F-SECURE RAPID DETECTION & RESPONSE SERVICE

F-Secure

F-SECURE RAPID DETECTION & RESPONSE SERVICE

“A LOT OF THE ATTACKS WE’RE SEEING NOWADAYS AREN’T ‘ADVANCED PERSISTENT THREATS’, THEY’RE SIMPLE HACKS PERFORMED BY ‘ADEQUATE PERNICIOUS TOERAGS’.”

Dr. Ian Levy
GCHQ

CYBER SECURITY EXPERTS WATCHING OVER YOUR ENVIRONMENT 365/24/7

MAX 30 MINUTES FROM DETECTION TO RESPONSE

IMMEDIATE RETURN ON INVESTMENT € AS A TURNKEY MANAGED SERVICE

From the field

Over the last few years, you’ve probably heard phrases such as “the tactics, techniques, and procedures crafted by highly resourced threat actors are falling into the hands of less skilled adversaries”. That’s long speak for “expect a lot more script kiddies to start pwning your systems”. As Dr. Ian Levy from GCHQ recently pointed out, a lot of the attacks we’re seeing nowadays aren’t “Advanced Persistent Threats”, they’re simple hacks performed by “Adequate Pernicious Toerags”.

Nothing illustrates this phenomenon better than the group we’ve dubbed “The Romanian Underground”. This is a group that we have had first-hand experience with on a number of occasions while performing incident response and forensics work.

The Romanian Underground are, simply put, a bunch of IRC chat room buddies who decided it would be cool to take up the hobby of “hacking”. Most of these kids, upon joining the collective, have little to no Unix skills to speak of. They probably know about five commands in total. Newcomers are taken

F-SECURE RAPID DETECTION & RESPONSE SERVICE

under the wing of a mentor who provides them with simple tools and training to get them started on their new hobby. These mentors are almost as unskilled as the newcomers - they probably know about five more Unix commands than their apprentices. But they've been in the game for a few weeks already, and have a wealth of experience.

As newcomers learn the ropes (which usually implies that they've learned to configure and use a couple of tools), they're promoted to mentors, and take on their own set of apprentices. This hierarchical model closely resembles popular pyramid selling schemes you might have had the misfortune to encounter. Of course, the guys involved in The Romanian Underground aren't looking to become millionaires by selling soap - the pyramid scheme is a form of gamification, where the goal is to collect as many owned systems as possible and move up the ranks.

Naturally, it's the guys at the top of the pyramid who really benefit from all of this. They're the ones providing the tools, and by pushing all their manual work downstream, they get access to thousands of compromised systems. Meanwhile, the newcomers are happy to proudly identify themselves as "hackers" on their Facebook pages (alongside other unrelated hobbies such as windsurfing or snowboarding).

The toolkits being pushed down the pyramid are usually designed to exploit or brute force common services such as SSH and webmail servers. What might surprise you (or not) is that these toolkits, in the hands of completely unskilled noobs, are being used to compromise even PCI-DSS-compliant organizations across the globe.

The Romanian Underground represent just one of many groups that form part of a growing trend of low-skilled hackers and cyber criminals. The motives of the masterminds behind these groups are, you guessed it, financial gain. Acquiring access to a large number of compromised company networks allows them to cherry-pick prime targets for cyber extortion and data exfiltration. And any company is a potential target.

The fact that these groups are able to compromise PCI-DSS-compliant organizations is a testament to the fact that purely preventative cyber security solutions simply aren't cutting it anymore. And the reason why so many companies are now being owned in this style is due to the fact that they simply don't have an ounce of visibility into post-breach activities on their networks.

That's not to say that skilled attackers aren't also out there. But, as a company that's been involved in more European cyber crime investigations than any other company in the world, we can tell you that there's no point in worrying about the NSA or APT28 until you know you can at least stop these guys.

"THAT'S NOT TO SAY THAT SKILLED ATTACKERS AREN'T ALSO OUT THERE. BUT, AS A COMPANY THAT'S BEEN INVOLVED IN MORE EUROPEAN CYBER CRIME INVESTIGATIONS THAN ANY OTHER COMPANY IN THE WORLD, WE CAN TELL YOU THAT THERE'S NO POINT IN WORRYING ABOUT THE NSA OR APT28 UNTIL YOU KNOW YOU CAN AT LEAST STOP THESE GUYS."

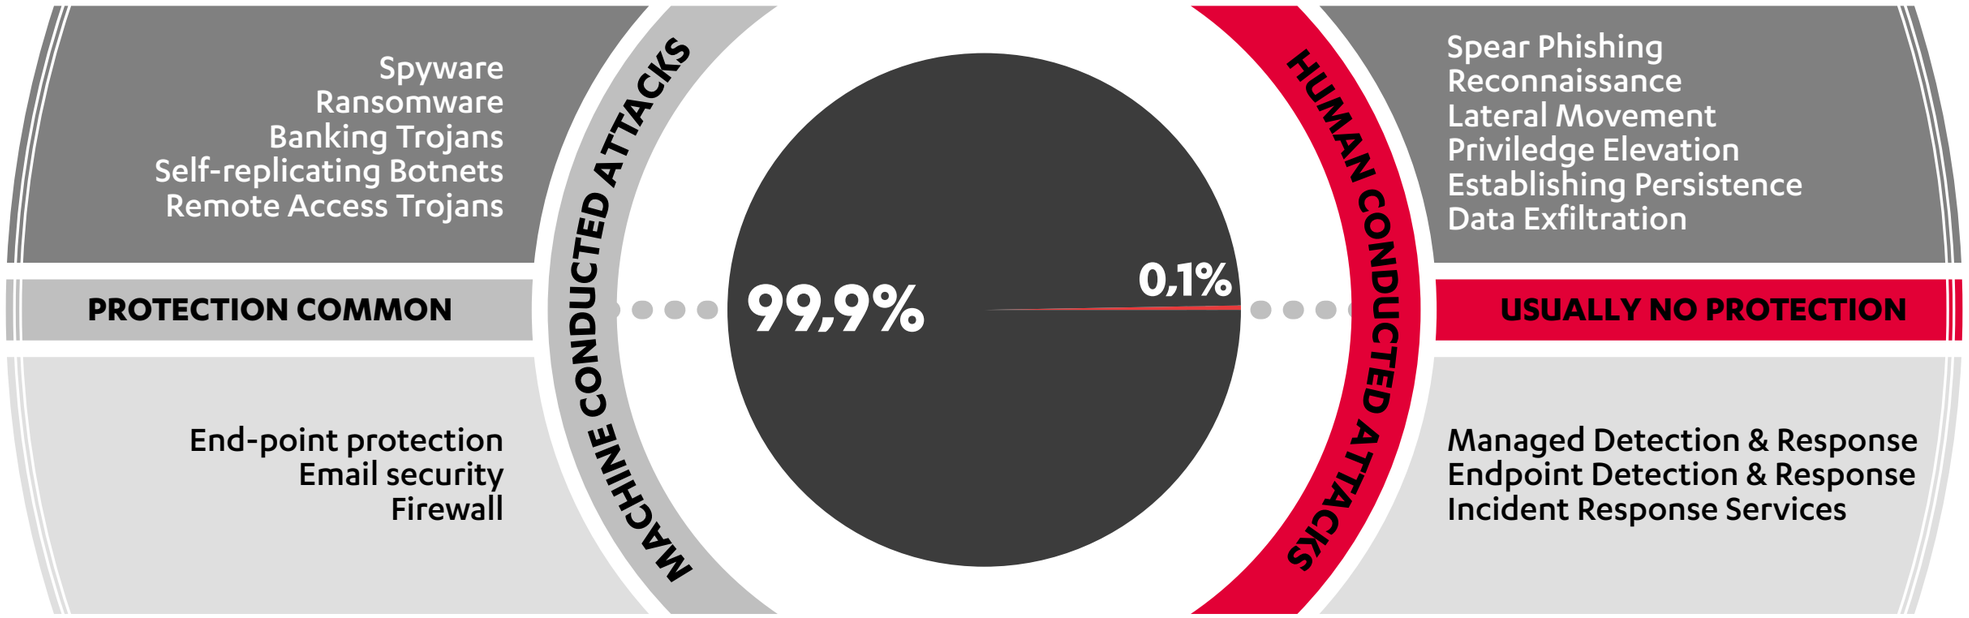

Commodity threats, and the solutions that protect against them are commonplace...

... But targeted attacks have the potential to be a lot more damaging. And most organizations aren't protected against those at all. Read on to learn more.

THE ANATOMY OF AN ATTACK

In our experience, most companies only discuss cyber security while on the broader topic of risk management. While performing risk analyses, companies identify threats or risks relevant to their organization and then prioritize them based on likelihood, impact, and cost to mitigate. When addressing cyber threats, we've noticed a potential disconnect between the risk that companies perceive and the reality of the situation. We'd like to help clear that up.

Sophisticated cyber attacks tend to start at the top and work their way down. It's the opposite of "low-hanging fruit". When new types of attacks are discovered, they're usually attributable to highly resourced threat actors (i.e. nation states). These adversaries, by default, go after the highest-value targets first. As the tactics, techniques, and procedures (TTPs) used in such attacks become public knowledge, they trickle down into the hands of less organized cyber criminals.

New TTPs first see use against governments, military targets, and defense contractors. Next on the ladder are usually banks and critical infrastructure providers (such as energy companies). The same TTPs then get used against heavy industry and finally, everyone else (manufacturing, retail, SMEs, etc.).

Threats to an organization aren't limited to attacks from the outside. Accidental and intentional leaks can and do originate from company insiders with enough access to critical or confidential assets. Upstream attacks, where a partner, supplier, or contractor are compromised by an attacker looking to establish a beachhead in an adjacent organization are also very common. In several incident response cases we've been involved with, even physical intrusion of a company's premises was used as part of the attack vector.

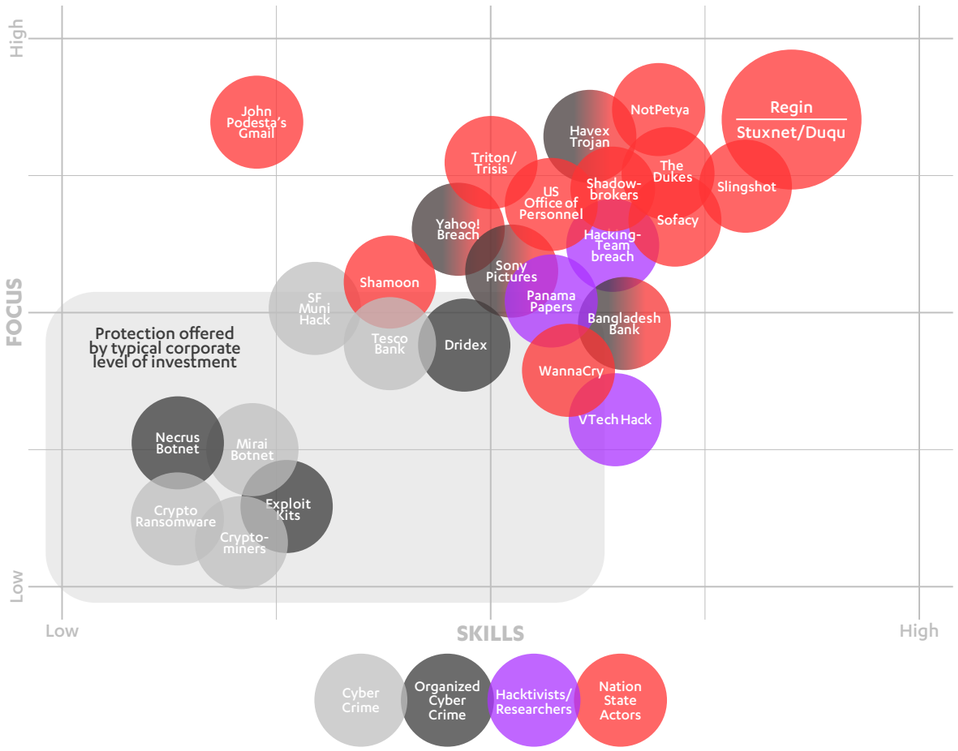

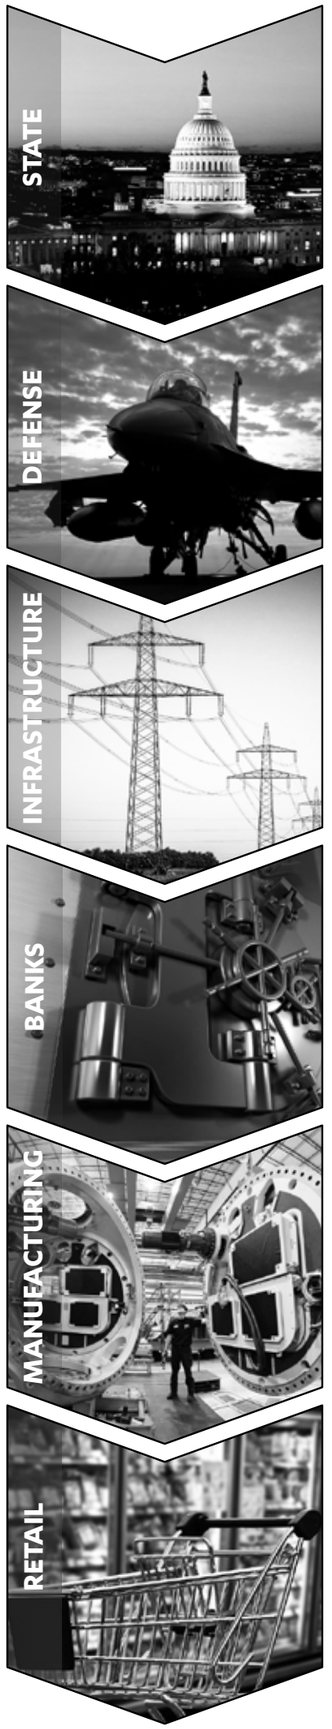

"WHEN GDPR AND NIS REGULATIONS ARE EFFECTIVE, IT WILL BE MANDATORY FOR ORGANIZATIONS TO HAVE AN ANSWER TO DATA BREACHES."

Cyber attacks come in many forms, ranging from commodity malware (such as ransomware) to highly skilled attacks performed by nation-state actors. We've broken these threats down into separate categories.

Commodity threats

Commodity threats are highly prevalent, and have been for decades. A company's chance of encountering commodity threats is, therefore, extremely high. However, due to their prevalence and long history, there are plenty of good software solutions available designed to protect against these threats. And these solutions work as intended. If a business is hit by a commodity threat (such as crypto-ransomware), the impact is usually fairly low. Most of the time it'll be blocked by endpoint protection software. If it does get through, there are two options – pay the ransom or fix the problem. Don't pay the ransom and a handful of staff will lose some productive work time. Pay and, most of the time, you'll get the data back. Ransom amounts are low by company standards. So, the likelihood of seeing a commodity threat is high, the impact tends to be low, and the mitigation cost is basically free (we assume you're smart enough to be running an endpoint protection solution already).

Cyber crime

Cyber crime represents the next category on our risk assessment scale. This category moves beyond the realm of commodity malware threats, and onto targeted attacks. Companies are selected as targets for various reasons. In some cases, a victim is chosen because they are "broadcasting" themselves via weak or vulnerable infrastructure. Other targets are selected simply because the attacker has taken interest in a particular organization, for one reason or another.

Cyber criminal attacks are often opportunistic - the attacker has an easy way in, sees an opportunity to make money, and takes it. Cyber crime is by-and-large financially motivated. Once the adversary has breached the target's network, systems or data will be held for ransom. We refer to this phenomenon as "cyber extortion". These types of attacks are very much on the rise, and can target organizations of any size, from SMEs to large enterprises.

We predict that the introduction of the NIS and GDPR regulations will further embolden cyber criminals and cyber extortion schemes. Once these regulations are in effect, companies may be more willing to fork over a ransom, in order to sweep the news of a breach under

the rug rather than face the expensive task of responding to and reporting the incident.

Cyber crime can be broken down into roughly two categories – organized and non-organized. Organized cyber criminal groups are very close, in terms of sophistication, to nation-state actors. The Bangladesh bank attacks of 2016 are a good example of organized cyber crime. Non-organized cyber criminals often run as lone wolves. They have less resources, and their skill can vary. The Romanian Underground falls into this category.

Two years ago, we'd have rated the likelihood of falling prey to cyber criminals as low. Today, the likelihood is medium, and on the rise. The financial and business impact of a targeted cyber crime attack can vary. In many of the cases we've responded to, ransoms demanded by non-organized cyber extortionists only ran into the tens of thousands of Euros - not a hefty sum for most organizations. But we don't imagine that any organization would simply pay the ransom and go about their business. The knowledge that an intruder is in their network is going to be enough to call in an incident response team to sort out the situation.

If an adversary manages to exfiltrate important data, the costs of a cyber crime incident can really start to skyrocket. This is especially true if customer data was involved. No matter what, a breach is most likely going to incur reputational, legal, PR, business, and internal productivity costs. And it's not anymore limited to protecting your business and its sensitive data, but regulatory bodies like NIS and European Union have made new requirements. For example, EU's General Data Protection Regulation (GDPR) requires organizations to be adequately prepared to detect, respond and report personal data breaches within 72 hours. We'd venture that it would be a good time to start thinking about that if you so far have not.

Nation state

Companies that worry about being targeted by nation-state attacks typically know who they are. They also know that defending against a nation-state attack is almost impossible. Regardless, they're forced to try (since they can't afford not to). The impact of nation-state attacks can vary from having top secret intellectual property stolen by overseas competitors or governments, to having your nuclear enrichment halted when centrifuges are destroyed.

"TARGETED ATTACKS DON'T CARE ABOUT YOUR 'NEXT GEN' PRODUCT, NO MATTER HOW SHINY THE VENDOR CLAIMS IT TO BE. TO BE BLUNT, THE SOLUTIONS THEY'RE SELLING ARE FIXING THE WRONG PROBLEMS."

DON'T BELIEVE THE HYPE

Getting to the point of why we presented this risk analysis - we've noticed that there's still a very strong marketing push towards endpoint protection solutions. We've seen "next gen" vendors claim that their solutions can prevent targeted attacks. Some even foolishly claim that breach detection is irrelevant, since it's already "game over" if a threat gets through perimeter defenses.

It's dangerously misleading.

Targeted attacks don't care about your "next gen" product, no matter how shiny the vendor claims it to be. To be blunt, the solutions they're selling are fixing the wrong problems.

Given this huge marketing push from "next gen", we're not really surprised to see that very few companies we've spoken to are aware of the need for breach detection and response capabilities. And therein lies the problem. Organizations are way too distracted to realize that they should start investing in breach detection and response, instead of another layer of protection against commodity threats (although the adversaries would love you to do this). Let's put it this way - would you rather have your next incident involve cleaning malware off a laptop in your sales department or dealing with a full-blown data breach?

But don't just take our word for it. Gartner predicts that by 2020, 60 percent of enterprise information security budgets will be allocated for rapid detection and response approaches , which is an increase from less than 30% in 2016. So, ask yourself this: how much of your budget have you allocated to breach detection and response right now? We're guessing it isn't close to 60 percent. In our experience, only 10% of companies we've talked to even had a budget allocated for breach detection and response.

BY 2020, 60 PERCENT OF ENTERPRISE INFORMATION SECURITY BUDGETS WILL BE ALLOCATED FOR RAPID DETECTION AND RESPONSE APPROACHES, WHICH IS AN INCREASE FROM LESS THAN 30% IN 2016.

**Gartner 'Special Report
Cybersecurity at the Speed of Digital Business'
Paul E. Proctor, Ray Wagner, 30 August 2016**

"ANY ORGANIZATION NOT RUNNING A BREACH DETECTION SOLUTION (OR NOT HAVING PERFORMED A RECENT INVESTIGATION) MUST, IN THIS DAY AND AGE, ASSUME THEY'RE IN A POST-BREACH STATE."

FROM A DEFENDER'S DILEMMA TO AN INTRUDER'S IMPASSE

Cyber threats are asymmetric in nature. An attacker only needs to succeed once to gain access to a network. Defenders must succeed one hundred percent of the time if they want to keep them out. You can't rely on being successful all the time.

And yet this is what most companies are doing. Traditional perimeter defense technologies, such as firewalls and endpoint protection software do a good job at what they're meant to do - namely detecting and blocking real-world and commodity threats. But you can't expect these solutions to stop advanced adversaries. Any adversary worth their salt will craft an attack designed to bypass those defenses. And they won't even need to use malware to gain a foothold in the organization (contrary to what you might have been told, skilled attackers rarely, if ever, use malware).

Cyber attacks commonly follow the same pattern. Attackers start by breaching the perimeter of an organization with spear-phishing, watering hole, or man-in-the-middle attacks. Sometimes attackers may gain entry by exploiting a vulnerability in a public-facing system, or even by purchasing access to an already compromised system. Once inside the perimeter, adversaries perform reconnaissance, elevate privileges (by exploiting misconfigured or vulnerable systems), hunt for domain admin passwords (using memory-scraping tools such as Mimikatz), and move laterally onto interesting systems. They'll often establish persistence using off-the-shelf RATs such as Orcus, Litemanager, or luminocityLink. They'll then exfiltrate data using subtle methods designed to mimic regular user behavior.

Most of the tools an attacker needs are built into the operating system itself. And attackers are adept at hiding from network-based IDS systems by hiding the command and control traffic. It's almost impossible to detect modern attack techniques simply by analyzing network traffic. In fact, there are too many ways for an attacker to hide. All of these techniques fly under the radar of traditional perimeter defenses such as firewalls, endpoint protection, and spam filtering.

In most cases, once a company has been breached, adversaries are able to act with impunity for as long as they wish. It's not uncommon for a company to find out they've been compromised from a third party (such as a CERT organization). In our experience in the field, on average the time between a breach happening and being discovered is 100 days. Think about that - it takes the majority of organizations months or even years to figure out they have been hacked.

Any organization not running a breach detection solution (or not having performed a recent investigation) must, in this day and age, assume they're in a post-breach state.

Breaches are becoming more and more commonplace. And this is because adversaries know that their targets have no idea they're being hacked. For many attackers, compromising systems is just as easy as a burglar walking into a house with the front door left wide open.

But here's something interesting - adversaries actually hate the idea of getting caught. And they hate operating in an environment where there's a chance they're being monitored. That's why most good attackers will exercise caution. Once inside a victim's network, a professional intruder will tread lightly, while constantly being on the lookout for signs that they've been detected.

Attackers know that a good defender won't react to signs of an intrusion in a panic – they'll watch the intruder, gather intel, and then act on the situation when they're good and ready. As a defender, successful detection and effective response will, in the eyes of the adversary, constitute a breach of their mission.

While it would seem that attackers have the advantage, there's actually a lot that defenders can do to turn the tables on them. Everything the attacker does is bound to leave a trail of evidence behind them. And while a compromised system may not be able to tell you when it's "owned", there's a chance that it logs some evidence. That evidence can be used to spot the intruder, or even travel back in time to reconstruct the adversary's movements.

Just imagine the frustration when an attacker realizes that every move they've made has been monitored, that they've exposed their entire toolchain, and that they've effectively been sent back to square one. Not a great feeling for the attacker. And a huge win for the defender.

This is, in our opinion, the best approach to cyber defense. Let's delve into how we achieve that goal.

"HIRING AND RETAINING CYBER SECURITY EXPERTS IS NOT EASY. IT IS ESTIMATED THAT, RIGHT NOW, THERE ARE AT LEAST TWO CYBER SECURITY JOBS FOR EVERY ONE PERSON WORKING IN THE FIELD. AND THIS PROBLEM IS EXPECTED TO BECOME EVEN MORE ACUTE IN THE FUTURE."

WHAT IS RDS?

F-Secure Rapid Detection & Response Service (RDS) is a managed breach detection and response service. What we mean by "managed" is that there's a minimal installation process on your side to get things up and running, and after that, everything from breach detection to response is handled by us.

We decided to take the managed service route after seeing the difficulties other companies had in building their own breach detection and response capabilities.

Of all the challenges that organizations face while building breach detection and response capabilities, nothing really compares to the difficulty they face when trying to hire and retain good cyber security expertise. We're lucky here at F-Secure - our line of work ensures that we already have plenty of in-house cyber security experts. We've been working with threats and building automation for decades. And we know how difficult it is to get things right. We realized that the best way to provide an unequalled breach detection and response service was not to build a product or rely solely on artificial intelligence, it was to provide both systems and expertise directly to our customers.

Hiring and retaining cyber security experts is not easy. It is estimated that, right now, there are at least two cyber security jobs for every one person working in the field. And this problem is expected to become even more acute in the future. The only way you're going to get valid data from an in-house IDS is by having experts on staff. The same goes for keeping up on threat intelligence, configuring systems, red teaming, and responding to incidents correctly. So, you're probably going to need more than one or two experts on your payroll.

RDS doesn't just provide human expertise, though. It's a service that's built on top of threat intelligence, sample analysis, and decision-making systems that have been developed in-house for over a decade. And while an organization could eventually develop their own in-house systems and expertise to the levels we've reached, it would take them a very long time.

OUR EXPERTS AT YOUR SERVICE 24/7

At the core of RDS is our Rapid Detection & Response Center, which is the base of operations for all of our detection and response services. At RDC, cyber security experts work on a 24/7 basis, where they hunt for threats, monitor data and alerts from our customer's environments, flag anomalies and signs of a breach, and then work with our customers to respond to real incidents as they take place.

RDC staff have access to our own in-house, world-class analytical tools, all of our threat intelligence data, and a wealth of information and knowledge from both our Cyber Security Services and F-Secure Labs organizations. In fact, all of these teams work closely in cooperation with each other.

Staff at our Rapid Detection & Response Center are trained to handle a variety of tasks. Their main tasks fall into roughly three different roles - threat hunters, incident responders, and forensics experts.

Threat hunters

Threat hunters are our first responders. They monitor the service and hunt for threats. When a threat hunter discovers something suspicious, evidence is collected to verify the incident. If a real incident is discovered, it is given a priority. High-priority alerts are generated when there's a strong indication of an ongoing breach, and in these cases, the customer is immediately contacted by phone. For non-critical cases, guidance is sent to the customer by email. Threat hunter also keep the customer up to date on any ongoing investigations.

Incident responders

Incident responders are assigned complex cases that customers are unable to handle on their own, and are usually sent out to assist the customer on-site. Incident response personnel can assist with a range of technical and non-technical response activities, depending on your needs. We are also familiar with collecting evidence for law enforcement purposes, should it be required.

Forensics experts

Forensics experts are specialists tasked with the most difficult of cases. We're one of the few organizations globally who can handle a very wide range of forensic tasks, ranging from internal network triage to deep reverse engineering of unique malware samples. This allows us to handle even the most complicated nation-state originated attacks.

FROM DIY TO ROI

Because expertise, monitoring, threat hunting, and response capabilities are covered by RDC, once you've decided to implement RDS in your organization, all you need to do is install simple sensors and devices on your network. The time from initial deployment and configuration to actual breach detection and response capabilities is less than a week. In fact, we've been told by several customers that we have the easiest system they've ever worked with.

The alternative to deploying a managed breach detection and response service is a lengthy (in most cases 3-5 years) and expensive (multi-million Euro) project of purchasing, deploying, and configuring dedicated systems, and hiring and training a sizeable staff. But RDS isn't just about a fast return on investment. We've seen many companies go to the trouble of building a SOC and setting up IDS and SIEM, only to still not catch threats. This is because, in our experience, finding actual threats is like finding a needle in a haystack.

To illustrate with a recent real-world example, in a 1300-node customer installation, our sensors collected around 2,000,000,000 events over a period of one month. Raw data analysis in our back end systems filtered that number down to 900,000. Our detection mechanisms and data analytics then narrowed that number to 25. Finally, those 25 events were analyzed by our staff at RDC, where 15 real threats were discovered (and verified by the customer).

The thing is, if you go with your own IDS/SIEM solution, it's your organization that will need to process those 900,000 events. And that's why we've gone to countless customer sites and found threats on their network, despite those customers already running very well-known IDS solutions. Combing through the noise and false positives is difficult, and can cause fatigue in even the most diligent of analysts.

In order to process this volume of events, you also need reliable, up-to-date threat intelligence. At F-Secure, we have our own in-house sources. And after over 25 years in the business, we also have a massive historical sample collection that even gives us the ability to find relevant threats left undiscovered from currently active threat actors. Our researchers do both threat intelligence investigations and reverse engineering. This gives us both high-level knowledge of the global threat landscape and in-depth technical knowledge of the threats themselves. Instead of studying each threat independently, we identify relationships between threats, allowing us to understand the capabilities and motives of an adversary. We focus on the puzzle and not just on the individual pieces.

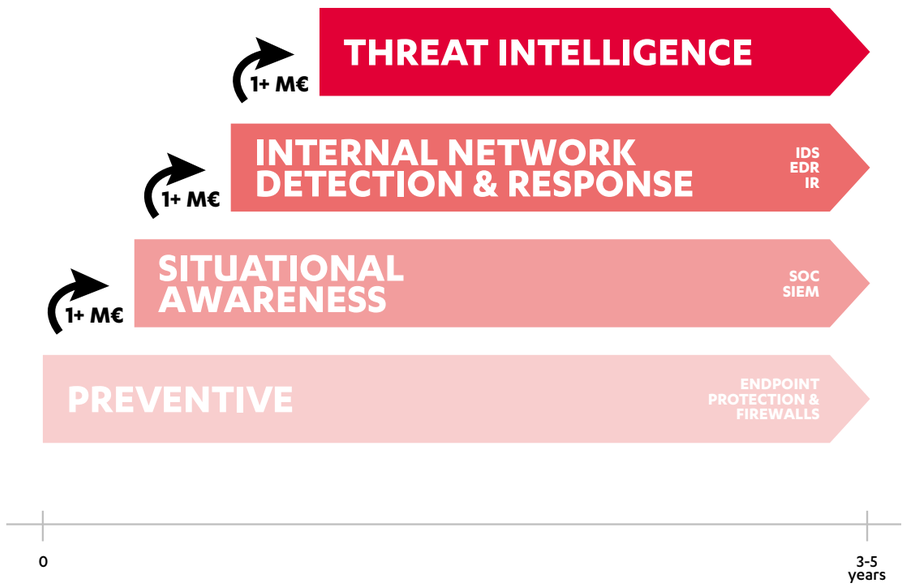

Building in-house breach detection and response capabilities is difficult

We've noticed that, for most organizations, setting up in-house breach detection and response capabilities tends to be a complicated, time-consuming, and expensive endeavor. There are multiple components that need deploying and configuring. All of them are expensive, so purchasing decisions take time and research. Different components may or may not interoperate well, so you have to figure that out, too. Then you need to select threat intelligence feeds, and there are dozens if not hundreds of those available. Deploying and configuring these systems is a complicated job. And at the end of all this, you'll be left wondering if you've got everything covered and whether or not all the pieces are talking to each other properly. And that's just the initial install. After that, systems, rules, and feeds need to be constantly improved and modified as the world changes.

Responding to a breach is usually also a lengthy and expensive process that requires expert data forensics and incident response work. A typical response scenario includes removing the adversary from the network, cleaning up or restoring affected systems, resetting compromised accounts, determining where the intruder has been, and determining what the intruder has done. Most companies don't have the in-house expertise or capabilities to perform these types of activities, and so must call on a third party to help.

Integrating your own systems

By the way, for organizations that have already invested in infrastructure such as SOC, SIEM, or IDS, our Rapid Detection & Response Service provides an additional layer of security that easily integrates into (via processes and APIs) and enhances any existing ecosystem. We have a REST API for detections and sensors, and we're building in more capabilities all the time. We also use a standard ticketing system that integrates easily into existing customer support processes.

''THE WEEKLY REPORTS FROM RAPID DETECTION & RESPONSE CENTER HAVE BEEN PROVIDING ME WITH GREAT COVERAGE OF THE LATEST SECURITY EVENTS. EVEN BETTER, I HAVE FOUND IT INTERESTING TO READ! I HAVE ALSO BEEN IMPRESSED BY THE RESPONSIVENESS OF RAPID DETECTION & RESPONSE CENTER. ''

Jukka Vallisto
IT Specialist
Amnesty International Finland

DETECTION AND DECEPTION

RDS was designed from the ground up to detect even the most skilled attackers using non-malware techniques, and to respond to those threats within a thirty-minute timeframe. This thirty minutes includes initial investigation, false-positive filtering, and prioritization. The goal is always to provide the customer with actionable guidance as part of the alert. Unlike many other solutions on the market, our customers always gain the ability to start incident response activities as soon as an anomaly is detected.

How do we achieve this?

We utilize a "detection and deception" approach that uses a combination of endpoint sensors and honeypots. Here's how it works:

The process starts when a sensor is installed in your network and starts looking for signs of compromise. Sensors collect and communicate events to our backend systems. This data is processed and matched against threat intelligence sources using user and entity behavioral analytics. Staff at our Rapid Detection & Response Center receive actionable alerts that make it through this automation.

Between the moment an RDC hunter receives an alert and the moment a customer is called, the hunter verifies the alert, decides on a priority, and determines what has occurred. Remediation steps are also formulated during this time. We then call the customer and advise them on how to respond to the situation.

Response actions are determined by the type of incident encountered. In the easiest of cases, one of our hunters can provide sufficient instructions over the phone. In more complex cases, we may need to send incident responders over to help. In the future, we expect to be able to automate more and more response activities.

When a breach is discovered, having access to historical data is the key to building a detailed post-breach event timeline. Since adversaries almost invariably wipe data to cover their tracks during an attack, having access to data that is stored off-premises means having a pretty much guaranteed tamper-proof source of evidence for incident response and forensic investigators. In the event of an incident, F-Secure Rapid Detection & Response Service helps the customer preserve any evidence that is essential in subsequent incident response actions.

RDS is also designed to look for the existence of newly discovered threats in historical data. Retrospective threat hunting is achieved when new detection algorithms are run against historical data collected from each of our customers. This mechanism is especially useful when dealing with attacks from more advanced adversaries (that may have gone hidden for some time).

RDS can be deployed during ongoing incident response work, and is used as a threat hunting service that can quickly gain visibility into a network that has already been breached.

Finally, RDS continues to work outside of the corporate network. In a world where the classical security perimeter is crumbling, traditional IDS approaches have become ineffective (since they typically only work on the edge of the network). These traditional approaches cannot track threats when devices are outside of the corporate network, or when people utilize cloud-based services. Our endpoint sensor approach solves this problem rather effectively. What's more, we've been working on extending RDS capabilities into cloud services, such as Salesforce.

Endpoint Sensor

F-Secure's Endpoint Sensors are lightweight, discreet monitoring tools designed to be deployed on all relevant Windows, MacOS, and Linux computers within your organization. Sensors are custom-configured for each organization and are easily deployed using standard IT remote administration tools. These components collect behavioral data from endpoint devices using well-documented mechanisms, and are specifically designed to withstand attacks from adversaries (we can detect tampering attempts against these components). Since they're just data collectors, they require very little maintenance.

Through a process of data normalization, we've managed to limit the amount of data upstreamed by each sensor to a few megabytes per day. Sensors have a low impact on performance and bandwidth utilization. By collecting such a small amount of data per day, we are able to save historical data for longer periods of time.

Note that endpoint sensors are also designed to function in Payment Card Industry Data Security Standard (PCI-DSS) compliant environments. Our sensors don't collect the sort of information that might jeopardize card-holder data, data is only ever relayed in one direction (from the endpoint to the back end), and it's not possible for human operators to directly interact with the sensors themselves.

Network and Decoy Sensors

Network and Decoy Sensors are designed to be deployed across your organization's network segments. Network Sensors analyze all connection attempts to and from your organization's network, and record selected network traffic, and analyze files that arrive on the systems. Data sent over the network reveals signs of potentially suspicious activity that otherwise would not be seen.

Decoy Sensors are honeypots working as an effective, low-noise method of identifying post-breach activity. Attackers typically perform a recon phase once they've gained access to a network (In order to identify easy targets for lateral movement, persistence, and privilege escalation). Network Decoy Sensors are designed to catch the scans associated with this sort of reconnaissance and provide easy targets for the attacker to focus on.

Any action the attacker performs on the active decoy is automatically detected and logged by our service. Furthermore, honeypots keep the adversary busy, reveal the tools they're using, and allow us to build a detailed base of forensic evidence, while the attack is in progress. We can actually observe adversaries "living off the land" by monitoring these environments.

Decoy Sensors emulate popular services including SSH, HTTP, and SMB, and are designed to mimic Windows servers, workstations, file server, and even VOIP servers. All connection attempts to and from network sensors are recorded, and any files that arrive on the systems are analyzed by F-Secure.

Reporting

You will be alerted whenever a real threat is flagged. With a dashboard you can stay on top of all alerts reported as suspected attacks. Actionable guidance provided by our service helps you respond promptly whenever under an attack, and the service helps you manage the verification process regarding less critical detections. The dashboard also provides continuous visibility into all installed sensors and hosts.

Privacy policy

All data collected from customer deployments is sent through secure, encrypted channels and stored on controlled, secured servers. Access to data is carefully restricted to authorized users and for authorized purposes only. All data is physically stored in Europe. We respect our users' privacy and our customers' need to protect sensitive data and corporate secrets. Data collected from one customer is never shared with other customers. You can find more information in our privacy and confidentiality policies, especially with regards to data handling.

THREAT HUNTING AND DATA SCIENCE

Unlike the traditional approach of creating and applying a set of detections based on known "bad" behavior, we run actual attacks against our systems and train them on what "good" behavior looks like. We then flag everything else for further analysis and false-positive filtering. This, we believe, is the approach that most other breach detection vendors will also settle on in the future.

Threat-hunting systems need to be able to adapt to changes quickly. Everything in a monitored environment is in flux. People and devices come and go. Operating systems and software get patched. New threats and TTPs emerge. Due to the nature of this flux, traditional IDS solutions tend to be "noisy" and prone to false alarms. These same traditional solutions are also always one step behind the threat landscape.

In order to tackle this problem, our data scientists, working alongside the experts at RDC, have designed and built a series of backend statistical analysis, machine learning, and expert systems to support our threat hunters. You may have noticed others in the industry referring to this approach as "Artificial Intelligence". The core of the RDS backend is very simple, and all of the complexity is embedded in surrounding algorithms. This approach enables very fast deployment times for new detection algorithms (in minutes) and allows us to adapt to changes quickly. With RDS in place, there's never a need to wait for the systems deployed on your own premises to receive updates – all the logic is in our backend systems.

Our analytics systems perform a number of tasks, from analyzing and learning behaviors in monitored environments to reducing false positives. Different analysis techniques are better suited for different tasks. For instance, an expert system is best suited to find the sort of behavior caused by common attack tools and by the TTPs employed by cyber criminals. These include PowerShell commands and malicious URLs and IP addresses. Machine learning systems are designed to spot previously unknown bad behavior, such as DHCP hijacks, spoofing, and other stealthy evasion tactics. We also utilize different multi-level combinations of expert systems, statistical analytics, and machine learning.

We've found that simple statistical analytics are best suited for eliminating false positives, and by applying these methods, we currently eliminate approximately 80% of all irrelevant alerts. The way we've built these systems and the way they interact with each other is quite unique, and something we've not seen elsewhere in the industry.

This combination of artificial intelligence and cyber security specialists is about the most efficient and accurate configuration we could come up with for working with the event data we receive. And it allows us to spot attacks before they have a chance to do damage or access business-critical data.

RED TEAMING

RDS capabilities are primarily developed using an iterative, red teaming approach. In short, we have our guys attack systems, figure out what RDS didn't catch, and make improvements. Some improvements are made by hand. Others are learned by our backend systems during the red teaming exercises. As part of this process, we document and visualize the various attack chains used, which allows the red teamers to come up with new, more devious attack methods.

Our first recommendation to customers who have just purchased RDS is to bring in a third party and run a red team exercise against our service. Not only does it help them verify that everything has been correctly set up, it allows them to see the process in action, which is a nice way to practice for a real incident.

On the subject of red teaming, we've challenged third parties to bypass RDS, but none have managed to do so yet. But there's more. There are at least seventy companies out there that claim they can detect and remediate any targeted attack. In our experience, there are very few that actually can. How do we know? Well, so far, we have a flawless success rate on corporate exposure assignments (where a customer ordered a targeted attack from us). In every single case, we successfully breached organizations running our competitors' products. And none of those products detected our attacks. We're not going to name any names.

SUMMARY

As we see it, today's cyber security situation can be summarized in the following few bullet-points:

- **Most organizations simply don't know if they've been breached or not.**
- **Static defenses aren't even close to being 100% successful against attackers.**
- **Attackers are overly cautious about being caught.**
- **Building good breach detection and response capabilities is difficult.**
- **Red teaming is the only way to properly test your defenses.**

What we've seen happening over the last few years reflects a new reality. And right now, building detection and response capabilities is a complex task involving many separate components and moving parts. And a lot of manual work.

We expect that down the road, the components and technologies designed to detect and stop cyber attacks will be pieced together to create self-adapting automated systems that are able to learn from any new stimuli they encounter. Such systems will automatically run network discovery, vulnerability assessments, patching, and perform post-breach response and remediation activities. When intrusion or post-breach TTPs are discovered, these systems will automatically reconfigure to prevent that mechanism from being used in the future. And in the event of a breach, affected systems, accounts, and access controls will be automatically remediated.

For now, though, if you are concerned about whether you're being hacked (and we think you probably should be), we highly recommend talking to us about RDS. Because we think a managed solution is the way to go. And here's why:

- **You'll have full detection and response capabilities up and running within days of initiating deployment.**
- **You won't need to hire your own cyber security experts, build your own systems, or run your own response operations – we've got that covered.**
- **We promise to contact you within thirty minutes of spotting any real incident on your network.**

If you're interested in reading more, we've got a three-part ebook series that starts with a detailed explanation of a real breach case, explains how companies are going about building their own breach detection and response capabilities, and concludes with our tips on implementing breach detection and response capabilities. We also have numerous blog posts on subjects including cyber crime, detection and response, and detailed explanations of the technologies and processes that F-Secure uses. Finally, we published State of Cyber Security 2017 which includes more stories and case studies from the field. All of these can be found from F-Secure's website.

F-SECURE RAPID DETECTION & RESPONSE SERVICE

CYBER SECURITY EXPERTS	WATCHING OVER YOUR ENVIRONMENT 365/24/7
MAX 30 MINUTES	FROM DETECTION TO RESPONSE *
IMMEDIATE RETURN ON INVESTMENT	AS A TURNKEY MANAGED SERVICE

* Our Service Level Agreement guarantees that no more than 30 minutes will elapse between detecting a real threat and communicating it with the customer.

FROM THREAT TO RESPONSE – HOW RAPID DETECTION & RESPONSE SERVICE WORKS

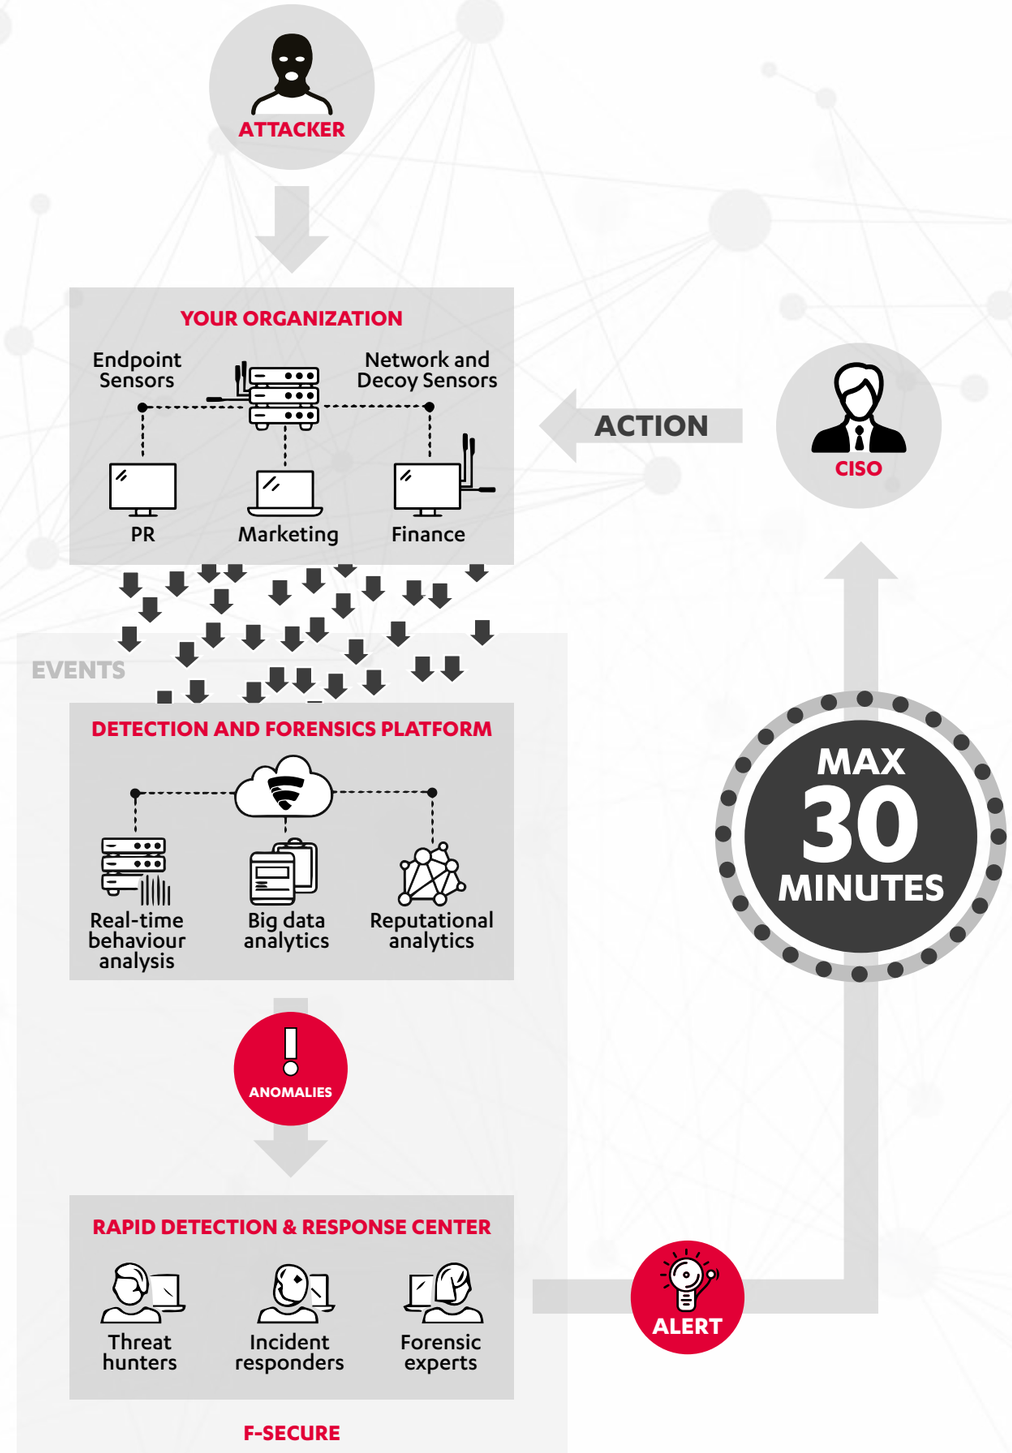

WE SEE THINGS OTHERS DON'T

About F-Secure

Nobody knows cyber security like F-Secure. For three decades, F-Secure has driven innovations in cyber security, defending tens of thousands of companies and millions of people. With unsurpassed experience in endpoint protection as well as detection and response, F-Secure shields enterprises and consumers against everything from advanced cyber attacks and data breaches to widespread ransomware infections. F-Secure's sophisticated technology combines the power of machine learning with the human expertise of its world-renowned security labs for a singular approach called Live Security. F-Secure's security experts have participated in more European cyber crime scene investigations than any other company in the market, and its products are sold all over the world by over 200 broadband and mobile operators and thousands of resellers.

Founded in 1988, F-Secure is listed on the NASDAQ OMX Helsinki Ltd.