F-Secure

# FIVE IMPERATIVES FOR ADVANCED CYBER SECURITY

What it takes to detect and respond to advanced cyber threats

CYBER SECURITY INSIDER – EBOOK 3/3

# ADVANCING THE STATE OF YOUR CYBER SECURITY

# ADVANCING THE STATE OF YOUR CYBER SECURITY

In the simplest possible terms, advanced cyber security is about stopping most breaches from happening, rapidly detecting the ones that do happen and then moving as quickly as possible to minimize their impact.

It's a far cry from the primarily defensive approach to cyber security that most companies are used to and set up for.

Instead, it's a proactive and dynamic discipline that's constantly evolving – just like the infrastructure you're protecting and the threat landscape you're protecting it from.

Getting this right comes down to two things: increasing the relevance of the information available to your team and reducing the amount of time spent between breach, detection and remediation.

In this eBook we'll be taking you through what we've found to be the five most important measures for advanced cyber security:

1. Collecting and managing forensic evidence
2. Using analytics to detect anomalies
3. Cutting through the noise with human expertise
4. Learning from threat intelligence
5. Creating (and rehearsing) incident response plans

We'll explain why these measures are so crucial and how to implement them in the real world.

**Let's start.**

## CATCH UP ON THE REST OF THE CYBER SECURITY INSIDER SERIES

In case you missed them, check out the last two eBooks from this series.

In the first eBook, 'The Chaos of a Corporate Attack', we looked at how one company was breached and what the impact of the breach was.

In the second, 'How CISOs deal with advanced cyber threats', we looked at the steps (both good and bad) European CISOs were taking to protect their companies.

# COLLECTING AND MANAGING FORENSIC EVIDENCE

# COLLECTING AND MANAGING FORENSIC EVIDENCE

When a breach is detected, the first step is a forensic investigation to figure out what the attacker did, what systems he had access to, how he got in, and how long he spent in the network. That takes evidence.

The better the evidence, the easier and quicker the investigation goes. More important, it's only once you've collected accurate forensic information that you can start to formulate an appropriate cleanup plan.

So evidence is key. But having the right evidence means having access to a long-term history of relevant events collected from endpoints, servers, and the network. And that can be challenging for a couple of reasons:

There's no definitive list of what evidence to collect. That's because the type of events you'll want to collect and store will depend largely on your infrastructure and the type of attacks you expect to be hit with.

At the same time, you can't possibly log everything either. Network bandwidth, application performance, storage constraints and the sheer complexity of processing huge volumes of data make 'store it all' an untenable strategy.

So you need to pick your sources wisely.

## LIVING OFF THE LAND

Most companies fail to detect breaches quickly enough. In fact, in most of the cases we've worked on, the intruders that gained access to the corporate network remained undetected for months before finally being hunted down.

# COLLECTING AND MANAGING FORENSIC EVIDENCE

## A PROCESS FOR PICKING THE RIGHT SOURCES

Some sources of information – like system logs, web logs, and firewall logs –are obvious choices for collection. But when it comes to picking which other events you should track, it helps to create a feedback loop.

In this process, your security experts perform an analysis of the situation, use that analysis to determine what data to collect, and then use that data to perform a subsequent analysis of the same situation.

This cyclical approach, which we recommend repeating on a regular basis, will allow your security experts to evaluate their coverage, check that they're gaining visibility into all levels of your infrastructure and see if they're collecting data relevant to the types of attacks you should expect. We've found it takes two full-time experts to run this process properly.

# COLLECTING AND MANAGING FORENSIC EVIDENCE

## BEST PRACTICES FOR SMARTER DATA COLLECTION

Most of the choices you make around evidence collection will need to be specific to your company and infrastructure. But there are some pieces of advice around data collection that are broadly applicable.

- Log outbound connections. It's surprising how many companies only log incoming connections. Outgoing connections can provide clues about command and control server connections and data exfiltration activities, as well as potential insider threats. Of course, your ability to do this may be constrained by the volume and type of network traffic your organization experiences. Bonus tip: remember to log IP addresses and not just reverse-DNS names.

- Collect access logs. This should include detailed logs from domain controllers and other authentication servers. These are extremely useful for forensic investigative work.

- Look out for Network Address Translations (NAT) routers. If you aren't aware of any NATs inside your organization, you may find that some of your system logs show all inbound traffic originating from a single IP address. As you'd imagine, this makes forensic investigations a lot harder.

- Log data during incidents. If you notice (or suspect) an attack is underway and your logging is broken or inadequate, all is not lost. Just fix the problem immediately. This is actually a great time to turn the dial up on what you're collecting and logging. The more data you can collect during an incident, the better. Once you're in the clear, you can always scale your data collection back down.

- Perform regular drills. Not only will you need to do this to check you're logging everything you need to, it'll also ensure you know where to find all the data you've been collecting when something goes wrong.

- Keep up with the threat landscape. Learn what attackers are doing. And make sure the data you're collecting is relevant to the type of attacks you're expecting to see

# COLLECTING AND MANAGING FORENSIC EVIDENCE

## PROTECTING YOUR EVIDENCE

Intruders will invariably try to cover their tracks. Often, this will mean tampering with the evidence you need to run a forensic investigation. So once you start collecting the right evidence, aim for a safe, tamper-proof evidence storage strategy and an adequate retention policy.

Collecting logs in a centralized storage platform, such as a Security Incident and Event Management system (SIEM), is a smart way to protect against tampering.

But whatever central location you choose to use, what's important is that you collect and store logs from hosts and devices across your network. And that you then retain those logs for a long enough period of time to ensure you have access to a copy of that data, should you need it for forensic evidence.

We recommend the following practices:

- Aim for a data retention policy of at least two years.

- Make sure you're synchronizing the time across all machines on your network (via NTP). And if your organization's spread across multiple geographic locations, make sure you're aware of time zone differences.

- Ensure your data storage has proper access controls surrounding it to reduce the chances of tampering.

- Make sure you have a solid, working backup policy in place for all collected and stored data.

# USING ANALYTICS TO DETECT ANOMALIES

# USING ANALYTICS TO DETECT ANOMALIES

Once you've defined a set of events and data worth collecting (and built the infrastructure to aggregate it), you can start the process of hunting for anomalies, signs of intrusion and suspicious behavior.

Ideally, you'll want to process incoming data on-the-fly. But you'll also want to process historical data against new Indicators of Compromise (IoCs) and rule-sets as dictated by changes in threat intelligence.

It's smarter to process large volumes of data like these with statistical analytics, rather than going line-by-line. So you'll need automation and big data analytics to perform these sorts of analyses well.

We've found that using freely available engines (such as Graylog and ELK Stack) makes this process far easier than attempting to build your own technology stack.

## THREAT INTELLIGENCE

Threat intelligence is a loosely defined term used to represent all sorts of information ranging from high-level political and strategic advice to streams of technical data designed for machines to process.

For the purposes of this eBook, when we talk about threat intelligence, we're referring to knowledge of the threat landscape and the tactics, techniques and procedures (TTPs) used by attackers, as well as low-level technical data such as IoCs.

# USING ANALYTICS TO DETECT ANOMALIES

When it comes to seeking out anomalies and malicious behavior, we'd recommend starting here:

Correlate collected data against threat intelligence feeds. Obvious feeds to start with include known-malicious IP addresses, domains, and file hashes.

Look for anomalous access management log patterns. Examples might include a single non-admin user attempting to log into multiple servers, or one machine attempting to log into a server under multiple different accounts.

Look for activity that might point to an intruder using brute force methods. This might include hundreds of thousands of login attempts in a suspiciously small time frame.

Seek out patterns of unusual user behavior. An example of this would be an SSH connection originating from a non-technical user's machine.

Flag activity that appears to happen at odd times. Such as in the middle of the night, on weekends, or during public holidays.

Look for signs of tunnelling software. A commonly used example is called Teredo.

The advice listed above is just a subset of all the possible ways to detect anomalies. But we've found these methods are particularly useful in the current threat landscape.

Of course, over time, you should continue to tune your own analytics models for the sorts of threats you expect to face based on a deeper understanding of the tactics, techniques and procedures (TTPs) used by attackers.

# CUTTING THROUGH THE NOISE WITH HUMAN EXPERTISE

# CUTTING THROUGH THE NOISE WITH HUMAN EXPERTISE

There are some incredibly sophisticated analytic methods for detecting anomalies. But they are by no means a viable replacement for real human intuition. Because no matter how well configured your analytics tools are, they will produce noise – false positives and unnecessary alerts.

Configure the rules to be too aggressive and a real incident might end up buried in noise. Configure the rules to be too relaxed and your tools might not flag suspicious activity at all.

So effective cyber security depends on the use of both analytics and human security experts.

This is perhaps the most important fact of advanced cyber security. Because there's a whole range of potential attacks on your infrastructure. And no automated tool can be prepared to tackle all of them without human expertise and intuition to guide your security strategy.

## THE BROAD SPECTRUM OF ADVANCED CYBER THREATS

On one end of the spectrum, there are uncoordinated attacks by cyber criminals that threaten your end-users with ransomware, bots and other crime ware and poison your externally facing web servers with exploit kits.

These aren't the most sophisticated attacks. But they can leave your organization exposed. For instance, it's not uncommon for more organized attackers to purchase access to compromised machines on corporate networks from other criminal organizations. It's an easy and cheap way in.

On the other end of the spectrum, you're also susceptible to advanced threat actors who'll target and access your internal network. These attacks can be as well planned as a military campaign. And you need experts who understand the details of how these attacks work to detect them.

In these cases, attackers usually get in through vulnerable software, SQL injection attacks or 'social engineering' tricks that get users to run a Trojan on a corporate machine.

# CUTTING THROUGH THE NOISE WITH HUMAN EXPERTISE

Once in, these intruders can map out your network and systems in search of data or opportunities for lateral movement, while laying low for weeks. They'll even install rootkits, destroy evidence of their presence, steal accounts and start to harvest your data. From there they can exfiltrate data using techniques that are hard to detect because they mimic normal user behavior.

The point of all this isn't to spook you. It's to explain the practical reality that sophisticated attackers know how to evade common detection methods. So it takes a combination of well-configured analytics tools and the keen, trained eyes of human experts to catch them.

## ONCE THEY'RE IN, THEY HAVE A NUMBER OF WAYS TO GET DATA OUT

Advanced exfiltration techniques include wrapping data in encrypted packets or password-protected zip files, using social media or webmail sites for uploads, or even use steganography (the practice of concealing secretive data in non-secretive text or image files) or peer-to-peer file transfer.

# LEARNING FROM THREAT INTELLIGENCE

# LEARNING FROM THREAT INTELLIGENCE

As we mentioned earlier, 'threat intelligence' is a pretty broad term that can be used to describe everything from strategic political advice to more tactical information on IoCs that's designed for machines.

## TECHNICAL THREAT INTELLIGENCE

You can buy technical threat intelligence feeds from a variety of third parties. These are ideally used as inputs to correlation rules in SIEM systems or other event analysis systems. Technical threat intelligence feeds often have short lifetimes so you need to keep refreshing them and running them against collected data in order to make sure they're still valid.

You can use these feeds in a variety of ways. For instance, a feed containing a list of malicious IP addresses can be parsed and used to automatically configure firewall or Intrusion Detection System (IDS) rules. Experts can, and often do, combine multiple technical threat intelligence feeds to create new types of intelligence that can be applied to your collected data.

Given the number and scope of technical threat intelligence feeds available, it makes sense to carefully select data that's relevant to your organization's specific needs. This is where tactical threat intelligence comes into play.

# LEARNING FROM THREAT INTELLIGENCE

## TACTICAL THREAT INTELLIGENCE

Tactical threat intelligence educates your security experts about the threat landscape. This sort of intelligence can be gleaned from a number of different sources like news feeds, social media, RSS feeds, chat forums and mailing lists.

These kinds of sources are great because they ensure your experts regularly gain knowledge and experience that can then be used to create rules, automation, and analytics models to detect potential threats. In fact, it actually helps them pick which technical threat intelligence feeds to source.

Tactical threat intelligence provides important context. It describes the 'who', 'why' and 'how' you need to understand and properly prepare against modern threats.

## UNDERSTANDING YOUR ATTACKERS

At the end of the day, you can only derive actionable threat intelligence from a deep understanding of how attackers operate. Good attackers understand the defenses they're going to face. They're stealthy, they know how to evade common detection methods and they understand how companies operate.

So if you're going to mount a solid defense, you and your team need to know as much about them, their methods and their motives as you can.

One way to look at threat intelligence is as a series of feeds you can buy and plug into your SIEM. But used well, they're a whole lot more. They're tools for education. It's about gaining the mindset and skillset of an attacker and using that knowledge to counter them.

### THINKING LIKE AN ATTACKER

To really get under the skin of the kind of people who attack corporate networks, it helps to know some of those people (or at least know people who know them). It's why we work so closely with people who've been on the other side.

# CREATING (AND REHEARSING) INCIDENT RESPONSE PLANS

# CREATING (AND REHEARSING) INCIDENT RESPONSE PLANS

In almost all cases, when a company realizes they've been breached, the responding team's instinct is to quickly block the attacker. That might mean taking a machine offline, closing down an account or adding a firewall rule.

But while these quick fix, knee-jerk reactions sound good at the time, all they end up doing is alerting the intruder to the fact that you're onto them. Invariably, the attacker stops what they're doing, lays low until things calm down – and then gets back to work.

A better approach is to observe the intruder, and carefully investigate the situation as thoroughly as possible without alerting them. By doing this, you'll gain insight into what they did, how they did it and why, while also being better placed to catch them in a way that blocks them completely.

We mention this because without a clear and coherent incident response plan, the chaos of the moment dictates how everyone reacts. That's a problem at the best of times and potentially disastrous in the middle of an incident.

Incident response plans are one of those things you never want to use. But when the time comes to use them, you'll really wish you had one. So here are some pieces of advice on what it takes to get incident response plans right.

# CREATING (AND REHEARSING) INCIDENT RESPONSE PLANS

Plan for the most relevant attack scenarios. By studying the threat landscape and reading about incidents other companies have faced, you should be able to assemble hypothetical attack plots that can then shape your incident response plans.

Define and train an incident response team. Make sure you know who's meant to do what, identify how communication channels should work and identify any secondary resources you might need for things like forensic investigation.

Practice your plans. Simulate different incidents on your scenario list. Make sure your plans work when they're run through. Record the amount of time it would take you to recover from the scenario. And adjust your plans depending on the results. We'd also highly recommend testing your plans with real attacks.

Prepare communications channels. Know who you need to notify about what, and when if something went wrong. Prepare a list of contacts and make sure it's kept up-to-date. Think about what you might need to tell your customers. In the chaos of a real incident, you'll have very little time to communicate detailed, well-groomed statements. So decide on the minimum communications needed and then plan for them.

Prepare "emergency first aid" guidelines for admins and staff. And then test them on those guidelines periodically. Train your response team on how to dive into forensics once they hit the scene. Train your whole staff to be aware of attack vectors relevant to their positions.

Make sure you have centralized logging set up. Make sure it works. Test it to make sure it's recording the amount and types of information you'll need for an investigation later.

Don't pull the power plug on any device you think is compromised. You're going to want to get memory forensics from affected devices so keep them going.

Train your staff to disconnect the network from compromised devices. This includes disabling Wi-Fi, Bluetooth, and mobile broadband connections.

Secure the logs from a compromised machine. Ideally, your centralized logging mechanisms should be doing this anyway. But it helps to train certain staff about how best to respond to a potentially compromised machine. Make sure they preserve logs from surrounding systems such as firewalls and active directory. This will help you verify how widespread the breach is.

# CREATING (AND REHEARSING) INCIDENT RESPONSE PLANS

Don’t trust anything on the compromised machine. Common tools are often “trojanized” and may give false readings.

Make sure your network diagrams are readily available and up-to-date. This speeds up response rates a whole lot more than you’d think.

Reset passwords and authentication. If you suspect that user passwords or other authentication factors were used in the attack, or if things like password hashes have been leaked, reset the passwords of affected users.

Figure out the root cause before you re-create breached systems. When the incident is over, you might be tempted to just re-create breached systems from backup. But without a clear understanding of the root cause, it’s very likely that the attack vector will still exist and the machines might soon be compromised again. So it’s crucial you understand the root cause, mitigate against it and then get the system back online.

Incident response is an essential aspect of modern crisis and risk management. And the best plans come with C-level buy-in.

This might sound like the kind of thing executives would roll their eyes at. But the reputational and professional implications of dealing with a breach have meant the smartest companies just want to get this kind of thing right.

If you’re still reading this eBook, you’re likely working for one of them.

## THE RESOURCES NEEDED TO RESPOND

For most organizations, it’d take a team of about ten trained experts to launch an appropriate investigation into an incident.

Since it doesn’t make sense to staff that many experts for every organization, it helps to know where you can find the external experts needed to handle incident escalations.

# TACKLING ADVANCED THREATS HEAD ON

# TACKLING ADVANCED THREATS HEAD ON

The attackers and tools your corporate infrastructure's up against have never been more sophisticated. So detecting them, responding to them and investigating them takes a deep understanding of how they work.

Of course, this is further complicated by the fact that attackers know how to evade prevention and detection measures and have their own in-depth knowledge of companies and their tactics. Which is why modern cyber security can't rely on static, defensive measures only.

It takes the combination of human expertise, analytics, and a constantly growing pool of relevant data and threat intelligence to get this right.

The good news is that the information, technology and skill needed to detect anomalies and react to them rapidly already exist.

So we hope you follow the tips and advice we've shared in this eBook. And crucially, we hope you won't stop learning about how to improve your approach to cyber security. It's just too important to ignore.

# WE'RE F-SECURE

And we've been a part of the security industry for over 25 years. It's why we've become a trusted advisor to both industries and EU law enforcement agencies across Europe.

In fact, we've been involved in more European crime scene investigations than any other company on the market.

Our Cyber Security Services help companies react faster, learn more and respond more intelligently to threats and breaches of all sizes. So if you're one of the smart ones and you're getting serious about cyber security, we should talk.

## THE CYBER SECURITY INSIDER SERIES

Read 'The Chaos of a Corporate Attack' to find out how one company was breached and how it impacted them.

Read 'How CISOs deal with advanced cyber threats' to find out what European CISOs are doing to protect their companies (and what they shouldn't be doing).