SEARCHMIDMARKETSECURITY.COM

technical guide on NETWORK SECURITY MONITORING

contents

5 Starting Points for Network Monitoring

8 How to Maintain Network Control Plane Security

12 Network-Based Integrity Monitoring Keeps Website Hacks in Check

16 Validate Your Perimeter Network Security Devices Are Working

19 Buying an IPS: Determine Why You Need Intrusion Prevention

22 Preparing for a Network Security Audit Starts with Monitoring and Remediation

TechTarget

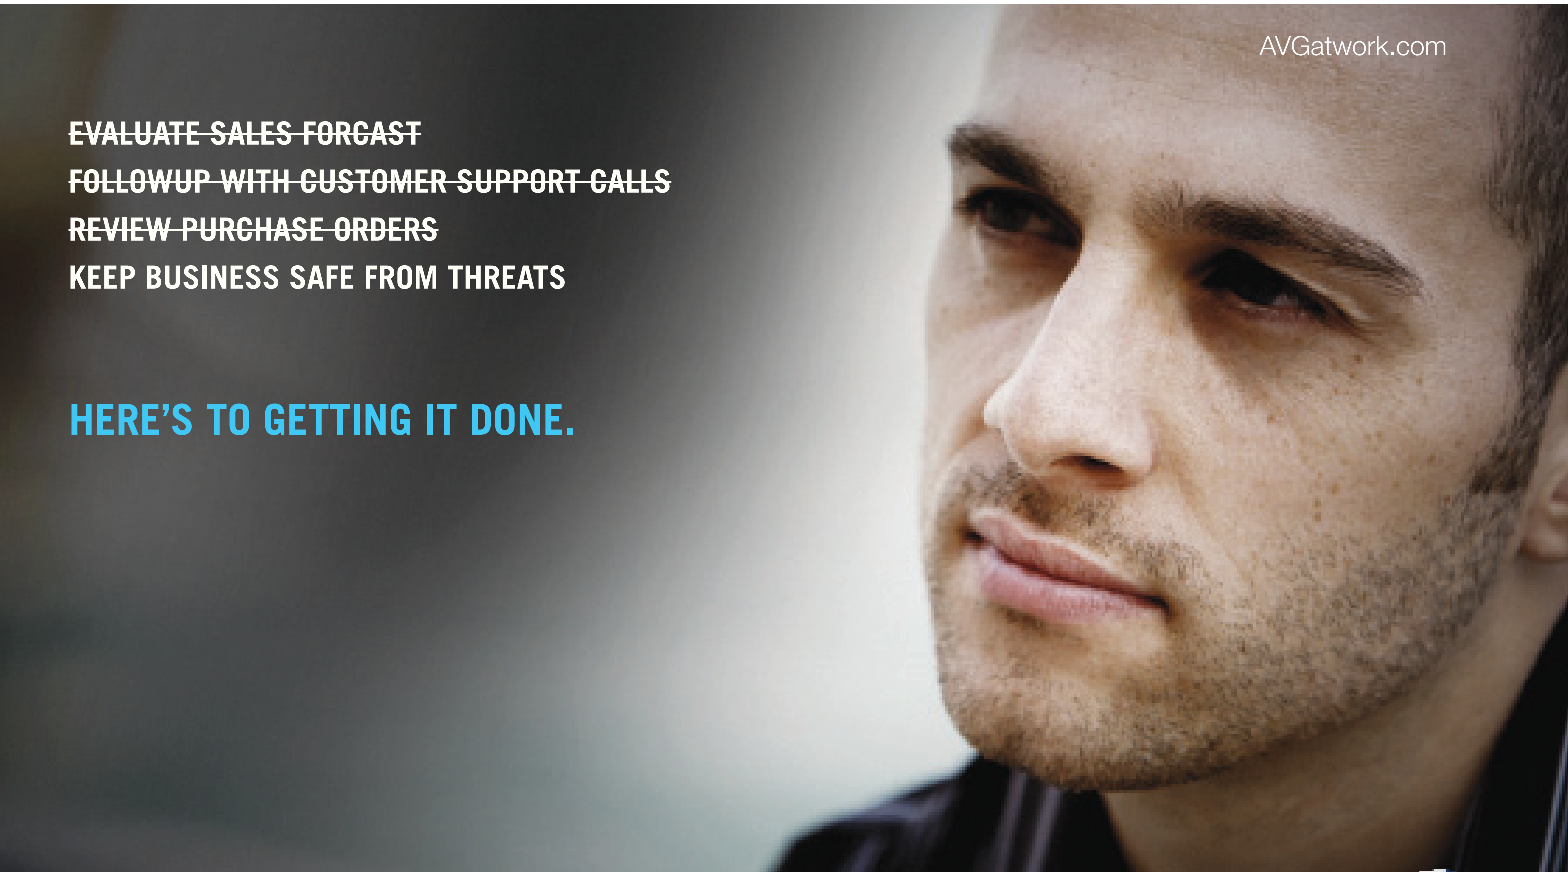
AVGatwork.com
~~EVALUATE SALES FORCAST~~
~~FOLLOWUP WITH CUSTOMER SUPPORT CALLS~~
~~REVIEW PURCHASE ORDERS~~
KEEP BUSINESS SAFE FROM THREATS
HERE'S TO GETTING IT DONE.

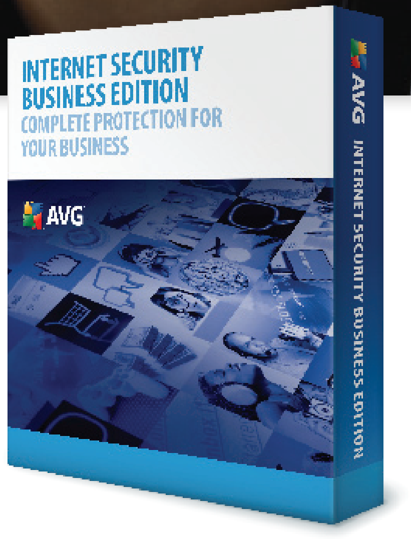
INTERNET SECURITY
BUSINESS EDITION
COMPLETE PROTECTION FOR YOUR BUSINESS
AVG
AVG INTERNET SECURITY BUSINESS EDITION

insight

Network Security Monitoring

Your organization's network is more than a collection of pipes through which traffic flows: it's your business' fundamental IT architecture. Your security devices keep a watchful eye on traffic and systems, and maintain the integrity of your data and systems.

SEARCHMIDMARKETSECURITY.COM presents a comprehensive guide to network security monitoring. Our experts cover all the angles with authoritative technical advice on: network security management; integrity monitoring; security device testing; intrusion prevention buying criteria; and how to prepare for your annual network audit.

contents

5 **Starting Points for Network Monitoring**
AUTOMATED MONITORING *When looking at automated network monitoring tools, your first considerations should be uptime and reachability of systems.* BY JOEL SNYDER

8 **How to Maintain Network Control Plane Security**
NETWORK SECURITY MANAGEMENT *Use access control lists and secure configurations to maintain the security of your organization network control plane.* BY JOEL SNYDER

12 **Network-Based Integrity Monitoring Keeps Website Hacks in Check**
INTEGRITY MONITORING *Network integrity monitoring brings the concept of file-based integrity monitoring to a company's online presence.* BY DAVID DAVIDSON

16 **Validate Your Perimeter Security Devices are Working**
SECURITY DEVICE TESTING *Validation tests on your perimeter network security tools such as antimalware can help identify security gaps and misconfigurations.* BY JOEL SNYDER

19 **Buying an IPS: Determine Why You Need Intrusion Prevention**
INTRUSION PREVENTION SYSTEMS *Develop the right IPS strategy by first asking why your organization needs intrusion prevention.* BY JOEL SNYDER

22 **Preparing for a Network Security Audit Starts with Monitoring and Remediation**
NETWORK AUDIT READINESS *A network security change-management and remediation process can make audit preparation easier.* BY MIKE CHAPPLE

26 **VENDOR RESOURCES**

TABLE OF CONTENTS

AUTOMATED MONITORING

NETWORK SECURITY MANAGEMENT

INTEGRITY MONITORING

SECURITY DEVICE TESTING

INTRUSION PREVENTION SYSTEMS

NETWORK AUDIT READINESS

SPONSOR RESOURCES

AUTOMATED MONITORING

Starting Points for Network Monitoring

When looking at automated network monitoring tools, your first considerations should be uptime and reachability of systems. BY JOEL SNYDER

ONCE YOUR network grows beyond a few systems, it's time to think about automated monitoring. There are tons of things you can look at, ranging from bandwidth and usage to application latency, but a good place to start is right in the middle: reachability and uptime.

The goal of reachability and uptime monitoring is to assure that the network and its applications are available when users need them, and that remediation occurs as quickly as possible when there is a problem.

You'll immediately discover some good news when you decide to implement reachability monitoring: you're not the first person to do this. In fact, a slew of outstanding open source and commercial tools are available to help you build the right solution for your network. I can't offer specific advice about which is best, because all excel in different ways. However, I can say that you'll have no problem finding a product that meets your needs because there are so many good choices out there.

No matter which product you choose, though, you'll want to keep in mind some important guidelines in designing your monitoring.

1. Mixing network monitoring and application monitoring is a good idea. Many of the tools, concepts, reports and alerting strategies overlap when you're looking at network monitoring versus application and server monitoring. Thus, plan to do all your reachability monitoring at the same time. A single monitoring system is generally capable of handling thousands of systems without breaking much of a sweat, so you can amortize the hardware, software and human resources required to put monitoring in place across multiple functional areas. However, don't be pedantic about doing everything on one platform if it doesn't do a good job. For example, we use a commercial network monitoring package and an open source bandwidth monitoring package for our network monitoring. The commercial package does do bandwidth monitoring, but it's clumsy and doesn't give us all the trivia we want. Thus, we added a second tool even though there's overlap.

2. Device uptime is not the same as application uptime. Using network reachability tools such as "ping" to determine if a system is up isn't a good idea, because systems can respond to pings yet be entirely dysfunctional. As you identify key systems in your network, focus on the applications running on those systems and making sure that

TABLE OF CONTENTS

AUTOMATED MONITORING

NETWORK SECURITY MANAGEMENT

INTEGRITY MONITORING

SECURITY DEVICE TESTING

INTRUSION PREVENTION SYSTEMS

NETWORK AUDIT READINESS

SPONSOR RESOURCES

each of the applications is running properly. For example, pinging a Web server tells you very little. Even connecting to port 80 of the Web server doesn't help much. What you want to do is connect to port 80 and retrieve a known document, validating that you're actually getting the document you want. That doesn't tell you that every part of the Web server is running properly, but it tells you a lot more than ping does. Email servers are the same way: ideally, you want to generate a message and send it in using SMTP, then use a protocol such as POP or IMAP to validate that the message was received. This gives you much more end-to-end assurance that things are working properly. For network devices, check metrics within the devices themselves. For example, we look at CPU usage, memory usage, and fan and power supply status in our switches.

3. Build a multitier alerting strategy. Your network monitoring system should generate reams of reports and display pretty Web pages, but one of the most important functions is alerting in the event of a problem. Step back and build a simple alerting strategy, then be rigorous in your use of templates (or some equivalent feature) to apply these alerts to your devices. Alerting occurs across two axes: first, who are you going to alert and how; second, what are the escalation points for alerting.

Your network monitoring system should generate reams of reports and display pretty Web pages, but one of the most important functions is alerting in the event of a problem.

The "who and how" of alerting will probably initially use SMTP email, but that won't work if your email server is down, or the network doesn't let the message through. Thus, you should get an out-of-band alerting method, such as dialing a phone line to touch-tone someone's pager or, preferably, using a wireless modem to send SMS messages to cell phones—something that completely bypasses the Internet. Not every alert has to go to a person. For example, if a system or application is down for only 30 seconds or so, you may want to send this to your central logging server rather than spamming yourself. At our company, we have set up four levels of "who" to alert, ranging from the lowest (a SYSLOG to a server), through two different types of email messages, all the way up to pager and SMS broadcasts.

Use different escalation policies in alerting as you decide how important an application or device is. For example, you will have some critical devices and applications that should escalate up through your various levels quite quickly. Other devices, such as printers, might fall into the "interesting but not urgent" category, with emailed alerts only after extended downtimes. And you might even have a third category of still less important devices. When I design monitoring systems, I sometimes even have a final category of devices that are being monitored, but which never send alerts—they simply show up on end-of-month reports.

Joel Snyder is a senior partner at Opus One, an IT consulting firm specializing in security and messaging.

SECURITY SO SIMPLE YOU FEEL

INVINCIBLE

WORRY LESS. ACCOMPLISH MORE. WWW.SOPHOS.COM

SOPHOS
simply secure

TABLE OF CONTENTS

AUTOMATED MONITORING

NETWORK SECURITY MANAGEMENT

INTEGRITY MONITORING

SECURITY DEVICE TESTING

INTRUSION PREVENTION SYSTEMS

NETWORK AUDIT READINESS

SPONSOR RESOURCES

NETWORK SECURITY MANAGEMENT

How to Maintain Network Control Plane Security

Use access control lists and secure configurations to maintain the security of your organization's network control plane.

BY JOEL SNYDER

IT DOESN'T happen very often, but when Cisco sends out a security advisory about one of their routing or security products, there's a big splash. Almost all of these advisories can be summarized like this: "If someone out on the Internet sends some bad packets to your Cisco device, and if your device is listening to them, then something bad will happen."

The phrase in that alert you need to pay attention to? "If your device is listening to them."

It shouldn't be.

Do you have SNMP enabled on edge devices? Fine…so long as you also have an access list saying that those SNMP packets can only come from your management station. Is the management interface, whether HTTP, HTTPS, SSH or (heaven forbid) Telnet, running?

Think of it as a different network that runs in parallel to your data network, and is used to control, monitor and manage the data network.

Fine…so long as no one outside our network can ever get there.

We call this the "control plane" or "management plane." Think of it as a different network that runs in parallel to your data network, and is used to control, monitor and manage the data network. In huge networks, there is a true network control plane that is completely separate from the data that the device sees. But in many smaller networks, control plane, management plane, and data plane run on the same wire.

You can, and should, secure your network control plane in many ways, but they mostly come down to two techniques: access control lists and self-protection.

ACCESS CONTROL LISTS MANAGE TRAFFIC TO EDGE DEVICES

Access control list protections usually occur when you put a block of some sort in non-firewall devices at the edge and core of your network. A good example is SNMP. Let's say you have an SNMP management station at 10.20.30.161. That represents the one valid flow to and from that management station to network and security devices. Now, any other SNMP traffic floating around on your network, or coming in from the edge, should be blocked. If you have intermediate routers in your network, and certainly if you have firewalls, you should use them to block SNMP traffic—and any

TABLE OF CONTENTS

AUTOMATED MONITORING

NETWORK SECURITY MANAGEMENT

INTEGRITY MONITORING

SECURITY DEVICE TESTING

INTRUSION PREVENTION SYSTEMS

NETWORK AUDIT READINESS

SPONSOR RESOURCES

other management traffic—to your security and network devices, except from authorized sources.

You can get as strict as you want. For example, you can simply block all SNMP anywhere in your network except to and from the official management station. Here's an example using Cisco Systems Inc. access list syntax (once you define these access lists, don't forget to apply them to the appropriate interfaces):

```
permit udp 10.20.30.161 any eq snmp
permit udp any 10.20.30.161 eq snmptrap
deny udp any any eq snmp log
deny udp any any eq snmptrap log
```

Or you could put a block in to just protect the network and security devices. Usually, stricter is better, but if you don't know who else might be using SNMP, then you can focus on the devices that run your network.

At the edge, a much stricter approach is appropriate. In this case, you should be blocking all traffic directed at your firewalls, load balancers, and routers on their management addresses. Remember: No one on the Internet needs to send packets to your firewall, or to your external router. They legitimately send packets through those devices all the time, but the packets are never destined (at the IP layer, anyway) directly to the device. They're always for some IP address behind the device. The only time you may want to consider letting traffic come directly to the management IP of your external security and network devices is for PING traffic—it's a very useful debugging tool and worth letting traffic come in.

With intelligent user rights, it has become important to understand the roles and responsibilities of an individual when determining his or her access to applications and services.

Here's an example, using Cisco syntax, of blocking access to a device 128.182.35.42:

```
deny IP any host 128.182.35.42
```

If you wanted to block all SNMP incoming, you could do something like this:

```
remark *** Deny all other SNMP incoming
deny udp any any eq snmp
deny udp any any eq snmptrap
```

If you're in a NAT environment and you're using the external IP address of your firewall or router both for management and NAT, here is some advice: Don't do that. You're asking for security trouble, because now you have the same IP address being used for two things. IP addresses may be in short supply, but they're not in that short supply. Here's an example in case you can't separate out NAT from other traffic, assuming you know which ports your router or firewall are listening to (not a very good assumption, as the Cisco advisories show):

```
remark *** Block obvious access to mgmt plane; allow others
deny tcp any host 128.182.35.42 eq 22
deny tcp any host 128.182.35.42 eq www
deny tcp any host 128.182.35.42 eq 443
permit ip any host 128.182.35.42
```

CONFIGURE SECURITY DEVICES TO PARTICULAR TRAFFIC

Another protective technique should be self protection. With self protection, you configure the network or security device so that it doesn't listen to traffic it shouldn't hear. On devices such as routers, you'll want to create a local access list that only allows management traffic from authorized management networks. If you can, also disable management protocols and interfaces you aren't using. On devices such as firewalls, there is often a series of check boxes that let you turn on or off management on certain interfaces. It doesn't need to be enabled on the outside, ever. That's what VPNs are for, if you really need external management.

Sometimes you want to disable protocols entirely. Most people, for example, do not manage Cisco routers using HTTP. Here's a configuration example that's double overkill: disabling the HTTP server, and then also putting an access list on it, just in case.

```
no ip http server
ip http access-class 21
ip http authentication local
no ip http secure-server
access-list 21 deny any
```

And even if you do have management enabled, you should also add lists of authorized management addresses. It shouldn't be possible for someone who happens to be inside your network to connect to the management IP of your firewalls, routers, or other security devices, unless they're on the official management network.

For example, again using Cisco syntax, here is what the SNMP part of the router configuration might look like in a self-protective mode of operation:

```
snmp-server community public RO 6
snmp-server community vewysekwitpassword RW 6
snmp-server location Opus One/Tucson, Arizona
access-list 6 permit 203.209.92.105
access-list 6 permit 192.245.12.0 0.0.0.255
access-list 6 deny any log
```

Joel Snyder is a senior partner at Opus One, an IT consulting firm specializing in security and messaging.

TABLE OF CONTENTS

AUTOMATED MONITORING

NETWORK SECURITY MANAGEMENT

INTEGRITY MONITORING

SECURITY DEVICE TESTING

INTRUSION PREVENTION SYSTEMS

NETWORK AUDIT READINESS

SPONSOR RESOURCES

Just downloaded the customer
database onto a thumb drive.

INTEGRITY MONITORING

Network-Based Integrity Monitoring Keeps Website Hacks in Check

Network integrity monitoring brings the concept of file-based integrity monitoring to a company's online presence.

BY DAVID DAVIDSON

WE ARE used to the concept of file-based integrity monitoring (FIM), where we monitor important files and binaries on internal servers, guaranteeing the integrity of the system if they are intact. If configuration files, binaries or the kernel is modified, that action is detected and traced to determine if it was authorized.

FIM is common and mandatory by compliance requirements such as PCI DSS and HIPAA, but we don't often see the network-based integrity monitoring, where the same concept is applied online to Whois or DNS information, for example. How do you know if Whois information has been altered, if your DNS has been tampered with and users are being redirected to a phishing site, or if your Web server has been hacked and its index page defaced?.

While there are tools available that monitor website availability, we don't see many applied to check their integrity. We need a reliable way to detect if a company's network presence, such as a website, Web applications, DNS or Whois has been altered.

While there are tools available that monitor website availability, we don't see many applied to check their integrity.

Website modifications may be detected by most FIM products running on the server. However, if the attack is more subtle, such as a DNS redirection or a modification of the Whois with the registrar, your FIM will not detect it and users could be redirected to a malicious site.

MANUAL NETWORK INTEGRITY MONITORING

Manual network integrity monitoring can be done with a handful of scripts and a daily (or hourly) *cron* job on most Linux systems. On Windows, it is also possible, but since the OS lacks some basic networking tools (such as Whois), we will focus on Linux. To start, you can setup *lynx* or *wget* to download your website pages and perform a md5/sha1 checksum to compare the outputs:

```
    mkdir /nim
```

TABLE OF CONTENTS
AUTOMATED MONITORING
NETWORK SECURITY MANAGEMENT
INTEGRITY MONITORING
SECURITY DEVICE TESTING
INTRUSION PREVENTION SYSTEMS
NETWORK AUDIT READINESS
SPONSOR RESOURCES

```
cd /nim
lynx —dump —source http:// yoursite .com > /nim/tmp-source.txt
lynx —dump http:// yoursite.com > /nim/tmp-dump.txt
md5sum /nim/*.txt > file-wish-hashes.txt
sha1sum /nim/*.txt >> file-wish-hashes.txt
md5sum -c /nbim/file-with-hashes.txt
sha1sum -c /nbim/file-with-hashes.txt
```

You can do the same to monitor the Whois and DNS:

```
Whois yourdomain.com > /nim/Whois.txt
host -t ANY yourdomain.com > /nin/dns.txt
md5sum /nim/*.txt > file-wish-hashes.txt
sha1sum /nim/*.txt >> file-wish-hashes.txt
md5sum -c /nim/file-with-hashes.txt
sha1sum -c /bim/file-with-hashes.txt
```

After this is done for the first time, you can edit the scripts to do only the md5sum/sha1sum compare (-c flag) and to run the diff command to see exactly what was modified:

```
diff /nim/Whois.txt /nim/Whois-old.txt | mail -s “Change detail” you @
email.com
cp -pr /nim/Whois.txt /nin/Whois-old.txt
md5sum /nin/Whois.txt > /nim/files-with-hashes.txt
```

This approach works well if you have a handful of systems to monitor, otherwise it can get complicated to keep track of all the scripts. Another issue is that if you are running it from within your company, you may not be seeing the same site as people on the outside. That’s why when you are monitoring your Internet presence, it is better to use an outside look.

AUTOMATED AND FREE NETWORK INTEGRITY MONITORING

To solve some of the issues with manual monitoring and provide a stable outside look at your Internet presence, we decided to develop a free network integrity monitoring application. It is called Sucuri NBIM and it simplifies all these steps for the user. It also provides a historic view of everything that changed, detailed diffs and availability information (if a resource was ever offline).

How powerful can it be? A few months back, during the development of this application, I got an email notifying me that the Whois information from one of my domains was modified. The alert was: *Sucuri nbim: www.xx.com (whois) modified*

```
Modifications:
  16,19c16,17
  < Status: clientDeleteProhibited
  < Status: clientTransferProhibited
  < Status: clientUpdateProhibited
```

TABLE OF CONTENTS

AUTOMATED MONITORING

NETWORK SECURITY MANAGEMENT

INTEGRITY MONITORING

SECURITY DEVICE TESTING

INTRUSION PREVENTION SYSTEMS

NETWORK AUDIT READINESS

SPONSOR RESOURCES

```
< Updated Date: 26-feb-2007
—- > Status: ok
> Updated Date: 07-jan-2009
End of Notification
```

As you can see, someone removed the lock flag from my domain, which is usually only done if you plan to transfer it to someone else. After a few minutes on the phone with the registrar and after all my passwords updated it was fixed. They also told me they are seeing lots of brute force attacks trying to get accounts in there.

Another example when Google's main website was modified for Mother's Day:

```
Sucuri nbim: www.google.com (whois) modified
Modifications:
6c6
< Google
—-
> Happy Mother's Day!
End of Notification
```

Not an attack, but this shows how powerful it can be if anyone outside your domain ever changes any of your sites. ▸

David Davidson, is a network security consultant, specializing in open source security and intrusion detection tools.

TABLE OF CONTENTS

AUTOMATED MONITORING

NETWORK SECURITY MANAGEMENT

INTEGRITY MONITORING

SECURITY DEVICE TESTING

INTRUSION PREVENTION SYSTEMS

NETWORK AUDIT READINESS

SPONSOR RESOURCES

SECURITY DEVICE TESTING

Validate Your Perimeter Network Security Devices Are Working

Validation tests on your perimeter network security tools such as antimalware can help identify security gaps and misconfigurations.

BY JOEL SNYDER

OUR COMPANY gets bids for penetration testing (the slang term is pen testing) all the time, and it's one of the least satisfying parts of the security business. But rather than riff about how bad pen testing is and how the results are so often misused, I'd like to encourage you to try some of your own penetration tests against your perimeter network security devices.

To get you started, let's pick one tiny piece of the picture: your antimalware tools. People usually start by asking "how well does my antimalware work?" That's a coverage test where you're looking to see how well your antimalware tool covers the attack space. Testing antimalware for coverage means throwing zillions of bits of badware against the defense and seeing what it catches. Most of us don't have the resources or patience for that kind of testing.

Let me suggest something different; step back and change the question to: Does my antimalware work at all? The answer may surprise you.

I like testing the effectiveness of antimalware in general, because this is a test we can do something about. If you find out that your antimalware tool only catches, say, 93% of the samples, how are you going to get that number up? There are not a lot of options short of changing vendors. However, if you do a validation test and discover that you've got a hole in your perimeter network security tools, this is often something under your control. You may find misconfigurations in your tools, or you may find that things don't behave exactly the way you thought.

Start your testing by getting a virus. If you're conservative, head to eicar.org and grab their sample test virus. Antimalware authors universally agree that they EICAR virus will detect as a virus — but it's also completely harmless, just a text string. Working with live samples is exceptionally dangerous and I can't encourage it, although I will admit that I often get different results.

You should run these tests with your desktop antivirus turned on, and then again with it turned off. You think you have defense-in-depth? Let's find out for sure.

To do this testing, you will need a small Web and email server sitting on the Internet. An easy approach is to download a Unix-based virtual appliance with these tools pre-

TABLE OF CONTENTS

AUTOMATED MONITORING

NETWORK SECURITY MANAGEMENT

INTEGRITY MONITORING

SECURITY DEVICE TESTING

INTRUSION PREVENTION SYSTEMS

NETWORK AUDIT READINESS

SPONSOR RESOURCES

TABLE OF CONTENTS

AUTOMATED MONITORING

NETWORK SECURITY MANAGEMENT

INTEGRITY MONITORING

SECURITY DEVICE TESTING

INTRUSION PREVENTION SYSTEMS

NETWORK AUDIT READINESS

SPONSOR RESOURCES

installed and leave it running at home with a static IP address. One very valuable technique I use is to have services running on both standard and non-standard ports. Since malware authors don't play by the rules, you shouldn't either. For example, when you test for POP protocol, run it on the normal port 110, but also on a non-standard port such as 1100 and on a port used by another protocol: 53 (the DNS port) is a particularly good choice, but so are 80 and 443 (the HTTP ports).

The goal here is to validate that the antimalware works, not how many viruses it can catch. So a single example virus is good enough. What we want to do is try and get that single virus into our organization through every single hole possible.

Now, think about all the vectors into your organization. Email is the obvious one, that's a push avenue. Try sending the virus to yourself from the server you set up at home to your normal corporate mail server. Now ZIP it, and retest. Now double-ZIP it. Now double-ZIP it and password protect the file. Try more unusual archive formats, such as RAR or GZIP. Renaming the file is a simple trick, but can help as well. It may be EICAR.COM, but try changing it to FOO.PDF, FOO.ZIP, FOO.DOC, FOO.CSV, FOO.TXT, FOO.JPG and FOO.DLL. You may end up with a couple dozen tests, but you may also discover some results that surprise you.

Continue testing any other push avenues. Can people on the Internet upload files to an FTP site at your company? Post them on a forum webpage (where others inside or outside could download them)? Attach them to an incoming Web-based customer support request?

Once you've exhausted push transmission from the Internet into your organization, look at the many ways your staff can pull data. The list is almost endless, but you should start with Web browsing. Put all of the test files, in all of their varying formats, onto a series of webpages. Now try and download them all. Run the Web server on port 80. Run it on port 443, unencrypted. Run it on ports 25, 110, 143, 53, 8080 and 7633.

Another nice testing strategy is to put the viruses in webmail. Most commercial webmail services will detect and block the virus, but if you install Squirrelmail or any other open source webmail tool, you can build your own webmail service and point it at the email server you just created. Try testing webmail—don't forget those non-standard ports—along with POP and IMAP as well as outbound SMTP on both standard and non-standard ports.

Some other hints on validating your own antimalware: Keep a notebook by your side (whether a paper one or electronic one) and make notes on each test you run. It's easy when you start talking about 50 or 100 tests to get lost in the details, but keeping notes on what did and didn't work can help you to see trends and analyze results.

If you find out something surprising, drop me a note and let me know what you learned about your own defenses! ▸

Joel Snyder is a senior partner at Opus One, an IT consulting firm specializing in security and messaging.

INTRUSION PREVENTION SYSTEMS

Buying an IPS: Determine Why You Need Intrusion Prevention

Develop the right IPS strategy by first asking why your organization needs intrusion prevention. BY JOEL SNYDER

TODAY'S THREAT landscape far exceeds the protection traditional signature-based products can offer, making a feature-rich network-based intrusion prevention system (IPS) a must for not only malware containment, but network activity monitoring and compliance.

Companies anxious to dip into these waters need to sidestep some traps. Avoid vendor marketing fluff, and spend only on what you need. Putting the wrong IPS into your network can be a costly error, both in terms of capital and operational expenditures.

This article helps you answer the question: "Why am I buying an IPS?"

IPS Drivers: DDoS, Compliance, Alerting, Forensics and More

Before you talk to vendors about IPS—or any network security products—you need to understand what you want to accomplish and why you're buying IPS. There are many good reasons to add an IPS into a network:

- You could be looking for extra protection at the perimeter that employs signature-based technology to trap some of the bad things that manage to make their way through the firewall.
- Or you could be focused on mitigation of denial-of-service attacks, and looking for products that employ rate-limiting security parameters to protect against these kinds of threats.
- With a new, onerous, load of regulation in many organizations and industries, you could be looking for tools to help in your compliance efforts.
- Or, perhaps you might be looking for a product that provides IDS-like alerting and forensics to help you get a better handle on what threats are trying and have been successful at hitting your network.
- You could be hoping to build more security into the core of the network, perhaps protecting a specific set of servers inside the network or even by wrapping an IPS around the entire network core.
- You could be worried about incoming threats—or just as worried about detecting

TABLE OF CONTENTS
AUTOMATED MONITORING
NETWORK SECURITY MANAGEMENT
INTEGRITY MONITORING
SECURITY DEVICE TESTING
INTRUSION PREVENTION SYSTEMS
NETWORK AUDIT READINESS
SPONSOR RESOURCES

TABLE OF CONTENTS

AUTOMATED MONITORING

NETWORK SECURITY MANAGEMENT

INTEGRITY MONITORING

SECURITY DEVICE TESTING

INTRUSION PREVENTION SYSTEMS

NETWORK AUDIT READINESS

SPONSOR RESOURCES

and blocking infected systems on your own network from attacking the rest of the world.

Note that this isn't a comprehensive list, but each can be equally valid in the right environment. But until you know which apply to you, you won't be able to select the proper IPS strategy or product. Every IPS has a different set of design goals and features targeted to address a limited set of the questions posed here.

It would be easier for all involved if you could simply reduce this list of implementation reasons and goals into a feature checklist, something you could throw into an RFP and subsequently pick the vendor with all of the right boxes checked. But, unfortunately, that's impossible, not so much because the appropriate features are not in place, but because of the disparate philosophies that go into the products' design.

For example, it's easy to put forensics onto your checklist as a feature—assuming that is something you care about. Unfortunately, listing "forensics" won't get you any closer to finding the right product; it will only help you to eliminate some products that don't have any forensics capabilities.

The more appropriate question is: Why do you want forensics? Are you really looking to comply with the classic definition for forensics in which you need to collect data that could be used in a courtroom to help prosecute an attacker? Or are you simply looking for data collected and stored over a period of time that will ultimately help you to understand how an attack actually happened? Will you need to tap into the forensics ability of the IPS daily or just once a month? If you expect to run daily forensics, the performance and design of the forensics interface is a huge issue. While they may not be as important if you only need to review on a monthly basis, knowing why you want forensics will help you to understand what products will work best for you.

If you expect to run daily forensics, the performance and design of the forensics interface is a huge issue.

Create an IPS Needs Statement

The IPS market is crowded on many levels. There are products ranging from high-performance standalone appliances to others shipped as add-ins to existing firewalls. After studying this product space for several years, it has become clear that while there are often common denominators between some products—for example, quite a few of the newer IPS products use Snort as their underlying detection engine—that help segment the market into broad, overlapping categories, the underlying design goals and capabilities still vary widely.

The table below is a list of reasons why corporations we've worked with in the past three years have implemented an IPS in their networks and the noted tradeoffs expected with each choice. This may guide you to your own IPS needs statement. No single IPS device is designed to operate in every environment and solve all problems, which means that you will have to make choices and weigh your own reasons to balance these tradeoffs.

WHY YOU NEED INTRUSION PREVENTION

Spectrum of Reasons for Implementing IPS		Design Characteristics of an Appropriate IPS
from:	*to:*	
You are focused on perimeter security	You want to protect the core of your network	The closer an IPS is to the core of your network, the more important issues such as performance, high availability and control of overflow become. IPS functions pushed out toward an Internet boundary don't necessarily operate under the same performance constraints, and may be designed to handle failure cases (such as too much traffic or too high latency) differently.
You want to protect your servers	You want to protect end users (clients) on your network	When protecting servers, an IPS can be tightly tuned to inspect particular incoming services and particular applications. To protect client desktops, the IPS must handle incoming and more importantly, outgoing traffic with twin goals: prevent incoming infection and attack by blocking packets, but also detect a compromised system by its outbound attacks.
You are looking for signature-focused IPS protection	You are looking for rate-focused IPS protection	While most IPSes have signature- and rate-based technologies, one or the other is generally the product's "sweet spot." For example, when your main concerns are denial-of-service attacks, a product architecture focused on rate-based IPS is needed. If you are more focused on break-ins through system vulnerabilities and reconnaissance, signature-based IPS is more appropriate.
You are most concerned about specific attacks, such as hacker break-ins or viruses.	You are most interested in detecting anomalous behavior, such as a normally unused server suddenly going active	Although these two capabilities are by no means exclusive, most products specialize in one or the other. Simple anomalies, such as protocol errors, are common across the board (even in rate-based products), but more sophisticated detection scenarios, such as behavior anomalies, call for a different architecture.
You want to be able to detect attacks and have some forensics evidence on how it happened	You want the IPS to operate on its own, but you are not interested in using it as a security console or as a primary tool in investigations	While an IPS can detect and prevent attacks, adding a full forensics capability of any sort dramatically changes product architecture, increases costs and impacts performance.
You want IPS in place for primary protection against attackers and break-in attempts	You want IPS as an additional layer in a defense-in-depth strategy	IPS products positioned as a primary protective layer, typically behind a firewall, may have other features such as "shunning" of known attackers. These bring additional security, but at considerable risk such as self-inflicted denial-of-service. When an IPS is part of a layered defense strategy, features such as shunning are often unnecessary.

To understand why you're looking for an IPS, write an IPS needs statement, a single paragraph that begins with this phrase: "What we're trying to accomplish is…" With this in place, you'll be in a much more informed position to correctly evaluate IPS products for your environment. Only after you understand why you want to add an IPS to your network, can you ask yourself about security and coverage, performance, management, and form factor—the other four main criteria for successfully selecting an IPS strategy for your network.

Joel Snyder is a senior partner at Opus One, an IT consulting firm specializing in security and messaging.

TABLE OF CONTENTS

AUTOMATED MONITORING

NETWORK SECURITY MANAGEMENT

INTEGRITY MONITORING

SECURITY DEVICE TESTING

INTRUSION PREVENTION SYSTEMS

NETWORK AUDIT READINESS

SPONSOR RESOURCES

NETWORK AUDIT READINESS

Preparing for A Network Security Audit Starts with Monitoring and Remediation

A network security change-management and remediation process can make audit preparation easier. BY MIKE CHAPPLE

THERE ARE a few people that come visiting once a year: some that we enjoy (Santa Claus and the Easter Bunny), and some that we'd rather avoid (the IRS agent and the network auditor). I can't help you prepare for the IRS, but I can certainly offer you some advice to help prepare for your next network audit.

I've had the opportunity to frequently observe how organizations handle annual audits, and the most successful security teams approach the process as a periodic review of the way business is conducted all year, with the goal of compiling a complete scope and picture of enterprise network processes. Organizations that encounter the most difficulty during audits are those that adopt the "cram for the exam" approach. An audit isn't something that can be crammed for, and those that think so likely miss the point of conducting a security audit in the first place.

I recently attended a conference where Investment Technology Group CISO David Drossman compared audit preparation to the training program followed by professional football teams, in that winning the Super Bowl requires year-round preparation. How can you apply this philosophy to your next network audit? I propose that you follow a lifecycle based upon the one illustrated below. It's a compilation of the best practices that I've seen throughout my 10 years of work with IT organizations.

NETWORK SECURITY AUDIT LIFECYCLE

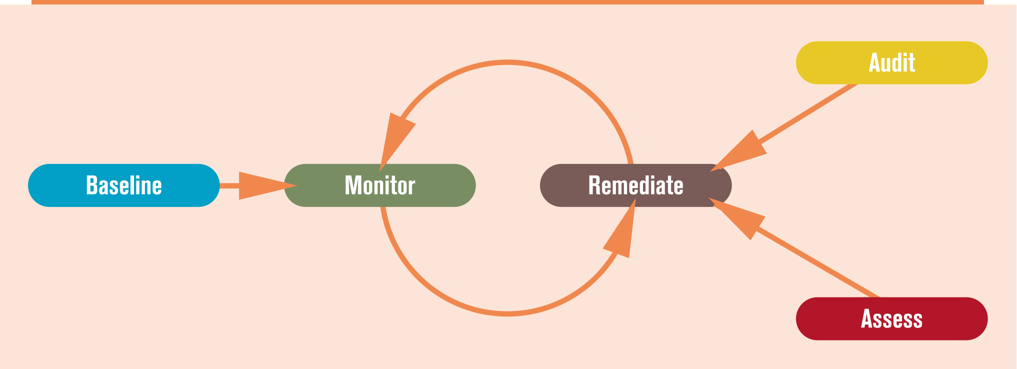

TABLE OF CONTENTS

AUTOMATED MONITORING

NETWORK SECURITY MANAGEMENT

INTEGRITY MONITORING

SECURITY DEVICE TESTING

INTRUSION PREVENTION SYSTEMS

NETWORK AUDIT READINESS

SPONSOR RESOURCES

The process involves a series of iterative steps:

1. Develop a well-documented snapshot of your network and how each device should perform. Repeat this process for each network element that you wish to manage. The number and types of network devices you monitor will depend upon your particular organization's infrastructure and the amount of time you have to dedicate to change monitoring. This snapshot will serve as a network "baseline" that you can use to measure changes against in the future. As an example, assume you are attempting to build a baseline of network firewall configurations. You can build this baseline by compiling the total picture of ports exposed by a data center's servers, as this information accurately reflects the performance of the firewall. Generally speaking, there's no need to go through the laborious process of building such a baseline very often; the frequency of repetition will depend upon how often the environment changes and the degree of compliance with the monitoring/remediation cycle.

2. Next, begin a continuous process that will monitor the network for changes. Continuing with the firewall example, you could use a daily Nmap scan to monitor for previously undetected firewall rules. A little Perl scripting and a simple database can automate this task, alerting you only when the environment changes. Alternatively, you may turn to advanced change management tools to detect and/or track changes to your environment. The degree of integration you achieve is limited only by your budget. For example, you might use an integrity monitoring tool and combine it with your organization's change management product. Such an arrangement will automatically reconcile identified configuration changes with change management records.

The degree of integration you achieve is limited only by your budget.

3. Remediate when you detect a change. Each time a network change is discovered, there are two possible reactions. If the change was expected, as reflected in your change management system, simply update your baseline and move along. If the change was unplanned and represents a potential risk to the network, begin investigating, documenting and remediating the problem.

These two options comprise the monitor/remediate cycle. Each time you encounter a change, you can either add it to your baseline or remediate, often by restoring the system to the pre-change state. By following this cycle, you'll stay abreast of network changes and maintain accurate network documentation. Essentially, your network will always be in audit-ready condition. If you use this cycle methodically, it will dramatically reduce the frequency at which you must rebuild your baselines. In an ideal environment where all changes are tracked and reconciled, you may only need to build a baseline when bringing a new system/device online. If changes slip by without proper tracking or remediation, you'll eventually need to rebuild your baseline to bring things back under control.

4. Prepare for the audit with an assessment. It's always a good idea to plan an internal assessment about 2-3 months prior to the actual audit. A practice test evaluates your lifecycle process and ensures that the controls you've put in place are functioning properly. Think of the assessment as a dress rehearsal for the audit.

Following a lifecycle approach to network auditing can't guarantee a perfect audit. (After all, auditors wouldn't exist if they couldn't find anything!) It will, however, ensure that you're as prepared as possible for your annual visit from the auditor.›

Mike Chapple, CISA, CISSP, is an IT security professional with the University of Notre Dame. He previously served as an information security researcher with the National Security Agency and the U.S. Air Force. Mike is a frequent contributor to SearchSecurity, a technical editor for Information Security *magazine and the author of several information security titles, including the* CISSP Prep Guide *and* Information Security Illuminated.

TABLE OF CONTENTS

AUTOMATED MONITORING

NETWORK SECURITY MANAGEMENT

INTEGRITY MONITORING

SECURITY DEVICE TESTING

INTRUSION PREVENTION SYSTEMS

NETWORK AUDIT READINESS

SPONSOR RESOURCES

TECHTARGET SECURITY MEDIA GROUP

SearchMidmarketSecurity.com

EDITORIAL DIRECTOR Michael S. Mimoso

SEARCHMIDMARKETSECURITY.COM
SENIOR SITE EDITOR Eric Parizo

NEWS EDITOR Robert Westervelt

SITE EDITOR William Hurley

ASSISTANT EDITOR Maggie Wright

ASSISTANT EDITOR Carolyn Gibney

ART & DESIGN
CREATIVE DIRECTOR Maureen Joyce

VICE PRESIDENT/GROUP PUBLISHER
Doug Olender

PUBLISHER Josh Garland

DIRECTOR OF PRODUCT MANAGEMENT
Susan Shaver

DIRECTOR OF MARKETING Kristin Hadley

SALES DIRECTOR Dara Such

CIRCULATION MANAGER Kate Sullivan

ASSOCIATE PROJECT MANAGER
Suzanne Jackson

PRODUCT MANAGEMENT & MARKETING
Corey Strader, Jennifer Labelle,
Andrew McHugh

SALES REPRESENTATIVES
Eric Belcher ebelcher@techtarget.com

Patrick Eichmann
peichmann@techtarget.com

Jason Olson jolson@techtarget.com

Jeff Tonello jtonello@techtarget.com

Nikki Wise nwise@techtarget.com

TECHTARGET INC.
CHIEF EXECUTIVE OFFICER Greg Strakosch

PRESIDENT Don Hawk

EXECUTIVE VICE PRESIDENT Kevin Beam

CHIEF FINANCIAL OFFICER Eric Sockol

EUROPEAN DISTRIBUTION
Parkway Gordon Phone 44-1491-875-386
www.parkway.co.uk

LIST RENTAL SERVICES
Julie Brown
Phone 781-657-1336 Fax 781-657-1100

REPRINTS
FosteReprints Rhonda Brown
Phone 866-879-9144 x194
rbrown@fostereprints.com

TechTarget
Where Serious Technology Buyers Decide

"Technical Guide on Network Security Monitoring" is published by TechTarget, 275 Grove Street, Newton, MA 02466 U.S.A.; Toll-Free 888-274-4111; Phone 617-431-9200; Fax 617-431-9201.

TABLE OF CONTENTS

AUTOMATED MONITORING

NETWORK SECURITY MANAGEMENT

INTEGRITY MONITORING

SECURITY DEVICE TESTING

INTRUSION PREVENTION SYSTEMS

NETWORK AUDIT READINESS

SPONSOR RESOURCES

SPONSOR RESOURCES

AVG Technologies

See ad page **2**

- How Malware Can Sneak Into Your Company Networks and How to Deal With It
- Why Traditional Anti-Malware Solutions Are No Longer Enough
- Social engineering: Hacking people, not machines

LogLogic, Inc.

See ad page **4**

- Enterprise Policy Management for Security and Compliance
- TechTarget *Information Security* Buyer's Guide — IT Decision Checklist: SIMs and Log Management
- Security and Compliance - It All Starts with Log Management

Sophos

See ad page **7**

- Security Threat Report: 2010
- How Unauthorized Applications Impact Security and How You Can Take Back Control
- Free tool: Free security scan - detect malware, devices and applications that can cause data loss and more.

ArcSight, Inc.

See ad page **11**

- SIMs and Identity Management: Creating a New Security Paradigm
- Combine SIM and IAM for Forensic Incident Response
- Demonstrating the ROI for SIEM: Tales from the Trenches

TABLE OF CONTENTS
AUTOMATED MONITORING
NETWORK SECURITY MANAGEMENT
INTEGRITY MONITORING
SECURITY DEVICE TESTING
INTRUSION PREVENTION SYSTEMS
NETWORK AUDIT READINESS
SPONSOR RESOURCES