



Network-Based Malware Detection: Filling the Gaps of AV

Version 1.3

Released: February 7, 2012

Author's Note

The content in this report was developed independently of any sponsors. It is based on material originally posted on [the Securosis blog](#), but has been enhanced, reviewed, and professionally edited.

Special thanks to Chris Pepper for editing and content support.

Licensed by Palo Alto Networks



Palo Alto Networks™ is the network security company. Its next-generation firewalls enable unprecedented visibility and granular policy control of applications and content – by user, not just IP address – at up to 20Gbps with no performance degradation. Based on patent-

pending App-ID™ technology, Palo Alto Networks firewalls accurately identify and control applications – regardless of port, protocol, evasive tactic or SSL encryption – and scan content to stop threats and prevent data leakage. Enterprises can for the first time embrace Web 2.0 and maintain complete visibility and control, while significantly reducing total cost of ownership through device consolidation. Most recently, Palo Alto Networks has enabled enterprises to extend this same network security to remote users with the release of GlobalProtect™ and to combat targeted malware with its WildFire™ service. For more information, visit www.paloaltonetworks.com.

Copyright

This report is licensed under Creative Commons Attribution-Noncommercial-No Derivative Works 3.0.



<http://creativecommons.org/licenses/by-nc-nd/3.0/us/>

Table of Contents

| | |
|---|-----------|
| Introduction | 4 |
| Current State = FAIL | 4 |
| The Value of Pattern Matching | 4 |
| Nothing Is Perfect | 5 |
| Identifying Today's Malware | 6 |
| Sandboxing and Evolving Heuristics | 6 |
| Network Impact | 7 |
| Where to Detect the Bad Stuff? | 8 |
| Traditional Endpoint-Centric Approaches | 8 |
| Content Security Gateways | 9 |
| Network Security Perimeter Gateways | 10 |
| The Impact of the Cloud on Malware Detection | 12 |
| AV on the Box | 12 |
| Sandboxing on the Box | 12 |
| Leveraging the Cloud for Malware Detection | 13 |
| No Silver Bullets | 14 |
| About the Analyst | 15 |
| About Securosis | 16 |

Introduction

Current State = FAIL

It's no secret that our existing malware defenses aren't getting the job done. Not by a long shot. Organizations large and small continue to be compromised by all sorts of issues. Application attacks. Drive-by downloads. Zero-day exploits. Phishing. But all these attack vectors have something in common: they are all means to an end.

It's no secret that our existing malware defenses aren't getting the job done. Not by a long shot.

That *end* is gaining a foothold in your organization, usually by installing some kind of malware on your devices. At that point – once the bad guys are in your house – they can steal data, compromise more devices, or launch other attacks. Or more likely all of the above. But most compromises nowadays start with an attack dropping some kind of malware on a device.

And it's going to get worse before it gets better – these online-fraud operations are increasingly sophisticated and scalable. Their software developers use cutting-edge development techniques. They test their code using services that run malware through various anti-malware

engines before deployment, to ensure they evade that low bar of defense. They use cutting-edge marketing tactics to achieve broad distribution and to reach as many devices as possible. All to further their objective: getting that foothold in your organization. So it's clear the *status quo* of anti-malware detection isn't cutting it, and will not moving forward. We know — that part is obvious.

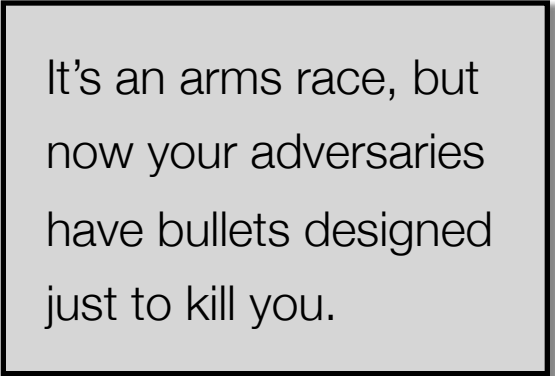
The first generation of anti-malware was based on signatures. You know, the traditional negative security model: building a list of what's bad and then looking for those bad files on each device. Whether deployed as endpoint anti-virus, content perimeter AV (typically inspecting email or web traffic), or network-based (IDS/IPS), the approach was largely the same. Look for *bad* and block it. Defense in depth meant using different lists of signatures and hoping you'd catch the bad stuff. But *hope* is not a strategy.

The Value of Pattern Matching

You may see that diatribe as an indictment of all pattern matching approaches – the basis of the negative security model. But that's not our position. Our point is that these outdated approaches look for the wrong

patterns in the wrong data sources. We need to evolve our detection tactics beyond what you see on *your* endpoints and *your* networks. We need to band together and get smarter. Leverage what we see collectively and do it *now*.

It's an arms race, but now your adversaries have bullets designed just to kill *you*, in the form of targeted malware expressly built to compromise *your* defenses. But malware works in a finite number of ways. There are only so many different registry keys or system files that can be tampered with. So if you can profile these proverbial *ways to die*, you can look for them regardless of the specific attack or how it is targeted.



It's an arms race, but
now your adversaries
have bullets designed
just to kill you.

Malware leaves tracks because it *must* impact the device and leverage the network in order to function. Maybe it's how the malware phones home. Perhaps it's the kind of network traffic that is sent, its frequency, or an encryption algorithm. It could be the type of files and/or behavior of devices compromised by this malware. Maybe it's how the malware was packed or how it proliferates. Most likely it's *all of the above*. You may need to recognize several possible indicators for a solid match. The point is that you can profile the malware and then look for those indicators in a number of places across your environment – including the network.

We have been doing anti-virus within email security gateways for years. But that was just moving the old approach to the perimeter. This is different. This is about really *understanding* what the files are doing, and then determining whether that behavior is bad. By leveraging the collective power of the network we can profile bad stuff much more quickly. With the advancement of network security technology we can start to analyze those files **before** they make their way onto our devices. Can we actually *prevent* an attack? Under the right circumstances, yes.

Nothing Is Perfect

Of course we cannot detect every attack before it does anything bad. We have never believed in 100% security, nor do we think any technology can protect an organization from a targeted and persistent attacker. But we certainly can (and need to) leverage some of these new technologies to react to these attacks more quickly.

In this paper we will talk about the tactics needed to detect today's malware attacks, and the kinds of tools and analysis required; then we will critically assess the best place to perform that analysis – whether on the endpoints, within the perimeter, or in the 'cloud' (whatever that means).

As always, we will evaluate the pros and cons of each alternative with our standard brutal candor. Our goal is to make sure you understand the up and down side of each approach and location for detecting malware, so you can make an informed decision about the best ways to fight malware moving forward.

Identifying Today's Malware

With rapidly morphing executables, increasingly sophisticated targeting, zero-day attacks, and innovative cloaking techniques, matching a file to a *known bad* AV signature is simply inadequate as a detection mechanism these days. We need to think differently about how to detect these attacks, which means working to figure out exactly what a file is doing and using that information to determine whether it's bad.

Sandboxing and Evolving Heuristics

We are talking about network-based malware detection, so we will assume you see all the streams coming into your network from the big bad Internet. With visibility into all ingress traffic, a perimeter device re-assembles the files from these streams and can analyze them. There are two main types of file-based analysis: static and dynamic. Static testing is basically taking a look at the file and looking for markers that indicate malware. This generally involves taking a file hash and matching it against a list of known bad files – effectively a signature – as well as identifying file packers and function calls that indicate badness.

Of course static analysis provides limited value, and we wouldn't want to bet on its findings – especially given that modern malware writers encrypt and otherwise obscure their files. This means you really need dynamic analysis: actually *executing* the file to see what it does. Yes, this is playing with live ammo – you need proper precautions to insulate your network and make sure that *running suspected malware* on your gear doesn't put you at risk.

Dynamic analysis effectively spins up an isolated vulnerable virtualized system – the sandbox – to host and execute the file; then you can observe its impact on the device and network. Clear indications of badness include configuration changes, registry tampering, installing other executables, buffer overflows, memory corruption, and a zillion other bad things malware can do. Based on this analysis, the perimeter gateway checks files against policies and may block bad files.

Given the real-time nature of network security it is not feasible to have a human review all the dynamic analysis results, so you are dependent on the detection algorithms and heuristics used by the security device to identify malware. The good news is that these capabilities are improving and reducing false positives. But innovative malware attacks (including zero-days) are not caught by perimeter gateways – at least not the first time – which is why multiple layers of defense still make a lot of sense. The death of defense in depth has been greatly exaggerated.

What's the catch? Clearly sandbox analysis is less effective at detecting advanced malware that is *VM-aware*. The malware writers aren't dummies, so they now check whether the OS is running in a virtual environment and act accordingly – typically going dormant.

Network Impact

Another aspect of dynamic malware analysis is profiling how the malware leverages the network. Remember the [Securosis data breach triangle](#): without exfiltration there is no breach. Malware depends on the network, both to get commands from the mother ship and to exfiltrate data. Dynamic analysis evaluates what networks the malware communicates with as another indication of badness.

Remember the
Securosis data
breach triangle:
without exfiltration
there is no breach.

But how can these network devices keep track of the millions of domains and billions of IP addresses which might be command and control targets? The good news is that we have seen this movie before. Reputation analysis has evolved to track these bad IP addresses and networks. The first incarnation of reputation data was URL blacklists maintained by web filtering gateways. That evolved to analysis of IP addresses, predominately to identify compromised mail relays for anti-spam purposes.

Now that model been has extended to analyze DNS traffic to isolate command and control (C&C) networks as they crop up. Malware writers constantly test malware and

new obfuscation approaches for their C&C traffic, but heuristic approaches can identify emerging C&C targets by analyzing DNS requests, exfiltration attempts, and network traffic. For example, if an IP address is the target of traffic that looks an awful lot like C&C traffic, it might be an emerging bot master ramping up operations. It's not brain surgery, and this type of analysis is increasingly common for network security gateway vendors. Obviously to keep current any vendor providing this kind of botnet tracking needs access to a huge amount of Internet traffic. So if your vendor claims to track botnets be sure to investigate how they track C&C networks, and substantiate their claims.

Why is isolating C&C traffic important? It all gets back to the detection window. Even with network-based malware gateways you *will* miss inbound malware. Endpoint devices may still get compromised, but any obfuscated communications you detect to known C&C targets will help identify compromised devices. This isn't going to be definitive but it's an excellent place to start.

Outside C&C traffic, analyzing the network characteristics of malware also provides insight into proliferation. How does the malware perform reconnaissance and then spread? What kind of devices does it target? You can glean a treasure trove of information from static and dynamic analysis of malware files. But that is only the beginning. Once you know what it does you need to block it *before* it damages your environment.

The ultimate goal of any malware analysis is to profile a malware file and then block it when it shows up again. That's what AV did in the early days, and what you need from your malware defenses. By understanding how malware uses the network you can design controls to block it on the perimeter. Of course that's all easier said than done, but first we need to look at our options for locations to perform malware detection.

Where to Detect the Bad Stuff?

Our research has assumed so far that the network is the right place to detect malware. But we all know what happens when you assume anything, so where *should* you detect? Let's make like Hollywood types and divulge the answer at the beginning, in a transparent plot ploy. Drum roll, please... You should detect malware *everywhere* you can. On the endpoints, at the content layer, and on the network. *It's not an either/or decision.* But of course each approach has strengths and weaknesses. Let's dig into those pros and cons to give you enough information to figure out what mix of these options makes sense for you.

We start with a malware profile of something bad. Now comes the fun part: you actually look for it, and perhaps even block it before it wreaks havoc in your environment. You also need to be sure you aren't flagging things unnecessarily (those dreaded false positives), so care is required when you decide to actually block something. Let's weigh the advantages and disadvantages of all the different places we can detect malware, and put together a plan to minimize the impact of malware attacks.

Traditional Endpoint-Centric Approaches

If we jump in the time machine and go back to the beginning of the Age of Computer Viruses (about 1991?), the main threat vector was 'sneakernet': viruses spreading via floppy disks. Then detection on actual endpoint devices made sense, as that's where viruses replicated. That started an almost 20-year fiesta (for endpoint protection vendors, anyway) of anti-virus technologies becoming increasingly entrenched on endpoints, always three or four steps behind the latest attacks. Now this type of endpoint protection is widely considered ineffective.

Does that mean it's not worth doing any more? Of course not, for a couple reasons. First and foremost, most organizations just *can't* ditch their endpoint protection because it's a mandated control in many regulatory hierarchies. Additionally, endpoints are not always connected to your network, so they can't count on protection from the mothership. At minimum you probably still need *some* kind of endpoint protection.

Detect malware
everywhere you can.
On the endpoints, at
the content layer, and
on the network. *It's
not an either/or
decision.*

Of course network-based controls (just like all other controls) aren't foolproof, so having another (even mostly ineffective) layer of protection generally doesn't hurt. But obviously there are issues with endpoint protection, including the complexity of keeping anything up to date on thousands of endpoints. And endpoint products run well inside your network, which means anything they detect (and hopefully stop) has already breached your network — not a position of strength for dealing with malware. Obviously the earlier — and closer to the perimeter — you can detect and stop malware, the better.

Detecting malware is one thing but how can you control it on endpoints? You have a few options:

- **Endpoint Protection Suite:** Traditional AV (and anti-spyware and anti-everything-else). Most of these tools already use some kind of advanced heuristics, reputation matching, and cloud assistance to help detect malware. But tests show they still don't catch enough, and even if the detection rate is 80% (which it probably isn't) across your 10,000 endpoints, you could easily spend 30-40 hours *per day* cleaning up infected endpoints.
- **Browser Isolation:** Running a protected browser logically isolated from the rest of the device basically puts the malware in a jail where it can't hurt your legitimate applications and data. When malware executes you just reset the browser without impacting the base OS or device. This is more customer-friendly than forcing users to browse in a full virtual machine, but can the browser ever be completely isolated? Of course not, but this helps prevent stupid user actions from hurting users (or the organization, or *you*).
- **Application Whitelisting:** A very useful option for truly locking down particular devices, application whitelisting implements a positive security model on an endpoint. Specify all the things that are permitted to run, and block everything else. Malware can't run because it's unauthorized, and alerts can be fired if the device attempts an action that smells like malware. For devices which can be subjected to draconian lockdown, AWL makes a difference. But those tend to be a small fraction of the devices in your environment, which relegates AWL to a niche.

Remember, this isn't an either/or decision. You'll use one or more of these options, regardless of what you do on the network for malware detection.

Content Security Gateways

The next layer we saw develop for malware detection was the content security gateway. This happened as LAN-based email was becoming pervasive, when folks realized that sneakernet was horribly inefficient, and the bad guys could just send viruses and spread their malware via email. Ah, the good old days of self-propagating worms. So a set of email (and subsequently web) gateway devices were developed, embedding anti-virus engines to move detection closer to the perimeter.

Many attacks continue to originate as email-based social engineering campaigns, in the form of phishing email — either with the payload attached to the message, more often as a link to a malware site, and sometimes even embedded within the HTML message body. Content security gateways can detect and block the malware at any point during the attack cycle by stopping attached malware, blocking users from navigating to compromised sites, or inspecting web content coming into the organization and detecting

attack code. Many of these gateways also use DLP-like techniques to ensure that sensitive files don't leave the network via email or web sessions, which is all good.

The weakness of content gateways is similar to the issues with endpoint-based techniques: keeping up with the rapid evolution of malware. Email and web gateways do have a positive impact by stopping the low-hanging fruit of malware — specimens which are easy to detect due to known signatures — by blocking spam to prevent users from clicking something stupid, and by preventing users from navigating to compromised sites. But these devices, along with email and web-based cloud services, don't stand much chance against sophisticated malware because their detection mechanisms are primarily based on old-school signatures. And once a gateway passes a message through or allows a connection to a web site, the gateway is largely blind. It has no way to detect a compromised device, or to cut off or clean such a device.

Network Security Perimeter Gateways

Next let's discuss detecting malware on the network. This generally means at the network perimeter, but some organizations deploy network security devices internally — you will need to pick an architecture that suits your requirements. Network security devices are ubiquitous, so performing some level of detection on them makes sense. Devices on the perimeter have access to ingress traffic, so they can detect malware at the perimeter, *before* it reaches anything vulnerable. These devices can and should also scrutinize egress traffic, scanning for sensitive data and command and control (C&C) traffic which indicates malware that has successfully eluded other defenses.

The second enabler
for network security
perimeter
consolidation is
technology evolution.

So what's the catch? Scalability — malware detection can be resource-intensive. This once again raises the almost religious battle about unified threat management (UTM) devices of the past few years. As you may recall, UTM was soundly thrashed in enterprise circles because people refused to believe gateway devices could scale to inspect traffic and do malware analysis at the (near) wire speeds required for network security devices. Three years ago these detractors were not wrong. But the only constant in the security business is change, and a few factors have turned the tide here. First is nomenclature. Some vendors have recast their enterprise UTM devices

as *next-generation firewalls* to capitalize on the hype around these new devices. Though the underlying technology between NGFWs on the market differs fundamentally (for more details check out our [Enterprise Firewall](#) paper), they all provide similar capabilities: the ability to enforce application-oriented policies on network traffic.

The second enabler for network security perimeter consolidation is technology evolution. Chips get faster, algorithms improve, and cloud resources provide both compute power and much greater scalability for analysis than was previously available on any standalone perimeter gateway. Times change, and it's time to reexamine the ability of network devices to detect malware.

We hate jumping on bandwagons, but we can't minimize the importance of the innovative technical architectures of NGFW gear hitting the market now. These purpose-built devices process packets differently and enable much faster and deeper analysis of network traffic, enabling many of these sophisticated functions – including malware detection. Combined with the ability to farm out some of the compute overhead, and the network effect of shared malware profiles in the cloud (which we will discuss later), we believe some sort of network-based malware detection will emerge as a key security control over the next 18-24 months. Mostly because we need all the help we can get, and as we discussed regarding endpoint protection, the farther out (meaning closer to the network perimeter) we can detect and block attacks, the better.

But we understand your (and our own) natural skepticism about adding yet another function to the same box, which didn't work well with UTM, as those devices dramatically slowed down when loaded with multiple functions. Fortunately you don't need to run malware detection in an existing box. As usual, some vendors are happy to provide one more device, to sit right next to your existing perimeter security gateway and do malware detection. For a small fee, of course.

But *can* you do malware detection on the same device? In good hedging form (there is a US Presidential election coming up, after all), the answer depends on how much traffic you have to analyze, and how the device leverages *cloud services* for malware analysis. So the answer is *maybe*, but that's not the end of the story. Everything new and shiny in technology uses the cloud in some way, shape, or form, and malware detection is no different. If there is a way to utilize the cloud to address some of these performance and leverage limitations, running these capabilities on the same box could become much more practical. Which brings us to our next topic...

The Impact of the Cloud on Malware Detection

So far we have made the case for considering gateway-based malware detection as one of the next key capabilities needed on your perimeter. Now we need to wade through the hyperbole and evaluate the strengths and weakness of each approach to detect malware on the network.

AV on the Box

A comprehensive rundown of all the alternatives should start with the status quo, which is a traditional AV engine (typically OEMed from an endpoint AV vendor) on your perimeter security gateway. This is basically what lower-end UTM devices do. This approach focuses on detecting malware within the content stream (think email/web filtering), and — just like traditional AV — isn't very effective for detecting modern malware. AV doesn't work very well on your endpoint, and alas it's no better on perimeter gateways. Moving right along...

Sandboxing on the Box

A new type of malware detection device has emerged that executes malware in a protected sandbox on the device and observes what it does. Depending on the behavior of the file, as we discussed above — basically, whether it does bad things — it can be blocked in real time. Of course virtualizing victim devices on a perimeter network security device to dynamically analyze malware at network speeds is a substantial technical advance. We have seen these devices provide a measurable improvement in the ability to block malware at the gateway, so large enterprises show great interest in these devices.

AV doesn't work very well on your endpoint, and alas it's no better on perimeter gateways.

Of course this entails trade-offs. First of all, do you really want to be executing malware within your production network? Of course it is designed as an isolated environment constrained to the malware detection device, but it's still a risk — even if a small one. The second trade-off is performance — you are limited to the performance of the perimeter device. Only so many virtual victims can be spun up on a given network device at a time, and at some point you will hit a scalability wall. You can throw bigger boxes at the problem but local analysis is inherently limited.

And remember these are additional dedicated devices. For some organizations that isn't a problem – they simply get a new box to solve a new problem. Perimeter sprawl isn't an issue for these companies. Others are more resistant to spending rack space on the perimeter on yet another niche device. Finally, this model provides no leverage. It requires you to execute every suspicious file locally — even if the malware was sent to every company in the world, every company would need to execute the malware locally to figure out what it did. And detecting malware is an inexact science given very capable adversaries, which means you will probably miss the first time something comes in, and suffer the consequences. So you need a feedback loop to take advantage of what you learned during incident response and malware analysis. Shame on you if you do all the work to analyze the malware but don't make sure you catch it the next time.

More sophisticated malware detection on the perimeter gateway represents a major advance, and has helped to detect a lot of the lower-hanging fruit missed by traditional AV.

To net this all out, more sophisticated malware detection on the perimeter gateway represents a major advance, and has helped to detect a lot of the lower-hanging fruit missed by traditional AV. But ultimately this approach does not scale and doesn't do much to protect you from reinfection, which remains the bane of many network security professionals.

Leveraging the Cloud for Malware Detection

We often point out there is rarely anything really new – just recycled ideas packaged a bit differently. We see this again with network-based malware detection, as we did for endpoint AV. When it became impractical to continue

pushing a billion malware signatures to each protected endpoint, AV vendors started leveraging the *cloud* to track the reputations of individual files, determine whether they are bad, and then tell endpoints to block them. The vendor's AV cloud analyzes unknown files to determine whether to allow or block them, depending on what the file does. Of course that analysis isn't real-time so each new malware attack tends to succeed until the cloud learns and protects subsequent targets.

This concept also applies to detecting malware on the perimeter security gateway. A preliminary list of bad files can be stored on the network device, but obviously it cannot be comprehensive. Unrecognized files can then be uploaded to the cloud service for automated analysis (using static and dynamic techniques), with the cloud service issuing an *approve* or *block* verdict. This addresses a number of the issues inherent to local analysis, as described above. Sending possible malware off to the vendors's cloud service rather than executing it locally avoids computational performance limitations and overhead — assuming a reasonably fast network. The analysis isn't on your hardware, which means both that it doesn't burden existing perimeter security gateways (which are likely already overburdened dealing with all these [new application-aware policies](#)) and has no opportunity to escape the container and run *inside* your network.

And the vendor's cloud service provides excellent leverage. If organization A sees a new malware file and the cloud service learns it's bad, all subscribers to the cloud service can automatically block that malware and

any recognizable cousins, even if they have never seen it before. So the larger the vendor's coverage umbrella, the better their network effect, and the less likely you are to see (and be infected by) the first specimen of any particular malware file – instead you can benefit from other people's misfortune and block the malware when it shows up.

So what's the catch? As with the latest generation of endpoint AV, there is critical latency between when you see the attack and when specific malware files are recognized as bad. That could be days at this point, but as the technology improves (and it will) the window will shrink to hours. But there will *always* be a window of exposure, since you aren't actually analyzing the malware at the perimeter.

And detection can never be perfect – malware writers already make it very hard to profile their wares exactly, and their obfuscation attempts improve daily. It's the same arms race we have been dealing with since the early days of virus attacks. The bad guys hide their stuff, the good guys figure out ways to (sort of) detect it, and the cycle repeats. So don't expect perfection from these devices. Not that you'd ever be that naive. Right?

The goal is to block malware as close to the perimeter as possible, preferably *before* it reaches any endpoints. We all know that once malware reaches the inside of your network your risk increases dramatically. The earlier you can take malware out of play the better.

No Silver Bullets

We need to be very clear that network-based malware detection cannot *solve* the malware epidemic. A targeted attacker (and even some more sophisticated script kiddies) can and will evade these defenses. You still need to execute on the fundamentals of security – things like egress filtering (preferably with application-aware policies) and endpoint hygiene to reduce your attack surface. We continue to advocate a “React Faster and Better” approach because you *will* be compromised, and you need the ability to detect the breach and respond effectively. The industry continues to get better at detecting today's malware but it's still not good enough – you need to supplement these advances with complementary controls. Okay – we can get off our soapbox now.

Network-Based Malware Detection is a definite step in the right direction. Being able to block the malware at the perimeter and leverage emerging cloud-based analysis environments improves detection, scalability, and leverage. When you are battling well-funded and patient attackers you need all the help you can get.

If you have any questions on this subject, or want to discuss your situation specifically, feel free to send us a note at info@securosis.com or ask us a question via the Securosis Nexus (<http://nexus.securosis.com>).

Blocking malware at the perimeter and leveraging emerging cloud-based analysis environments improves detection, scalability, and leverage.

About the Analyst

Mike Rothman, Analyst/President

Mike's bold perspectives and irreverent style are invaluable as companies determine effective strategies to grapple with the dynamic security threatscape. Mike specializes in the sexy aspects of security — such as protecting networks and endpoints, security management, and compliance. Mike is one of the most sought-after speakers and commentators in the security business, and brings a deep background in information security. After 20 years in and around security, he's one of the guys who “knows where the bodies are buried” in the space.

Starting his career as a programmer and networking consultant, Mike joined META Group in 1993 and spearheaded META's initial foray into information security research. Mike left META in 1998 to found SHYM Technology, a pioneer in the PKI software market, and then held executive roles at CipherTrust and TruSecure. After getting fed up with vendor life, Mike started Security Incite in 2006 to provide a voice of reason in an over-hyped yet underwhelming security industry. After taking a short detour as Senior VP, Strategy at eIQnetworks to chase shiny objects in security and compliance management, Mike joined Securosis with a rejuvenated cynicism about the state of security and what it takes to survive as a security professional.

Mike published [The Pragmatic CSO](http://www.pragmaticcso.com/) <http://www.pragmaticcso.com/> in 2007 to introduce technically oriented security professionals to the nuances of what is required to be a senior security professional. He also possesses a very expensive engineering degree in Operations Research and Industrial Engineering from Cornell University. His folks are overjoyed that he uses literally zero percent of his education on a daily basis. He can be reached at mrothman (at) securosis (dot) com.

About Securosis

Securosis, LLC is an independent research and analysis firm dedicated to thought leadership, objectivity, and transparency. Our analysts have all held executive level positions and are dedicated to providing high-value, pragmatic advisory services.

Our services include:

- **The Securosis Nexus:** The Securosis Nexus is an online environment to help you get your job done better and faster. It provides pragmatic research on security topics that tells you exactly what you need to know, backed with industry-leading expert advice to answer your questions. The Nexus was designed to be fast and easy to use, and to get you the information you need as quickly as possible. Access it at <https://nexus.securosis.com>.
- **Primary research publishing:** We currently release the vast majority of our research for free through our blog, and archive it in our Research Library. Most of these research documents can be sponsored for distribution on an annual basis. All published materials and presentations meet our strict objectivity requirements and conform to our Totally Transparent Research policy.
- **Research products and strategic advisory services for end users:** Securosis will be introducing a line of research products and inquiry-based subscription services designed to assist end user organizations in accelerating project and program success. Additional advisory projects are also available, including product selection assistance, technology and architecture strategy, education, security management evaluations, and risk assessment.
- **Retainer services for vendors:** Although we will accept briefings from anyone, some vendors opt for a tighter, ongoing relationship. We offer a number of flexible retainer packages. Services available as part of a retainer package include market and product analysis and strategy, technology guidance, product evaluation, and merger and acquisition assessment. Even with paid clients, we maintain our strict objectivity and confidentiality requirements. More information on our retainer services (PDF) is available.
- **External speaking and editorial:** Securosis analysts frequently speak at industry events, give online presentations, and write and/or speak for a variety of publications and media.
- **Other expert services:** Securosis analysts are available for other services as well, including Strategic Advisory Days, Strategy Consulting engagements, and Investor Services. These tend to be customized to meet a client's particular requirements.

Our clients range from stealth startups to some of the best known technology vendors and end users. Clients include large financial institutions, institutional investors, mid-sized enterprises, and major security vendors.

Additionally, Securosis partners with security testing labs to provide unique product evaluations that combine in-depth technical analysis with high-level product, architecture, and market analysis. For more information about Securosis, visit our website: <http://securosis.com/>.