The CISO's Guide to Advanced Attackers

Version 1.5
Released: May 28, 2013

Securosis, L.L.C. 515 E. Carefree Highway Suite #766 Phoenix, AZ 85085 **T** 602-412-3051
info@securosis.com www.securosis.com

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 Dell SecureWorks

Dell Inc. (NASDAQ: DELL) listens to customers and delivers innovative technology and services that give them the power to do more. Recognized as an industry leader by top analysts, Dell SecureWorks provides world-class information security services to help organizations of all sizes protect their IT assets, comply with regulations and reduce security costs. For more information, visit www.dell.com/secureworks.

Copyright

Table of Contents

Sizing up the Adversary

Every year some new shiny object seems to be working security marketers into a frenzy. The Advanced Persistent Threat hype continues to run amok 3 years in, and doesn't seem to be abating at all. Of course there is still lot of confusion about what the APT is, and a Securosis post from early 2010 explains our view pretty well.

That said, security vendors are predictable animals, and they adhere to the classic maxim "If all you have is a hammer, everything looks like an APT." So it makes no difference what a security product or service does — they are *all* positioned as the only viable answer to stop the APT. Of course this isn't useful to security professionals who actually need to protect important things. And it's definitely not helpful to Chief Information Security Officers (CISOs) who have to explain their organization's security programs, set realistic objectives, and manage expectations to senior management and the Board of Directors.

So as usual your friends at Securosis are here to help you focus on what's important and enable you to wade through the hyperbole to understand what's hype and what's real. This paper provides a high-level view of these "advanced attackers" designed to help a CISO-level audience understand what they need to know, and maps out a clear 4-step process for dealing with advanced attackers and their innovative techniques.

Let's dismiss the common belief that advanced attackers always use "advanced attacks". That is simply not the case.

Defining Advanced Attacks

First let's dismiss the common belief that advanced attackers always use "advanced attacks". That is simply not the case. Of course there are innovative attacks such as Stuxnet, stealing the RSA token seeds to attack US defense sector organizations, and compromising Windows Update using stolen Certificate Authority signing keys. But those are exceptions, not the rule. These attackers are very business-like and do not waste valuable advanced (0-day) attacks. They would just as soon get an unsuspecting office worker to click a phishing email and subsequently use a known Adobe Reader exploit to provide a foothold in the environment. They get no prizes for using 0-day attacks.

This understanding changes the way we need to think about adversaries. The attacks we see vary greatly with the attacker's mission and assessment of the most likely (and easiest) way to compromise your environment. A good way to get your arms around potential advanced attacks is to understand the objectives of the attackers and infer targets for the attacks. Then you can assess the likelihood of certain attack vectors based on the target and the profile of the adversary. With this information you'll get a feel for the tactics you are likely to face and then you can evaluate controls to deter them — or at least slow them down.

The security industry would have you believe that a magic malware detection box on your perimeter or locking down your endpoints will block advanced attackers. You can't afford to believe everything you hear at security conferences.

The security industry would have you believe that a magic malware detection box on your perimeter or locking down your endpoints will block advanced attackers. Of course you can't afford to believe everything you hear at security conferences, so let's break down exactly how to determine what kind of threat you are facing.

Evaluate the Mission

Having the senior security role in an organization (yes, Mr./Ms. CISO, we're talking to you) means your job is less about *doing stuff* and more about defining the security program and evangelizing the need for security with senior management and peers. To start the process you need to learn what's important in your environment, which leads you to identify interesting targets for advanced attackers. But you don't have unlimited resources or capabilities to protect against every attack, so you need to prioritize your defenses.

When trying to understand what an advanced attacker will probably be looking for, you get a pretty short list:

1. Intellectual property
2. Protected customer data
3. Business operations (proposals, logistics, etc.)
4. Everything else

It is unlikely that you can really understand what's important to your organization by sitting in your office. A big part of the job is to talk to senior management and your peers to get a feel for what is important to them. After a few of these conversations it should be pretty clear what's really important

and what's not. Once you understand the likely targets for advanced attackers — the important stuff — you can make educated guesses at the adversaries you will face.

Profile the Adversary

You know the old saying about assuming anything, right? We understand that it is simplistic to make generic assumptions about the kinds of attackers you will face, but you need to start somewhere. So let's quickly develop capsule descriptions of some adversaries you may face. Keep in mind that many security researchers (and research organizations) have assembled dossiers on potential attackers, which we will discuss later in this paper.

1. **Unsophisticated:** These folks tend to favor smash and grab attacks, where they use publicly available exploits (perhaps leveraging attack tools such as Metasploit and the Social Engineer's Toolkit) or some kind of packaged attack kit they buy on the Internet. They are opportunists who take what they can get.

2. **Organized Crime:** The next step up the food chain is organized criminals. They invest in security research, test their exploits, and have a plan to exfiltrate and monetize what they find. They are also opportunistic but can be quite sophisticated in attacking payment processors and large-scale retailers. They tend to be most interested in financial data but have been known to steal intellectual property if they can sell it and/or use brute force approaches like DDoS threats for extortion.

3. **Competitor:** Competitors sometimes use underhanded means to gain advantage in product development and competitive bids. They tend to be most interested in intellectual property and business operations.

4. **State-sponsored:** Of course we all hear the familiar fretting about alleged Chinese military attackers but you can bet *every* large nation-state has a team practicing offensive tactics. As we've written on the Securosis blog, the Chinese are a bit different in using military resources to gain economic advantage, but all these nations using offensive tactics are interested in stealing all sorts of data — from both commercial and government entities. And some of them don't care much about concealing their presence.

Of course there are many other kinds of adversaries. The value of broader or deeper profiling depends entirely on your situation. But the process is the same and the list above offers a decent start on the kinds of folks you will see trying to get into your stuff.

Identifying the Most Likely Adversaries

Let's work through the process to identify your most likely targets and then back into who your adversaries are most likely to be. Then you should plan to be wrong. In security only fools think they have all the answers.

Once you have implemented sufficient controls to protect your important assets, you need to ensure you monitor extensively to detect things not covered in your plans.

Once you have implemented sufficient controls to protect your important assets, you need to ensure you monitor extensively to detect things that are not covered in your plans. You cannot eliminate surprises in this business, but you can lessen the impact of an unexpected attack from a different adversary targeting a lower-value (to your thinking, anyway) target. This paper focuses on advanced attackers, but keep in mind everything you do is also applicable to the unsophisticated attacker (and attack).

The Process for Advanced Attacks

At Securosis we tend to be process centric. So let's establish a high-level process to deal with advanced attackers. This paper will dig into each step with specifics about what you need to do.

1. **Threat Intelligence/Information Sharing:** A key defensive capability for dealing with advanced attackers is knowing who they are, where they are coming from, and what attacks they are using. This entails leveraging external threat intelligence to learn from the misfortune of others.

2. **Data Collection and Mining:** The next step is to implement a comprehensive monitoring initiative which instruments networks, systems, applications, and data with sensors to collect data that can be mined for indications (provided by threat intelligence) of imminent attack.

3. **Verification:** When you believe you are being targeted you need to do an initial damage assessment and kick your incident response process into gear. This involves verifying, validating, and ultimately figuring out the root cause, the degree of compromise, and any damage resulting from the attack.

4. **Breaking the Kill Chain:** Once the attack has been verified and the root cause has been identified you need to decide how to "break the kill chain" or remediate the issue. This is non-trivial decision requires feedback from senior management, legal counsel, and likely law enforcement and government.

Ultimately all these functions need to become systemized parts of your security program. So it is critical to ensure your program evolves to handle these advanced attackers while continuing to paying attention to other stuff (hygiene, everyday attacks, and compliance).

Master the Basics

Dealing with advanced attackers is not for unsophisticated or immature security organizations. What does that mean? You need to master security fundamentals and have good security practices in place. We will not go into detail here — you can check out the Securosis Research Library for chapter and verse on all sorts of security practices. But we must point out that you need to have already hardened key devices, implemented a strong device hygiene (patch and configuration management) program, and properly segmented your network to make it difficult for attackers to get at important data *before* you start worrying about advanced attackers. We can laugh about the futility of traditional endpoint protection but you still need some level of protection on key devices with access to sensitive data.

For the rest of this paper we will assume that you are ready to deal with an advanced attacker — you have a relatively mature security program in place with adequate control sets.

Intelligence, the Crystal Ball of Security

Armed with context on likely adversaries, we can move into the intelligence gathering phase. This involves learning everything we can about possible and likely adversaries, profiling probable behaviors, and determining which kinds of defenses and controls make sense to address higher probability attacks.

Most intelligence activities produce educated guesses. That's why we keep using the word 'likely'. But these guesses can provide a very useful head start on detecting advanced attacks. When you are racing the clock with an adversary in your environment, that head start might enable you to stop exfiltration of key data.

Profiling the Adversary

For better or worse the industry seems to believe that all intelligence is "threat intelligence." This generalization holds for most immature security organizations, which still struggle with the basics. But as your security organization matures and increases its competence you will see that threat intelligence is only a subset of broader intelligence gathering. To really understand your adversaries you need to go deeper than the indicators of compromise found in the latest attack.

"Zoltar" originally uploaded by Jonathan Reyes

You will want to learn what they do, how they do it, where they live, what they like to do, where they were trained, the tools they use, the attacks they have undertaken, the nuances of their attack code, and their motives. Yes, that is a big list, and not many organizations are positioned to gather that kind of real intelligence on adversaries. You can check out some of the publicly available information in the APT1 report, which provides unprecedented detail about a group of apparently state-sponsored Chinese hackers. That report provides a feel for the depth of intelligence needed to seriously combat advanced attackers.

In light of limited resources and even more limited intelligence expertise, you are likely to buy this kind of intelligence, utilize an information sharing and analysis center (ISAC), or get it from buddies with more resources and expertise. You can gather a lot of intelligence by asking the right questions within your information sharing community or talking to researchers at your strategic information security vendors. Depending on how the vendor packages intelligence you may need to purchase an additional service, or perhaps you can negotiate for access to their security researchers when you renew or upgrade your support/license/service agreement.

> The adversary intelligence you need goes well beyond what's published in the quarterly threat reports from all the security vendors.

Many vendors produce periodic research reports detailing attack and data breach trends. That's all good and well, but the adversary intelligence you need goes well beyond what's published in these quarterly threat reports. They tend to give away their least interesting data as marketing bait, but generally have much more interesting data for their own work. We expect advanced intelligence to be increasingly packaged as a value-add offering from security vendors.

Threat Intelligence Indicators

Now that we have defined the intelligence terminology we can get into the stuff that directly impacts security activity: the threat intelligence that has become such a hot topic in security circles. We have recently covered this topic extensively so we will stick to the high points in this paper, but we recommend you read our research on Building an Early Warning System, Network-based Threat Intelligence, and Email-based Threat Intelligence for much more detail on specific data sources and indicators to look for.

Here is a high-level overview of the general kinds of threat intelligence you are likely to leverage in your efforts to deal with advanced attackers.

Malware

Malware analysis is maturing rapidly and it is now possible to quickly and thoroughly understand exactly what a malicious code sample does, and define both technical and behavioral indicators which you can seek out within your environment, as described in gory detail in Malware Analysis Quant. More sophisticated malware analysis is required because the old methods of AV blacklisting are not good enough, given how polymorphic malware and other attacker tactics make it hard to match files exactly. Instead you will identify indicators of what the malware did to a device. To put it succinctly, malware identification is no longer about what file looks like — it is now about what it *does*.

A number of companies offer data on specific pieces of malware. You can upload a hash of a malware file; if the recipient has seen that file already they will match the hash and send back their analysis; otherwise you upload the whole file for a fresh analysis. The services run malware samples

through proprietary sandbox environments and other analysis engines to figure out what they do, build a detailed profile, and return a comprehensive report which includes specific behaviors and indicators. You can then search for those indicators within your environment to pinpoint devices possibly compromised.

You can also draw conclusions from the kinds of indicators you find. Have those tactics been tied to specific adversaries? Do you see these kinds of activities during reconnassaince, exploitation, or exfiltration. Your analysis can enrich these indicators with additional context to make better decisions about what the best next step will be.

Reputation

Since its emergence as a primary data source in the battle against spam, reputation data seems to have become a component of every security control. The most common reputation data is based on IP addresses, and provides a dynamic list of known bad and/or suspicious IP addresses. It has a variety of uses — for example learning that a partner's IP address has been compromised should set off alarms, especially if that partner has a direct connection to your network.

Besides IP addresses, pretty much everything within your environment can (and should) have a reputation. Devices, URLs, domains, and files, for starters. If you have traffic going to a known bad site, weird traffic coming from a vulnerable contractor-owned device, or even a known bad file showing up when a salesperson connects to the corporate network, you know it might be a problem. If something in your environment develops a bad reputation — perhaps as a spam relay or DoS attacker — you need to know ASAP, hopefully before your entire network gets blacklisted.

C&C Traffic Patterns

One specialized type of reputation which is now often a separate feed is intelligence on command and control (C&C) networks. These feeds track global C&C traffic and use it to pinpoint malware originators, botnet controllers, and other IP addresses and sites your devices should avoid. They also help to identify likely compromised devices within your network that communicate with malware controllers. Integrating this kind of network-based threat intelligence with an egress firewall or web filter might enable you to prevent exfiltration, or enable a more aggressive monitoring stance to identify what attackers are doing.

Of course advanced attackers do not make analyzing C&C traffic easy. They work hard to obscure their communications, including using compromised devices with "good" reputations as C&C nodes to confuse reputation filtering, and frequently changing the locations of their nodes using a variety of sophisticated Domain Generating Algorithms (DGA). So accurately identifying C&C traffic is currently a kind of black magic, but it is a critical aspect of intelligence.

Intelligence for Sale

You might think that if you could bottle and sell intelligence there would be an infinite market. And indeed we are starting to see development of a market for stand-alone security/threat intelligence services offering information on attacks and malware. We also expect a thriving market for detailed research about specific attackers to develop, sold to larger companies with the sophistication to take advantage of this intelligence.

We expect a thriving market for detailed research about specific attackers to develop, sold to larger companies with the sophistication to use the intelligence.

We also see significant value in getting intelligence focused on monitoring both high-visibility people such as executives, as well as, your critical intellectual property on the Internet. These kinds of services are a combination of Big Brother watching your executives to ensure they haven't been compromised, and Internet-based DLP capability looking for instances of your brand or intellectual property being misused.

Mining for Indicators

It is fortunate that you do not need to seal off every window of vulnerability to defend against advanced attackers, because that would be not only impractical but impossible. Advanced attackers *will* figure out a way to gain a foothold in your environment — actually they will find *multiple* ways. So if you hope for any kind of success your goal cannot be to simply stop them — you need to work on shortening the window between compromise and detection. We have called that Reacting Faster and Better for years. 5 years, to be exact, but who's counting?

If you hope for any kind of success your goal cannot be to simply stop them — you need to work on shortening the window between compromise and detection.

The idea is to monitor your environment, gathering key security information that can either identify typical attack patterns as they are happening (a SIEM-like capability), or more likely searching for technical indicators identified via intelligence activities.

Collecting All the Security Data

We say "all the security data" with tongues in cheeks, but we still strive for completeness. We have been saying Monitor Everything almost as long as we have been talking about Reacting Faster, because if you fail to collect data you won't get another opportunity to capture it later. Unfortunately most organizations don't realize their existing security data collection leaves huge gaps, until the high-priced forensics consultants explain that they cannot truly isolate the attack, or the perpetrator, or the malware, or much of anything, because they just don't have the data.

Most folks only need to learn that lesson once.

So it's critical to have a robust collection infrastructure to store all your security data. The good news is that you have likely been collecting security data for quite some time, and your existing investment and infrastructure should be directly useful for dealing with advanced attackers. Your existing log management system might be useful after all. Now let's consider the data you need to collect:

1. **Network Security Data:** Firewalls and IPS devices can generate huge logs of what's blocked, what's not, and which rules are effective. Their feeds that generally include source, destination, port, and protocol, or application identifiers from next-generation firewalls; whatever the specifics, they may identify attack traffic.

2. **Configuration Data:** One key resource to mine for indicators is device configuration data. It enables you to look for specific files and configurations that have been identified as indicators of compromise.

3. **Identity:** Information about logins, authentication failures, and other identity-related data is useful for matching against attack profiles from third-party threat intelligence providers.

4. **NetFlow:** Another data type commonly used in SIEM environments, NetFlow data includes information on protocols, sources, and destinations for network traffic as it traverses devices. NetFlow records are similar to firewall logs but far smaller, making them more useful for high-speed networks. Flows can identify lateral movement by attackers, as well as large exfiltration file transfers.

5. **Network Packet Capture:** The next frontier for security data collection is actually to capture all network traffic on key segments. Forensics folks have been doing this for years during investigations, but proactive continuous full packet capture — for the inevitable incident responses which haven't even started yet — is still an early market. For details on how full packet capture impacts security operations, check out our Network Security Analytics research.

6. **Application/Database Logs:** Application and database logs are generally less relevant unless they come from standard applications or components likely to be particularly targeted by attackers. But you might be able to discover unusual application and/or database transactions — which might represent bulk data removal, injection attempts, or efforts to attack your critical data.

7. **Vulnerability Scans:** These detail which devices are vulnerable to specific attacks, and help filter out devices that could not be exploited by certain attacks, in order to filter and focus searches.

Capturing data within the context of a compliance audit is fundamentally different than trying to detect advanced attacker activity.

Of course this isn't an exhaustive list, and fortunately you are likely already capturing much of this data. That is good but capturing data within the context of a compliance audit is fundamentally different than trying to detect advanced attacker activity.

We are sticking to the CISO view in this paper, so we won't dig into the technical nuances of collection infrastructure. But it must be built on a strong analytical foundation, with a threat-centric

view of the world rather than one focused on compliance reporting. More advanced organizations may already have a Security Operations Center (SOC) leveraging a SIEM platform for more security-oriented correlation and forensics to pinpoint and investigate attacks. That's a start but you are likely to need a more advanced analysis platform to handle the volume and complexity of today's security data.

Attack Patterns FTW

We have repeatedly mentioned the impossibility of blocking every advanced attack, that doesn't mean we shouldn't learn from the past and benefit from the misfortune of others. We discussed sizing up the adversary at the beginning of this paper — for insight into what is likely to be attacked and perhaps even how. That insight enables you to look for those attack patterns within your security data — the unmet promise of SIEM technology.

The ultimate disconnect with SIEM was the hard truth that you need to know what you are looking for. Far too many vendors forgot to mention that little requirement when selling a bill of goods. Perhaps they expected attackers to post their plans on Facebook or something? But once you do the work to model likely attacks on your key information, and then enumerate those attack patterns in your tool, you can get tremendous value. Just don't expect it to be fully automated.

The quickest way to get killed is to plan for the best case. You also need to ensure you're ready for the worst case. That is advanced attackers using attacks you haven't seen before in ways you don't expect.

In the best case you will receive an alert about a very likely attack because it's something you were looking for. But the quickest way to get killed is to plan for the best case. So you also need to ensure you're ready for the worst case. That is advanced attackers using attacks you haven't seen before in ways you don't expect. That is when all your gathered intelligence comes into play.

Mining for Indicator Gold

We have already listed a number of different threat intelligence feeds which can be used to search for specific malware files, command and control traffic, DNS request patterns, and other indicators. Mining for indicators isn't that much different from early gold prospecting. They were trying to find gold among millions of rocks moving down the stream. The main tool was a metal strainer — a filter.

The advantage we have today in security is that we can tune our filters to search through billions of 'rocks' at a pretty good clip. So you can search your security data infrastructure for a series of events and/or files within your environment — quickly and accurately to narrow down your focus to the most likely attacks.

Which brings us to the overhyped talk of using Big Data Analytics for mining attack indicators. Not that it's *all* hype. Innovative new technology for effectively indexing and searching huge security data sets is essential to finding advanced attackers quickly. But like its predecessor SIEM, Big Data vendors are particularly prone to hyperbole about the immediate value of their security analytics platforms.

We recently summarized how Big Data will impact security analytics:

> *We have every confidence that big data holds promise for security intelligence, both because we have witnessed attacker behavior captured in event data just waiting to be pulled out, and because we have also seen miraculous ideas sprout from people just playing around with database queries. In the same way hackers stumble on vulnerabilities while playing with protocols, security analysts stumble on interesting data just by asking the same question (query) different ways. The data holds promise. The mining of most data, and all of the work that will be required in writing M-R (MapReduce) scripts to locate actionable intelligence, is not yet here. It will take years of dedicated work — and it will take development against different data types for different NoSQL varieties.*

It still early days for Big Data technology to address these security problems. You are clearly constrained on internal capabilities (you will look for a lot of data scientists over the next few years), as well as by the immaturity of technologies such as Hadoop, MapReduce, Pig, Hive, and a variety of others for use in a security context. So remain skeptical about a magic "Big Data" box that ingests scads of security data and pops out geographic coordinates for advanced attackers.

Companies seriously trying to detect advanced attackers will be capturing packets (network full packet capture) to supplement the security data they already collect, and subsequently using Big Data technologies to mine it all. Sounds easy, right? Unfortunately it is thankless work, so make sure you swing by the cubes of your forensics folks to give them a big thank-you. They spend a lot of time chasing down false positives, all for those times they do find an active attack.

Verify the Alert

All our discussion so far has been about preparation for the main event, when an alert fires and it's time for your incident response process to kick in. But "advanced attackers" present some unique challenges. In particular, they devote significant resources and time to achieving their mission, which makes them difficult to deter — even if you successfully block one attack or stop an exfiltration, there will be more. A lot more.

This class of adversaries requires you to put a premium on analyzing malware to isolate the root cause of the attack, look for indicators to identify additional compromised devices, and then try to piece together a bigger picture of the attack.

If you weren't worried enough about this, remember that your perceived success as CISO is directly correlated to your ability to respond effectively to incidents and keep your organization out of the headlines.

React Faster and Better, CISO Style

Let's turn back the clock to review some of the Incident Response Fundamentals we introduced a few years ago. The process remains largely the same, but you are likely to need some of the data sources covered in React Faster and Better and analysis techniques presented in Malware Analysis Quant's process maps to deal with advanced attacks.

If you weren't worried enough about this, remember that your perceived success as CISO is directly correlated to your ability to respond effectively to incidents and keep your organization out of the headlines. You don't need a SIEM to make that specific correlation, by the way.

During the Attack

Once the alert sounds it is time to figure out whether the attack is legitimate, what it looks like, and the proper escalation path if necessary. The basic steps are:

1. **Gather information:** Before an investigator can make heads or tails of anything your first tier responders (typically the help desk) needs to collect some information. Who triggered the alert? What systems and devices were involved? Were you notified by a third party (not a good sign)? Could you find any related alerts (perhaps that were

ignored) around the time of the attack? You need a feel for whether it was an innocent operational failure or something designed to evade your defenses.

2. **Escalate:** Next you decide how far up the chain of command this needs to go — for now at least. If critical systems are involved (those on your list of things whose compromise would be *bad*), then your Spidey senses need to start tingling and you must involve the big guns both on the technical and management sides. You need to define and agree on escalation scenarios ahead of time so your first tier responders know what to do and when.

3. **Size up:** Once your second and/or third tier responders (and executives) are involved the key is to determine the scope of the situation. Was this a total compromise? Does extensive lateral movement indicate potential exfiltration? You need to know what you might be dealing with, and to assemble a list of the stuff you need in order to continue investigating.

4. **Initial Containment:** Depending on your initial assessment of the situation you may need to quarantine devices, step up monitoring, or remove devices' access to sensitive data — or likely all of the above. As with escalation the initial set of containment actions should be documented in a playbook with documented approval from all stakeholders to ensure containment is not held up by bureaucracy.

If the attack doesn't seem sophisticated or coordinated you can probably just wipe the machine and move on, hopefully using it as a teaching moment so the user doesn't do something stupid again.

At this point you should have initial defenses in place and a feel for whether your adversaries know what they are doing. If the attack doesn't seem sophisticated or coordinated you can probably just wipe the machine and move on, hopefully using it as a teaching moment so the user doesn't do something stupid again. Is it a risk to just wipe and move on? Of course! You give up any chance to analyze the attack in depth, but part of the CISO's job is to allocate resources to the stuff that matters. Being able to tell the difference between an advanced attacker, an operational failure, and a stupid user error is a key success determinant — as is resource allocation.

If there is a chance you are dealing with an advanced attacker (or something else is pushing you into a broader investigation), you will start working through a more detailed forensic process. That means quarantining affected devices, taking forensic images, and working to determine why the attack

succeeded. That requires you to dig into the malware and determine how the devices were compromised, then assess the extent of the damage.

Digging for the Root (Cause)

Malware analysis is a discipline all its own. We documented the entire process in Malware Analysis Quant, but we understand CISO types rarely fire up BackTrack or ship files up to malware analysis sandboxes. So here is what you need to make sure you can identify the root cause of a compromise.

1. **Build Testbed:** It is rarely wise to analyze malware on production devices connected to production networks. So your first step is to build a testbed to analyze what you find. This is mostly a one-time effort but you will always be adding to the testbed as your attack surface evolves. There are services that offer access to a testbed without the hardware investment.

2. **Static Analysis:** The first actual analysis step is static analysis of the malware file to identify things like packers, compile dates, and functions used by the program.

3. **Dynamic Analysis:** There are three aspects to what we call Dynamic Analysis: device analysis, network analysis, and proliferation analysis. To continue the investigation, observe the impact of the malware on the specific device. Here you seek insight into memory usage, configuration, persistence, new executables, and anything else interesting about execution of the malware. This typically involves running the malware in a sandbox. Once you understand what the malware does to a device, you can begin to figure out its communication paths. This includes command and control traffic, DNS tactics, exfiltration paths, network traffic patterns, and other clues to identify the attack. Finally you need to understand how the malware spreads, which we call proliferation analysis. You'll look at the kind of reconnaissance it performs, and try to find other clues that might indicate the malware running rampant in your environment.

4. **The Malware Profile:** Finally we need to document what we learned from the analyses, which we package up into a malware profile. The profile lists the indicators of compromise we won't shut up about.

Next you need to determine the extent of the damage, which means taking the technical indicators defined in the malware profile and mining your security data to see how many devices have similar characteristics (infections). Then you can start determining how deep the adversary is into your environment and whether you are dealing with an exfiltration.

As we described above, over the next few years Big Data Analytics will have a major impact on how we search for these indicators within our environments. But its short-term impact will be minimal. So in the meantime you need to use multiple tools (endpoint protection consoles, SIEM, network behavior tools, configuration management systems, vulnerability scanners, etc.) to search for specific indicators. Having to integrate the data from all these tools manually is a pain in the backside, but remember that priority #1 is to determine the root cause and extent of the compromise. As messy as that may be.

Having to integrate the data from all these tools manually is a pain in the backside, but remember that priority #1 is to determine the root cause and extent of the compromise. As messy as that may be.

Not to beat a dead horse, but all the stuff mentioned above is focused on figuring out how the malware in your environment works, what it does, and how to find it. Effectively leveraging threat intelligence in your security program can get you ahead of these reactive investigation fire drills by keeping a lookout for indicators *before* you have a widespread outbreak. In the worst case you will find compromised devices before you get called by the FBI or your payment processor. In the best case you will have mitigations in place to actually block the attack. That is why we are such fans of threat intelligence for mature security organizations.

Breaking the Kill Chain

Before you roll up your sleeves to fix things, you need to take a step back and assess the bigger picture, and decide whether you are dealing with a coordinated attack. We know it's hard to maintain sufficient discipline to look strategically and consider non-obvious links between various attacks rather than just fixing things and moving on. But in our experience playing "response whack-a-mole" can be very dangerous when dealing with advanced attackers, because once they know you found them they generally burrow deeper into your environment with urgency to maintain presence. Deciding how best to break the kill chain, to provide the best opportunity for successful remediation, is non-trivial.

Playing "response whack-a-mole" can be very dangerous when dealing with advanced attackers, because once they know you found them they generally burrow deeper into your environment.

Let's work through the steps involved in breaking the kill chain, disrupting the attackers, taking countermeasures, and/or getting law enforcement involved. Incident response needs to be a structured and conditioned response. Work to avoid setting policies during firefights, even though it is impossible to model every potential threat or gain consensus on every possible countermeasure. What you can do is try to define the most likely scenarios and get everyone on board with appropriate tactics for containment and remediation. Those scenarios provide a basis for making decisions even in scenarios that don't quite match your models.

Contain the Damage

As we described in Incident Response Fundamentals, containment can be challenging because you don't exactly know what's going on, but you need to intervene as quickly as practical. The core imperative is very clear: *do not make things worse*. Make sure you provide the best opportunity for your investigators (both internal and external) to isolate and study the incident. Be careful not to destroy data by turning off and/or unplugging machines without first taking appropriate forensic images.

At a high level containment involves two main parts:

1. **Quarantine the device:** Isolate the device quickly so it doesn't continue to perform reconnaissance, move laterally within your network, infect other devices, or progress toward completing its mission and stealing your data. You may monitor the device as you figure out the best option for remediation, but start by making sure it doesn't cause further harm.

2. **Protect critical data:** One reason to quarantine is to ensure that the device cannot continue to mine your network and possibly exfiltrate data. But you cannot assume the compromised device you identified is the only one. Go back to the potential targets you outlined when you sized up the adversary and take extra care to protect the critical data most interesting to your adversary.

One thing we know about advanced attackers is that they generally have multiple paths to accomplish their mission. You may have discovered one — the compromised device — but there are likely more. So be extra diligent monitoring data access and egress points to be sure you still disrupt the kill chain in case of multiple compromises.

Investigate and Mitigate

Now you identify the attack vectors and determine appropriate remediation plans. You want to be sure to gather just as much information as you need to mitigate the problem (stop the bad guys), in a way that doesn't preclude subsequent legal or other action in the future. This includes trying to get a sense of the adversary before moving forward with the remediation. This is where leveraging the adversary intelligence gathered early in the process becomes critical. If you have an understanding of the tactics of the likely adversary, you are far more likely to be success in remediating the attack.

When it comes to mitigation you will set a series of discreet achievable goals and assign resources to handle them. Just like any other project, right? But when dealing with advanced attackers you have a few remediation scenarios to consider:

1. **Clean:** People also call this the Big Bang approach because you need to do it quickly and completely. If you leave the attacker with any foothold in your environment you will start all over again — sooner rather than later. Most organizations opt for this approach — the more quickly you can clean your environment the better.

2. **Observe:** In certain cases, such as when dealing with an inside job or working with law enforcement, you may be asked not to clean all the compromised machines. In this case, you need to take extra care to ensure you don't suffer further losses while observing the attackers. That involves deep monitoring (likely network full packet capture and memory forensics) on critical data stores — as well as tightening controls on egress filters and DLP gateways.

3. **Disinformation:** Another less common alternative is to actively provide disinformation to adversaries. That might involve dummy bids, incorrect schematics, or files with tracking data which can help identify the attacker. This is a very advanced tactic, generally performed with the guidance of law enforcement or an elite incident response firm.

Executing the Big Bang

To get rid an advanced attacker you need to find *all* compromised devices. We have been talking about how to do that by searching for indicators of compromise, but you cannot assume you have seen and profiled all the malware they've used. Those pesky advanced attackers might even be throwing 0-day attacks at you. This, again, is where threat intelligence comes in — to look for patterns others have seen. Once you have identified **all** the affected devices they need to go dark at the same time. You cannot leave the adversary any opportunity to compromise other devices or execute a contingency plan to retain a foothold while you cleanup your machines. This probably entails wiping them down to bare metal — even if that means losing data. Given the capabilities of advanced attackers, you cannot be sure of totally eliminating the compromise any other way.

Do not underestimate the challenge of truly eradicating the adversary from your environment. We know of situations where the attackers used more than a dozen separate attacks on dozens of devices to ensure they maintain presence in the environment. If you miss even one, the attackers can maintain their foothold in your stuff.

You won't be perfect. At some point an advanced adversary will get back in — your job is to make sure they have to work for it and that you can respond effectively.

Once the affected devices are wiped and rebuilt you need to monitor them and capture egress traffic during a burn-in period to make sure you didn't miss anything. That means scrutinizing all configuration changes for indications that the attacker is back and finding victims again, as well as looking for command and control indicators. The moment the adversary is blown out they will start working double-time to get back in. *You are never finished.* You need to ensure your defenses have evolved to deal with these kinds of attacks. And no — you won't be perfect. At some point an advanced adversary will get back in — your job is to make sure they have to work for it and that you can respond effectively.

Advanced Attacker Complications

Incident response is hard enough. But when you factor in a well-funded, capable, and advanced attackers, things get more complicated. The first realization is that many decisions relating to mitigation, remediation, and prosecution are not up to security. These are executive decisions — the impact could ripple throughout the organization and might involve customer disclosure. This is why a team approach which encompasses all the important stakeholders is so important.

That doesn't mean you shouldn't offer recommendations and build a case to support your line of thinking. When deciding between cleaning the affected devices and observing the attackers, factor in the cost of complete and total cleanup and the likelihood of success. Many organizations (and senior executives) prefer to clean up but it is very difficult to keep advanced attackers out forever, so you are likely to be doing the same dance again — sooner than later. When advocating cleanup be sure you have details about which business operations might be disrupted, along with a realistic timeline for eradication. Your job longevity is directly related to how well you manage these expectations.

Another area complicated by advanced attackers is disclosure. If sensitive data was lost you will likely to need to make a best effort at full cleanup. It is hard to explain a decision to let attackers remain in your environment once they have stolen sensitive data. As with most security situations, PII (Personally Identifiable Information) changes everything.

Evolving the Security Program

The tactics described so far are very useful for detecting and disrupting advanced attackers — even if used only on a one-off basis. But you can and should establish a more structured and repeatable process, especially if you expect to be an ongoing target for advanced attackers. So you need to evolve your existing security program, including incident response capabilities. But what exactly does that mean?

Success requires empowering your folks to rise to the challenge of advanced attackers. This provides an opportunity for some staff to take on more important responsibilities and ensures everyone is on the hook to get things done.

You need to factor in the tactics you will see from advanced attackers and increase the sophistication of your intelligence gathering, active controls, and incident response. Change is hard — we understand. Unless you have just had a recent breach. Then, instead of budget pressure, you get a mandate to fix things no matter the cost, and you will face little resistance to changing processes to ensure success with the next response. Even without a breach catalyst you can make this kind of changes, but you will need some budgetary *kung fu* and strategic use of recent high-profile attack examples to make your point.

But even leveraging a breach doesn't necessarily result in sustainable change, regardless of how much money you throw at the problem. Evolving these processes entails not only figuring out what to do now and in the future — those are short term band-aids. Success requires empowering your folks to rise to the challenge of advanced attackers. But this requires management chops. You need to ensure your staff understands the additional effort expected and what's in it for them. More importantly, ensure you have a way to recognize them for stepping up. When managed correctly, dealing with advanced attackers provides an opportunity for some staff to take on more important responsibilities and ensures everyone is on the hook to get things done. Just updating processes and printing out new workflows won't change much without adequate resources and clear accountability to ensure change takes place.

Identify Gaps

Start evolving your program by identifying gaps in the status quo. That is easiest when you are cleaning up a breach because it is usually pretty obvious what worked, what didn't, and what needs to change. Without a breach you can use periodic risk assessment or penetration testing to pinpoint issues. But regardless of the details of your gaps or how you find them, it is essential that you, the senior security professional, drive process changes to address those gaps. *Accountability starts and ends with the senior security professional — with or without the CISO title.* Be candid about what went wrong and right with senior management and your team, and frame the discussion in terms of improving your overall capability to defend against advanced attackers.

Intelligence Gathering

A key aspect of detecting advanced attackers is building a repeatable intelligence gathering program to follow what is happening out there. Benefit from the misfortune of others, remember? Larger organizations tend to formalize an intelligence group, while smaller entities need to add intelligence gathering and analysis to the task lists of existing staff. Of all the things that could land on a security professional, an intelligence research responsibility isn't bad. It provides exposure to cutting-edge attacks and makes a difference in your defenses, so that's how you should sell it.

Once you have sorted organizational structure and accountability for intelligence gathering you need to focus on integration points with the rest of your active (defensive) and passive (monitoring) controls. Is the intelligence you receive formatted to integrate directly into your firewall, IPS, and WAF? What about your SIEM or forensics tools? Don't forget about analyzing malware — isolating and searching for malware indicators is essential for detecting advanced attackers. More sophisticated and mature environments should evolve beyond just searching for technical indicators of compromise and include proactive intelligence gathering about potential and active adversaries as described earlier. If you don't have those capabilities internally, which of your service providers can offer them, and how can you take advantage of them?

The key to information sharing networks (aside from trust) is reducing the signal-to-noise ratio — it is easy for active networks to generate lots of chatter that isn't relevant to you.

Finally you need to figure out your stance on information sharing. We are big fans of sharing what you see with folks like you (same industry, similar company size, geographical neighbors, or whatever criteria you like) to learn from each other. The key to information sharing networks (aside from trust) is reducing the signal-to-noise ratio — it is easy for active networks to generate lots of chatter that isn't relevant to you. As with figuring out integration points, you need accountability and structure for collecting and using information from these networks.

Tracking Innovation

Another aspect of dealing with advanced attackers is tracking industry innovation and the new technologies and services coming to market. We have done considerable research into evolving endpoint controls and network-based advanced malware detection, and the application of intelligence (Early Warning, Network-based Threat Intelligence, and Email-based Threat Intelligence) to understand how these technologies can help you.

As always, someone needs to be accountable. Who in your organization will be responsible for evaluating new technologies? How often? You might not have budget for all the latest and greatest shiny objects to hit the market but you still need to know what's out there, and you might need to find money to buy something that solves a sufficiently serious problem.

We have seen organizations assemble a new technology task force, comprised of promising individual contributors in each of the key security disciplines. These folks monitor their areas of expertise, meet with innovative start-ups and other companies, go to security conferences, and leverage research services to evaluate new technologies. They meet periodically to present their findings. Not just what the shiny object does, but how it could change what the organization does and why that would be better. They need to go beyond parroting back what vendors tell them to figure out how to apply that capability to existing control sets.

Evolving DFIR

A key aspect of detecting advanced attackers is digital forensics and incident response (DFIR). You need to ensure responders have adequate tools to determine what happened and analyze attacks. So revisit your data collection infrastructure, and consider capturing more detailed information at both network and device levels. Evaluate full packet capture technologies and possibly endpoint forensics. We are evolving the security program — it is not only about selecting and deploying tools, but also how will they will be used in your program and who will be responsible for deploying and managing them.

DFIR tools are just a means to an end. More important is how your incident response process is evolving to handle new attacker capabilities. Do you need to procure a sandboxing capability and build a malware analysis testbed? What kinds of organizational changes are required? Do you need multiple playbooks for different adversaries? For financial fraud you would deal with a predominately finance-driven oversight team. But an intellectual property risk would warrant CEO involvement. Those are just examples — your CEO might well want to be hands-on with every incident.

There are no right or wrong answers, but you need to be sure to ask the right questions, evolving every aspect of your security program.

Summary

Advanced attackers are not fundamentally different from the adversaries you have been dealing with for years. But advanced attackers are better at what they do and their greater capabilities cause additional stresses. You need to be ready, and no shiny object or widget from a vendor or service provider can get you there. To resist advanced attackers you need to evolve your security processes, increase the capabilities of your staff, improve the technical controls in place for deterrence, and ensure your incident response function can detect and remediate an advanced malware attack. But not necessarily in that order.

Intelligence is key in this battle — accessing new indicators of compromise and for tracking your adversaries. The default mode for a security professional is to not talk about what they are doing or why, but advanced attackers are rendering that attitude hopelessly behind the times. We all need each other's wisdom, experience, and mistakes to protect against these new attackers.

Don't expect to win every battle — work to win the war, which means keeping the really important stuff safe, and continually improving to shorten the window between exploitation and detection.

But leveraging intelligence to focus on the most likely attacks isn't a silver bullet. You still need to look critically at your existing controls and keep abreast of the latest innovations hitting the market to deal with advanced attacks. We recommend strong skepticism — especially anything that sounds too good to be true. Weigh new technology options against your existing controls. If there isn't a clear and measurable reduction of attack surface and risk, there is no reason to change. Advanced attackers particularly stress forensics and incident response capabilities. You need the ability to analyze malware or have it analyzed, as well as to search for indicators within your environment, quickly.

Finally, never forget you are in an arms race. Once you implement new defenses your adversaries will adapt their tactics to evade them. And the game goes on and on. Don't expect to win every battle — work to win the war, which means keeping the really important stuff safe, and continually improving to shorten the window between exploitation and detection. In our book, that's a win.

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

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/> 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/>.