic information; real understanding requires experts.

Finally, the response must be integrated with the business needs of the organization. Security engineers only see half the information. They understand attacks and their security significance, but they don't understand the business ramifications. A large e-business might keep its Web site up and running even if it is being attacked; preventing the loss of revenue may be more important than the site's immediate security. On the other hand, a law firm may have the exact opposite response; the sanctity of its customers' data might be more important than having its Web site available.

This is detection and response as applied to computer networks. Network devices produce megabytes of audit information daily. Automatic search tools sift through those megabytes, looking for telltale signs of attacks. Expert analysts examine those signs, understanding what they mean and determining how to respond. The owner of the network, the organization, makes security decisions based on ongoing business concerns.

To make network monitoring work, people are needed every step of the way. Software is just too easy to fool. It doesn't think, doesn't question, doesn't adapt. Without people, computer security software is just a static defence. Marry software with experts, and you have a whole different level of security.

Outsourcing Monitoring

The key to a successful detection and response system is vigilance: attacks can happen at any time of the day and any day of the year. While it is possible for companies to build detection and response services for their own networks, it's rarely cost-effective. Staffing for security expertise 24 hours a day and 365 days a year requires five full-time employees; more, if you include supervisors and back-up personnel with more specialized skills. Even if an organization could find the budget for all of these people, it would be very difficult to hire them in today's job market.

Retaining them would be even harder. Security monitoring is inherently bursty: six weeks of boredom followed by eight hours of panic, then seven weeks of boredom followed by six hours of panic. Attacks against a single organization don't happen often enough to keep a team of this calibre engaged and interested.

In the real world, this kind of expertise is always outsourced. It's the only cost-effective way to satisfy the requirements. I may only need a doctor twice in the coming year, but when I need one I may need him immediately. I may need specialists. Out of a hundred possible specialities, I may need two of them—and I have no idea beforehand which ones. I would never consider hiring a team of doctors to wait around until I happen to get sick. I outsource my medical needs to my clinic, my emergency room, my hospital. Similarly, a network needs to outsource its security monitoring to an MSM service.

Aside from the aggregation of expertise, an MSM service has other economies of scale. It can more easily hire and train its personnel, simply because it needs more of them. And it can build an infrastructure to support them. Vigilant monitoring means keeping up to date on new vulnerabilities, new hacker tools, new security products, and new software releases. An MSM service can spread these costs among all of its customers.

An MSM provider also has a much broader view of the Internet. It can learn from attacks against one

customer, and use that knowledge to protect all of its customers. And, from its point of view, attacks are frequent. There's a reason you don't have your own fire department, even if you can afford one. When the fire department comes to your house, you want them to have practiced on the rest of the neighbourhood. To an MSM company, network attacks are everyday occurrences; as experts, they know exactly how to respond to any given attack, because in all likelihood they have already seen the same attack many times before.

In the real world, security is always outsourced. Every building hires another company to put guards in its lobby. Every bank hires another company to drive its money around town. Security is important, complex, and distasteful; it is smarter to outsource than to do it yourself.

Monitoring First

You have a safe in a dilapidated building, and you need to secure it. What's the first thing you do? Inventory the safe? Assess the security of the building? Install better locks on the doors and bars on the windows? Probably not. The first thing you do, as quickly as possible, is put an alarm on the safe and post a guard. Once the safe is being monitored, you can then afford the time and attention needed to inventory the stock, analyze the environment, and improve the security. Without monitoring, you're vulnerable until your security is perfect. If you monitor first, you're immediately more secure.

In network security, people have been doing this backwards. Companies see monitoring as something to do after their security products are in place. First they develop a security policy. Then they do a vulnerability analysis. Then they install a firewall, and maybe an intrusion detection system. And finally they think about monitoring. This makes no sense.

Monitoring should be the first step in any network security plan. It's something that a network administrator can do today to provide immediate value. Policy analysis and vulnerability assessments take time, and don't actually improve a network's security until they're acted upon. Installing security products improves security, but only if they are installed correctly and in the right places. How does a CIO know what products to install, and whether they are actually working—as installed, not as they worked in the lab? The only way he can know is to monitor the network. Monitoring ensures that security products are providing the type of security they were intended to provide.

This kind of thinking is especially important in dynamic environments like company networks. The network changes every day: new applications, new servers, new vulnerabilities. A CIO can go to sleep one night confident that his network is secure, and wake up the following morning to read about a major vulnerability in the newspaper. Suddenly his network is wide open, even though nothing changed. A CIO can reconfigure his network to increase productivity, or add a new network service, or simply upgrade a software package, and suddenly the security of his environment is completely different. Networks are extremely complex—nonlinear and tightly coupled—and it's impossible to predict how different subsystems interact. How does he know the security ramifications of what he does? The only way is to monitor security.

It's specious logic for a CIO to decide to wait until his network is stable, he understands his security, and all his patches are up to date. It'll never happen. Monitoring's best value is when a network is in flux—as all large networks always are—due to internal and external factors. Monitoring provides immediate security in a way that neither doing a vulnerability assessment nor dropping a firewall into a network never can provide. Monitoring provides dynamic security in a way that yet another security product can never provide. And as security products are added into a network—firewalls, IDSs, specialized security devices—monitoring only gets better.

In engineering, control theory is based on the concept of monitoring. An engineer might want to be able to tune his factory: "How can I control this plastic film extruder to ensure a uniform thickness of plastic?" This is a real question, and a complicated one. The plastics extruder might have a dozen different

dials controlling things like temperature, pressure, and speed. You can adjust the amount and force of the air being blown, the amount of plastic bead material in the machine, or how rapidly the film is pulled out of the machine. All of these controls affect the thickness of the plastic. What you really want is to turn a dial marked "thickness" to a setting labelled "4 mil plastic." But since each dial affects the others, and can even cause time-dependent feedback loops, it's not nearly that simple. So what do you do? You monitor the system; not just at the output, but internally. Then, based on what you've observed, you establish feedback loops to create a closed-loop system (I am ignoring a library's worth of advanced mathematics here,) and apply the mathematics of control theory to get what you want. It might take hundreds of pages of analysis, but that's how control theory works. But first you need to monitor so you know what's going on. Monitor, and you gain control of the system.

Security is no different. Monitoring is what gives companies a window into their security. Did you install a firewall? An IDS? Why? Did it increase security or not? Did you configure it right? Did you install it at the right place in your network? How do you know? Monitoring is the only way you can really know. Once you know, you can start making changes. If you make changes without monitoring, you're just guessing.

Monitoring is the feedback loop that makes all the other network security activities more effective. It's how you determine where to install security devices, and whether or not they're doing any good. It's how you know if your security devices are configured correctly. It's how you ensure that your security doesn't degrade over time. It needs to be done first.

Counterpane's Monitoring

The fundamental truism behind Counterpane's value to its clients is that all security products can be significantly improved by real-time vigilant monitoring. Moreover, it is not enough to have an automatic software routine watch for anomalies and sound an alarm. Any response must include an expert response, including detailed information on how to repel the intruder and fix the problem. Automatic programs are plagued by false positives that substantially reduce their value; Counterpane relies on expert human analysts to separate the real attacks from the false positives, and to contact the customers and provide them with expert assistance in stopping the attack and closing the vulnerability. To us, the network architecture and security products are just the terrain we fight on.

Counterpane's value resides primarily in our analysts. They are experts in network security, experienced from defending against countless attacks against customer networks, and reinforced by a fast escalation procedure that can call in the world's best experts when required. They are vigilant, constantly watching customer networks for signs of attack. They are adaptive, able to see new attacks that automatic products miss, and respond to them. And they are relentless in their pursuit of security.

Counterpane's analysts are aided by several processes. First, there is the process that collects information from the client's networks and presents it in a form that the analysts can efficiently understand and analyze; this is our Sentry. Second, there is the process by which the network information is matched with a diagnosis or series of diagnoses; this is the Socrates system. Third, there is the process within Socrates by which the network information and diagnosis are matched with client-specific information about the seriousness of different sorts of network events. And fourth, there is the process by which both the Sentry and Socrates are updated with new information about attacks, vulnerabilities, products, etc. This is our Network Intelligence function. We constantly scour security databases and research institutions, and even the hacker underground, looking for new attacks, new vulnerabilities, and new hacker tools, modifying Socrates to take this information into account.

It's not just smart people and smart processes make Counterpane's MSM work; it's the way we create our analysts. We have exacting hiring requirements. We run background checks and perform rigorous personal screening. Our training program resembles boot camp: it's arduous, it's demanding, and not everyone makes it. Experienced analysts mentor new analysts.

In the end, there is an esprit de corps that will do everything possible to defend the customer. And in the end, Counterpane is hired because companies want that esprit de corps defending them.

Counterpane's analysts work in our SOCs. We currently have two SOCs: a primary facility in Chantilly, VA, and a secondary facility in Mountain View, CA. Having multiple SOCs is key to guaranteeing uninterrupted service. These are both physically hardened facilities: entrance requires a password, a biometric, and a physical access token. The SOCs are under constant surveillance—audio, video, and clickstream—to provide an unambiguous audit trail. In fact, the two SOCs monitor each other, and can force a switchover if there is a problem. The internal network is physically isolated from the general Internet, but distributed among the SOCs. None of the workstations can create, or read, removable media. This level of security may seem excessive, but to us it is vital to ensure the security of our customers' networks.

The eyes and ears of Counterpane's MSM service are the sensors on a network. Some sensors are already in place, and can immediately provide information to Counterpane. These are routers, servers, firewalls, and other network devices. Other sensors are security-specific, and must be specially installed on a network: IDSs, honeypots and other Internet burglar alarms, content-protection systems, and HTML-based security devices.

Counterpane's unique security value comes from correlating information from divergent sensors. Often attacks are not obvious from a single log entry, but only from a pattern of activity. Because Counterpane monitors all sensors on a company's network, it can see things that no single sensor can. This is a critical difference between Counterpane and other monitoring companies: we don't just watch a firewall and an IDS, we don't just watch a limited menu of security products, we watch everything on a customer's network.

A company could, in theory, attempt to replicate Counterpane's monitoring in their own network. This would require at least five people to staff a single monitoring station 24x7. It would require additional experts for specific products and problems, a team to monitor the security newsgroups and hacker underground, more people to constantly train and drill the team, and a secure facility to house all of this. And even then, the in-house group would not be as qualified as Counterpane. They would not have the experience of monitoring hundreds of different customers. They would not have Counterpane's training. They would not have access to Counterpane's Network Intelligence, or its Socrates system. They would not have access to Counterpane's experts. When a company is attacked, it needs the best possible people working to defend it. Counterpane's MSM service can protect a network better than doing it yourself, and at a fraction of the cost.

Conclusion

Network security risks will always be with us. The downside of being in a highly connected network is that we are all connected with the best and worst of society. Security products will not "solve" the problems of Internet security, any more than they "solve" the security problems in the real world. The best we can do is to manage the risks: employ technological and procedural mitigation while at the same time allowing businesses to thrive.

Computer security equals vigilance, a day-to-day process. It's been thousands of years, and the world still isn't a safe place. And no matter how fast technology advances, alarms and security services are still state-of-the-art.

The key to effective security is human intervention. Automatic security is necessarily flawed. Smart attackers bypass the security, and new attacks fool products. People are needed to recognize, and respond to, new attacks and new threats. It's a simple matter of regaining a balance of power: human minds are the attackers, so human minds need to be the defenders as well.

On the day you're attacked, you want the best possible defence. It's not enough to give the job to

your overworked system administrators or in-house security staff. You need the best, and you need them immediately. This is why you hire an MSM company, because they can defend you better.

MSM combines people, processes, and products to create a security environment for the chaos of modern business networks. The reality of today's Internet makes MSM the most cost-effective way to provide resilient security, and Counterpane Internet Security, Inc. is the MSM industry leader.

Anyone can install a security product, carry a pager, and pretend to monitor a network; Counterpane marries the best people and the best processes to provide complete security, resilient security, a level of security unmatched anywhere.

Monitoring in the news

Monitoring and Resilient Security

During the course of the year 2000, several groups of Eastern European hackers broke into at least 40 companies' Web sites, stole credit card numbers, and in some cases tried to extort money from their victims. The network vulnerabilities exploited by these criminals were known, and patches that closed them were available—but none of the companies had installed them. In January 2001, the Ramen worm targeted known vulnerabilities in several versions of Red Hat Linux. None of the thousands of infected systems had their patches up to date. In October 2000, Microsoft was molested by unknown hackers who wandered unchallenged through their network, accessing intellectual property, for weeks or months. According to reports, the attackers would not have been able to break in if Microsoft's patches had been up to date. The series of high-profile credit card thefts in January 2000, including the CD Universe incident, were also the result of uninstalled patches. A patch issued eighteen months previously would have protected these companies.

What's going on here? Isn't anyone installing security patches anymore? Doesn't anyone care?

What's going on is that there are just too damn many patches. It's simply impossible to keep up. I get weekly summaries of new vulnerabilities and patches. One alert service listed 19 new patches in a variety of products in the first week of March 2001. That was an average week. Some of the listings affected my network, and many of them did not. Microsoft Outlook alone had over a dozen security patches in the year 2000. I don't know how the average user can possibly install them all; he'd never get anything else done.

Security professionals are quick to blame system administrators who don't install every patch: "They should have updated their systems; it's their own fault when they get hacked." This is beginning to feel a lot like blaming the victim: "He should have known not to walk down that deserted street; it's his own fault he was mugged."; "She should never have dressed that provocatively; it's her own fault she was attacked." Perhaps such precautions should have been taken, but the real blame lies elsewhere.

Those who manage computer networks are people too, and people don't always do the smartest thing. They know they're supposed to install all patches. But sometimes they can't take critical systems off-line. Sometimes they don't have the staffing available to patch every system on their network. Sometimes applying a patch breaks something else on their network. I think it's time the industry realized that expecting the patch process to improve network security just doesn't work.

Security based on patches is inherently fragile. Any large network is going to have hundreds of vulnerabilities. If there's a vulnerability in your system, you can be attacked successfully and there's nothing you can do about it. Even if you manage to install every patch you know about, what about the vulnerabilities that haven't been patched yet? (That same alert service listed 10 new vulnerabilities for which there are no patches). Or the vulnerabilities discovered but not reported yet? Or the ones still undiscovered?

Good security is resilient. It's resilient to user errors. It's resilient to network changes. And it's resilient to

administrators not installing every patch. Managed Security Monitoring is an important part of that resilience. Monitoring makes a network less dependent on keeping patches up to date; it's a process that provides security even in the face of ever-present vulnerabilities, uninstalled patches, and imperfect products.

In a perfect world, systems would rarely need security patches. The few patches they did need would automatically download, be easy to install, and always work. But we don't live in a perfect world. Network administrators are busy people, and networks are constantly changing. Vigilant monitoring is by no means a panacea, but it is a much more realistic way of providing resilient security.

Managed Security Monitoring vs. Managed Security Services

There has been a lot of confusion in the press about managed security. There are Managed Security Services (MSS), Managed Security Monitoring (MSM), managed firewalls, and managed IDSs. Dozens of companies are offering different services, and they're not even naming them consistently. Here's my quick guide to the differences:

MSS is basically the firewall and IDS equivalent of managed PBX. Years ago, outsourcers sprang up to manage a company's PBX telephone switches. What they did was known as: 'moves, adds, and changes.' They would add phone numbers, move people from one extension to another, set up voice mail...that sort of thing. Today, MSS companies do the same thing for firewalls, IDSs, and VPNs. They will set up the devices, configure them, update them as required, change their rulesets if required, etc. Sometimes there is a monitoring component to MSS—checking whether the device is working and whether there are any problems—but it's often no more than a guy with a pager. The customer still manages the process; there's not much outsourced management in MSS.

MSM is a very similar term, but a very different experience. MSM is a true outsourced service, but it's outsourced security monitoring. Counterpane doesn't just monitor individual products; we monitor a company's entire network. This difference is critical. One of the most serious problems with network security today is that it is fragile. Every security device has vulnerabilities, blind spots, and problems—and new problems are discovered every week. MSM does a far better job at security monitoring than MSS can ever do, simply because it has a broader view of a company's security. By monitoring a company's entire network, MSM provides resilient security. And by choosing an outsourcer focused on monitoring and security instead of management, you get more expertise and more attention to response.

MSM is a set of services centred around monitoring. At Counterpane, our analysts are experts at detecting network attacks and providing timely response. Our Socrates system, our Sentry, and our SOCs are all optimized to monitor network security. Our Network Intelligence experts continually enhance our monitoring, in a world where security never stands still. MSS companies don't do all of this; they may do pieces of it, but it's half-baked. When you think vigilent, adaptive, and relentless, you don't think management, you think expert monitoring. Managed Security Monitoring is its own thing.

MSM is about providing expertise to the customer when there is a problem. MSS is more focused on helping the customer when the customer asks. At Counterpane, we don't wait until the customer contacts us. We monitor the customer's network and alert the appropriate people when there are problems. We can find problems that a firewall or an IDS misses. An MSS company can't possibly do that.

MSS is a commodity business; many companies provide it, and they don't charge very much. MSM is a whole different level of security. Vigilant monitoring of a customer's entire network provides a level of protection that MSS just can't match.

Intrusion Detection Systems and Monitoring

Recently I've been seeing several articles foretelling the death of Intrusion Detection Systems (IDS). Supposedly, changes in the way networks work will make them an obsolete relic of simpler times. While I agree that the challenges IDSs face are serious, and that they will always have limitations, I am more optimistic about their future.

IDSs are the network equivalent of virus scanners. IDSs look at network traffic, or processes running on hosts, for signs of attack. If they see one, they sound an alarm. In Secrets and Lies, I spent several pages on IDSs (pp. 194–197): how they work, how they fail, the problems of false alarms. Suffice it to say that the two problems IDSs have are: 1)Failing to detect real attacks; and 2)Failing to ignore false alarms. Smart administrators regularly ignore "chatty" IDS systems because they are getting so many false positive messages. And hacker tools specifically designed to bypass IDSs are common.

These two problems are nothing new, but several recent developments threaten to undermine IDSs completely.

First is the rise of IPsec. IPsec is a security protocol that encrypts IP traffic. An IDS can't detect what it can't understand, and is useless against encrypted network traffic. (Similarly, an anti-virus program can't find viruses in encrypted e-mail attachments.) As encryption becomes more widespread on a network, an IDS becomes less useful.

Second is the emergence of Unicode. In the July 2000 Crypto-Gram, I talked about security problems associated with Unicode. One problem is the ability to disguise character strings in various ways. Since most IDSs look for character strings in packets indicating certain network attacks, Unicode threatens to make this job unfeasible.

Third is the increased distribution of networks. Today's traffic isn't just coming through one firewall, but also through hundreds of different direct external links to customers, suppliers, joint venture partners, outsourcing companies, IPsec gateways for telecommuters and road warriors, etc.

This makes it very hard to monitor the traffic.

And fourth is the sheer speed of networks. For an IDS to be effective, it has to examine every packet. This slows down an Ethernet software switch or router, but completely stalls a gigabit hardware device. Data transmission rates are getting so fast that no IDS can possibly keep up.

Some security experts are predicting the death of IDSs, but I don't agree. Even with all of the drawbacks, an IDS is still the most effective tool for detecting certain network attacks. But it is not a panacea. I think of IDSs as network sensors, similar to a burglar alarm on a house. It won't detect every attack against the house, it can be bypassed by a sufficiently skilled burglar, but it is an effective security countermeasure.

And just as door and window alarms are more effective when combined with motion sensors and electric eyes, IDSs are more effective when combined with other network sensors. Tripwire, for example, is a network sensor that alarms if critical files are modified. Honeypots include network sensors that alarm if attacked.

The missing piece is a way to interpret and respond to these alarms. The whole point of building Counterpane was to deal with the problem of these sensors going off. Someone has to watch these sensors 24x7. Someone has to correlate information from a variety of sensors, and figure out what's a false alarm and what's real. Someone has to know how to respond, and to coach the network administrator through the process. I think MSM will finally make IDSs look good.

Military History and Network Security

Military strategists call it "the position of the interior." The defender has to defend against every possible attack. The attacker, on the other hand, only has to choose one attack, and he can concentrate his forces

on that one attack. This puts the attacker at a natural advantage.

Despite this, in almost every age of warfare the attacker is at a disadvantage. More people are required to attack a city (or castle, or house, or foxhole) than are required to defend it. The ratios change over history—the defence's enormous advantage in WWI trench warfare lessened with the advent of the tank, for example—but the basic truth remains: all other things being equal, the military defender has a considerable advantage over the attacker.

This has never been true on the Internet. There, the attacker has an advantage. He can choose when and how to attack. He knows what particular products the defender is using (or even if he doesn't, it is usually one of a small handful of possibilities). The defender is forced to constantly upgrade his system to eliminate new vulnerabilities and watch every possible attack, and he can still get whacked when an attacker tries something new, or exploits a new weakness that can't easily be patched. The position of the interior is a difficult position indeed.

A student of military history might be tempted to look at the Internet and wonder: "What is it about warfare in the real world that aids the defender, and can it apply to network security?"

Good question.

The defender's military advantage comes from two broad strengths: the ability to quickly react to an attack, and the ability to control the terrain.

The first strength is probably the most important; a defender can more quickly shift forces to resupply existing forces, shore up defence where it is needed, and counterattack. Here we see the same themes from elsewhere in this booklet: how detection and response are critical, the need for trained experts to quickly analyze and react to attacks, and the importance of vigilance. I've built Counterpane's MSM service around these very principles, precisely because it can dramatically shift the balance from attacker to defender.

The defender's second strength also gives him a strong advantage. He has better knowledge of the terrain: where the good hiding places are, where the mountain passes are, how to sneak through the caves. He can modify the terrain: building castles or SAM batteries, digging trenches or tunnels, erecting guard towers or pillboxes. He can choose the terrain on which to stand and defend: behind the stone wall, atop the hill, on the far side of the bridge, in the dense jungle. The defender can use terrain to his maximum advantage; the attacker is stuck with whatever terrain he is forced to traverse.

On the Internet, this second advantage is one that network defenders seldom take advantage of: knowledge of the network. The network administrator knows exactly how his network is built (or, at least, he should), what it is supposed to do, and how it is supposed to do it. Any attacker except a knowledgeable insider has no choice but to stumble around, trying this and that, trying to figure out what's where and who's connected to whom. And it's about time we exploited this advantage.

Think about burglar alarms. The reason they work is that the attacker doesn't know they're there. He might successfully bypass a door lock, or sneak in through a second-story window, but he doesn't know that there is a pressure plate under this particular rug, or an electric eye across this particular doorway. MacGyver-like antics aside, any burglar wandering through a building wired with alarms is guaranteed to trip something sooner or later.

Traditional computer security has been static: install a firewall, configure a PKI, add access-control measures, and you're done. Real security is dynamic. The defence has to be continuously vigilant, always ready for the attack. The defence has to be able to detect attacks quickly, before serious damage is done. The defence has to be able to respond to attacks effectively, repelling the attacker and restoring order.

This kind of defence is possible in computer networks. It starts with effective sensors: firewalls,

well-audited servers and routers, intrusion-detection products, network burglar alarms. But it also includes people: trained security experts that can quickly separate the false alarms from the real attacks, and who know how to respond. It includes a MSM service. This is security through process. This is security that recognizes that human intelligence is vital for a strong defence, and that automatic software programs just don't cut it.

It's a military axiom that eventually a determined attacker can defeat any static defence. In World War II, the British flew out to engage the Luftwaffe, in contrast to the French who waited to meet the Wehrmacht at the Maginot Line. The ability to react quickly to an attack, and intimate knowledge of the terrain: these are the advantages the position of the interior brings. A good general knows how to take advantage of them. We must leverage them effectively for computer security.