



DETECTING TARGETED ATTACKS WITH BROAD CONTEXT DETECTION™

TECHNICAL WHITEPAPER

INTRODUCTION

Cyber security is in the midst of a paradigm shift. Targeted attacks are outmaneuvering the prevention and detection mechanisms companies have in place. Endpoint protection solutions are incapable of detecting fileless attacks that are defined by behavior and the use of legitimate OS tools, rather than by a malicious program being installed on a machine. Detection technologies certainly detect suspicious events, but too often they fail to filter out noise from critical incidents, generating overwhelming numbers of alerts that have no hope of being processed.

According to a 2017 EMA study,¹ 79% of security teams report being overwhelmed by high numbers of threat alerts. And it's no wonder: for example, a study by Ovum found that 37% of banks receive more than 200,000 alerts per day, and 61% receive over 100,000². The Ponemon Institute reports that nearly half of all security alerts are false positives³. Of the rest, a large share are inconsequential and easily remedied.

With the possibility to examine only a tiny fraction of alerts, overstretched security teams are forced to let the majority of alerts triggered on a daily basis go without attention. Teams are left frustrated. EMA found that 52% of operations personnel feel high levels of stress, with 21% of them stating that "not enough manpower" is a stress driver.¹ The cyber security skills shortage itself is well-documented, with a 2017 ESG/ISSA finding the cyber security skills shortage worsening and impacting 70% of organizations.⁴

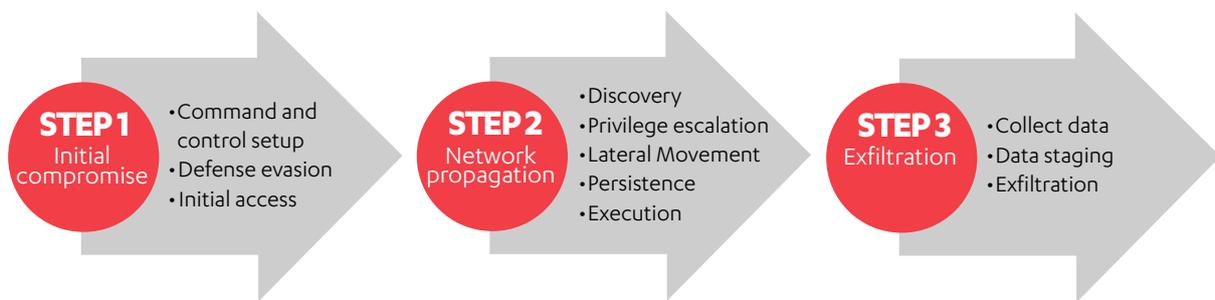
So here we are in 2018, with cyber security high in our collective awareness, and companies are still struggling with breaches. The average breach dwell time is reported to be 100 days, or more depending on the industry and study⁵. Companies are still being caught off-guard with breaches exposing their networks, and their customers.

All the while, the intruders continue, concealed by a sea of alerts.



How a real targeted attack looks: Case Gothic Panda

To demonstrate an advanced and targeted cyber attack, we use the example of the advanced persistent threat (APT) group known as Gothic Panda from MITRE's Adversarial Tactics, Techniques, and Common Knowledge (ATT&CK™) knowledge base and model for cyber adversary behavior.⁶ The adversaries in this example attack are interested in exfiltration of documents and intellectual property, often industrial in nature.



Gothic Panda's process can be broken down into three main phases as shown above. In the initial compromise phase, attackers aim to achieve successful code execution and control of a system within the target environment. The goal of the second phase, network propagation, is to identify and move to desired systems within the target environment with the intention of discovering credentials and documents for exfiltration. In the final exfiltration phase, the goal is to collect the data, compress it into an easy-to-transmit package while in the target environment, and then proceed with exfiltration by hiding in other outbound network traffic. Depending on the defensive setup, exfiltration can be much noisier and more noticeable than attempting to hide in the noise with lay-of-the-land tools.

From a detection point of view, the most important phase is naturally the first one, before the attacker gains persistence and moves to high value systems within the target environment. Well-prepared organizations use a preventive layer like endpoint protection platforms to block commodity malware threats such as and ransomware, which prevents most malicious code execution in the target environment. However, advanced attackers are able to remain undetected by using low and slow attacks, eventually finding a way around the preventive layer. That's when detection and response comes into the picture.

CONTEXT IS EVERYTHING

Everything. In life, and in cyber security. A key turns the lock on your front door. It matters whether the person holding the key is your spouse, or a burglar. A person emerges from a department store carrying a large bag of merchandise. It matters whether or not that person has completed a payment transaction before doing so. A complicated Powershell command is executed from a user's machine. It matters whether it is executed as a part of system maintenance, or by Microsoft Word.

Lack of context, conversely, is lack of well, almost everything one needs to know to make a judgment. Without context, single isolated events are meaningless. It is only when connecting the dots between related events that a full picture can emerge.

Lack of context is a primary contributor to alert fatigue. Many intrusion detection systems today still produce isolated alerts that are in and of themselves anomalous, but when connected with other related events, are found to be innocuous. So false alarms clutter the stack, increasing the workload for security teams who are already strained to the max, and decreasing the likelihood that actual incidents will be uncovered.

MAN & MACHINE

From a mindset perspective, we must let go of the notion that cyber security is about products and services. It's about skills. Experience. Competence. No matter what products and solutions you have in place, it comes down to proactively managing a highly complex and ever-changing environment and threat landscape. That requires skills. Unfortunately, the availability of the skills needed is in shortage, and this shortage will become ever more pressing as digital capabilities become a more important part of value creation for all.

But perhaps even more important is that skills alone are not enough. Just as technology has boosted our capabilities when it comes to everything from productivity work to engineering, we must develop technologies that boost our skills and help us scale our efforts when it comes to cyber security. We must develop technologies that can learn to do what human security analysts do, only at lightning pace – connect the dots, placing events into a proper overall picture in order to make an accurate judgment.



Detective-like piecing together the facts

Detecting broader context is perhaps most easily explained for non-technical people in plain English with a real-world analogy. Imagine a car is found, crashed at the base of a cliff. Was it an accident, or was a crime involved? Was someone in the car? These are important questions that need to be answered in order to understand how to respond to the discovery.

Forensics investigators are called to the scene of the crash. They study the site to piece together the sequence of events that led up to the crash. They study the tire tracks at the top of the cliff to determine whether the car accelerated, or whether brakes were applied. They check the speedometer to see if they can determine how fast the car was traveling. They check the temperature of the engine to try to determine how long the car has been there. They run a license plate check to find out who the car is registered to and scour the site for any sign of a human being. They look into events that happened in preceding weeks, such as whether the vehicle owner received suspicious phone calls or whether his or her browsing history reveals a map service search around the cliff, to root out anomalies in normal behavior.

On their own, each of these factors - speedometer, engine temperature, events of weeks prior the crash, etc. - would appear to be meaningless, but when placed in proper context, a story will emerge that will help investigators determine what has happened and whether a crime has been committed.

GAME CHANGE

From the defender's point of view, this paradigm shift changes the game. It's not an option to give a boost to an existing endpoint protection solution and market it as a new technology. Cyber security vendors need to build new solutions from the ground up that are specifically targeted for the new era and new problems we face. This means a shift from single-shot, point detections and binary on/off responses, to event flow and context-based detections and multi-faceted, risk-based responses.

To give a bit of perspective of the enormity of this change, in the traditional endpoint protection world our backend systems analyze over a million samples every day to decide whether the samples are malicious or not. This is an impressive number, thanks to us already having tens of millions of endpoint clients out there that send these samples.

However, in the new era of trying to detect malicious, hidden attacker activity from the small individual events attackers trigger when executing their tactics, techniques and procedures, the game is totally different. In a single midsized customer environment – with 1,300 endpoints – we have to analyze 70 million behavioral events a day.

Artificial intelligence and machine learning is obviously the only scalable solution that can be applied. But again, AI alone is not the answer – by itself, AI is little more than a glorified false positive generator. Rather, what's needed is the perfect combination of data science and cyber security experts.

TAKING SHAPE

The idea for a context-aware, expert-tuned technology came about after discussions with our customers. We asked them what they were missing in their organization. They told us they have systems in place to stop the commodity malware files that make up the 99.9% of threat volume an organization faces. What they needed was a tool to stop the other .1% of threats that use nontraditional means to infiltrate an organization.

These are the threats that do the most damage, fileless threats that generate events that are almost indistinguishable from events an ordinary user would generate. It is only by connecting the dots between events that a malicious pattern emerges. Connecting the dots is the point where the security analyst is usually needed to step in.

Our Rapid Detection & Response Service, launched two years ago, puts our world-class cyber security experts at the service of organizations. Working from our Rapid Detection & Response Center, they monitor our customers' environments 24/7. When an anomaly is detected, our experts investigate it and once they determine it's a real threat, they alert the customer within 30 minutes of detection.

There's just one problem with this service: these highly skilled experts are in limited supply. We realized that we needed to find a way to bring the knowledge and skills of our Rapid Detection Center experts to any company. So we got to work on technology that comes as close as we possibly can to what our human experts do: investigate the context of an alert to determine if it's a real incident.

The result is something we call Broad Context Detection™.

INNOVATION AT ITS BEST

Spotting misuse from proper use is like looking for a needle in a haystack. It requires collecting vast amounts of behavioral events. Broad Context Detection™ is designed to take this sea of events and filter it down to a trickle of meaningful incidents.

For example, a midsize organization with 650 sensors typically generates roughly one billion events every month, but only about ten detections require containment and remediation actions. The role of Broad Context Detection™ is to enable us to hone in on the few incidents that matter. It does this by analyzing myriad numbers of events, flagging suspicious ones, then relating similar events together and classifying related events into a group that pertains to an incident. Broad Context Detection™ then displays the group events in a chronological timeline to give a full picture of the incident that has taken place.

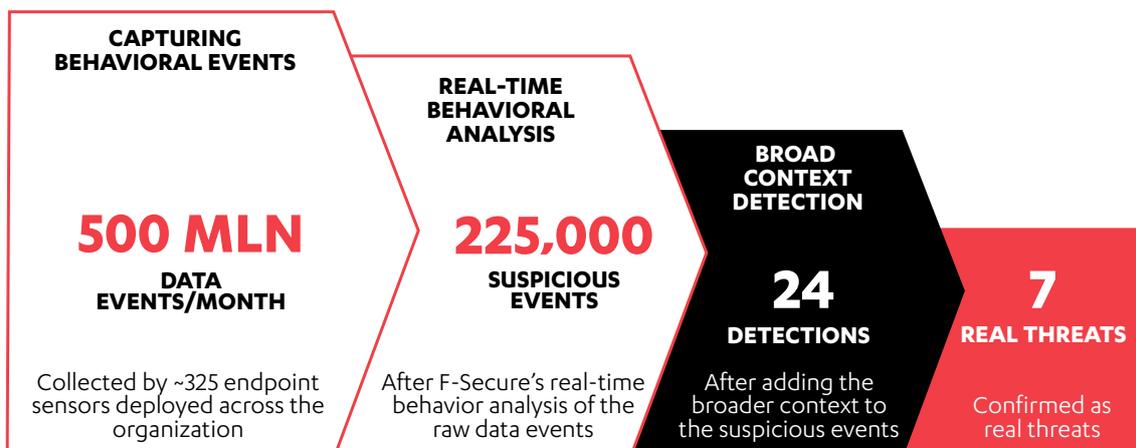
With Broad Context Detection™, as the risk score increases with each detection based on the adversary's actions, the captured behavioral events of the initial compromise will be revealed and the manager of the target environment alerted. The broader context of the attack becomes instantly visible on a timeline displaying all impacted hosts and relevant events, along with recommended response actions.

As result, the attacker can be isolated from the network early, before propagating into the network and exfiltrating data from servers storing personal data, sensitive business documents and intellectual property.

HOW IT WORKS

When incident detection is based on providing alerts, it's no wonder the false positive rate is so high. F-Secure's Broad Context Detection™ takes detection technology further than ever, by filtering alerts down into actual incidents. Using context, we narrow a long list of alerts into a shorter list of detections, then we narrow down the detections to an even shorter list of actual incidents that's clear and actionable for security specialists to respond to.

First, simple events are streamed to our complex event processing engine, which monitors suspicious behavior and moves processes into more detailed monitoring. In this monitoring stage, we take into account the greater context, allowing us to identify events that would be too prone to false alarms if analyzed by themselves. We can now identify suspicious and hostile behavior reliably, and ignore behavior that is acceptable.



Board Context Detection™ in action

In a 325-node customer installation, our sensors collected around 500 million events over a period of one month. Raw data analysis in our back end systems filtered that number down to 225,000 suspicious events.

Suspicious events were further analyzed by our Broad Context Detection™ mechanisms to narrow the number of detections down to just 24. Finally, those 24 detections were reviewed in detail by human experts and only 7 detections were confirmed as real threats.

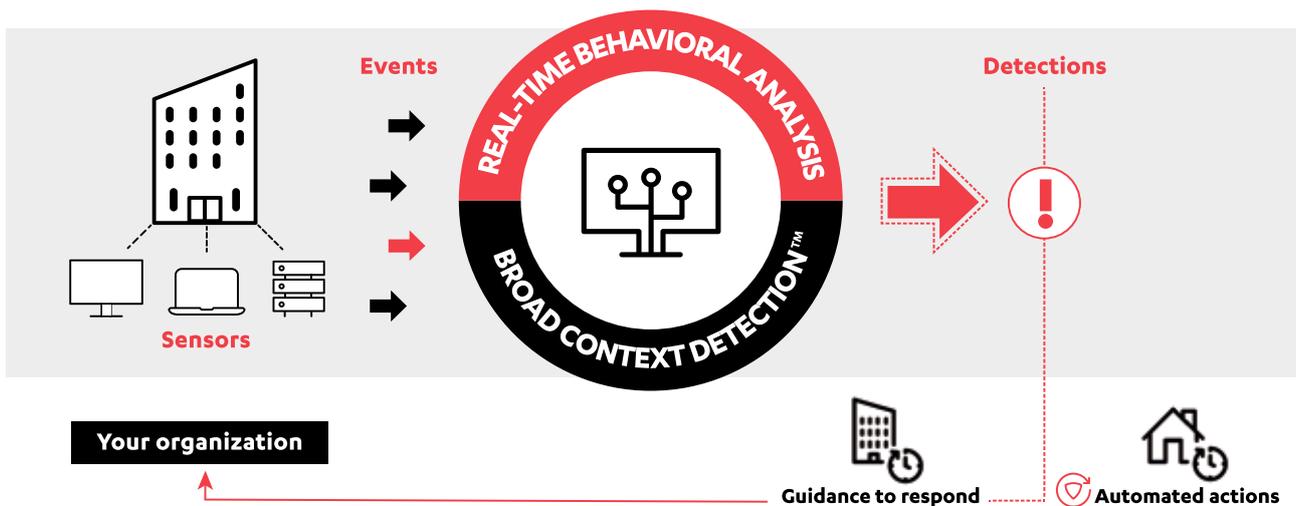
Focusing on fewer and highly accurate detections allows response to be faster and more effective when under an actual attack.

The stream of alerts then enters an alert aggregation engine, which begins to build the contextual picture by grouping together multiple related alerts. Based on the more comprehensive picture that is formed when grouping alerts together, judgments are made for actual detections, with a near zero probability of false positives. Finally, these detections are streamed to an incident detection engine that confirms real incidents. Again, the false positive rate is near zero.

From the now much more manageable, yet comprehensive list of detections, a timeline is built, placing the detections chronologically so analysts can see the full set of circumstances surrounding the incident. Detections are also prioritized by severity according to risk level, criticality of host, and the prevailing threat landscape.

With this approach, IT teams are provided with a relatively short list of confirmed detections, each flagged with distinct priority levels and recommended response actions. So not only do teams know what to focus on first, but they also know how to respond and can do so quickly and decisively.

F-Secure Rapid Detection & Response



1. Lightweight sensors monitor endpoint activities and stream behavioral events to our cloud in real-time.
 2. Real-time behavioral data analytics processes the events, flagging and monitoring both the processes and other user behaviors that have triggered the events.
 3. Broad Context Detection™ mechanisms further narrow down the data, placing related events in context with one another, quickly identifying real attacks and prioritizing them with respect to risk level, host criticality, and prevailing threat landscape.
 4. Following a confirmed detection, the solution guides IT and security teams through the necessary steps to contain and remediate the threat.
-

DETECTIONS AND BEHAVIORS

Broad Context Detection™ flags indications of possible breaches by alerting admins of tactics, techniques and procedures (TTPs) used in targeted attacks. For example, this can include the following possibly suspicious actions:

- Abnormal activity of standard programs
- Calls to running processes from non-standard executables
- Running of unexpected scripts
- Unexpected running of system tools from standard processes

Broad Context Detection™ flags TTPs used to achieve the following objectives:

- Persistence
- Privilege Escalation
- Defense Evasion
- Credential Access
- Discovery
- Lateral Movement
- Execution
- Exfiltration
- Command and Control

Broad Context Detection™ makes understanding the scope of a targeted attack easy by :

1. Combining real-time behavioral, reputational and big data analysis with machine learning
 2. Taking into account risk levels, affected host criticality and the prevailing threat landscape to provide comprehensive insight into an incident and its severity
 3. Presenting only relevant detections with actionable visualization for risk-based and multi-faceted responses
-

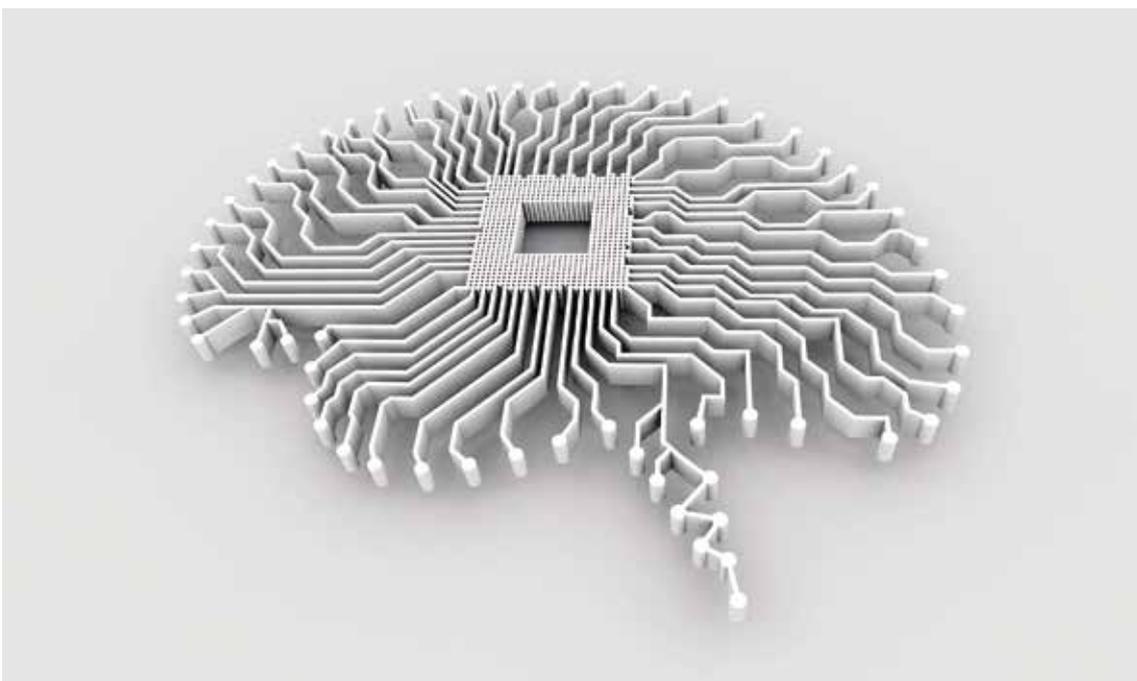
LEVERAGING MACHINE LEARNING TO THE MAX

The use of machine learning and AI that's constantly tuned by human experts means that our systems just keep getting smarter. Unlike traditional approaches of merely training the machine what bad behavior looks like, our basic approach is centered around atypicality modeling. That is, we train our systems what normal, "good" behavior looks like, and then we flag everything that's different from what we expect. This allows us to be open to a much wider array of possible malicious behavior.

This approach means we are not limited to typical attack detection methods, which are based on factors such as unusually high permissions and rapid pattern of operations. Competent attackers have caught on, choosing to employ minimal permissions and so-called "low and slow" attacks, meaning they carry out attacks gradually, in steps paced out over time. This way they can fly under the radar of conventional monitoring tools, because these attacks do not match the patterns that were used to teach the tools.

The advantage of machine learning is that we can train machines to learn from everything, even from their mistakes. (That's something even humans don't always do.) When the machine pinpoints an alert that turns out to be a false positive, the machine learns why it was a false positive, and takes this into account next time, ensuring it won't again flag the same type of alert. This is one of the reasons our rate of false positives is so low.

While real-time detection is the basis of our solution, sometimes it's necessary to detect something after the fact. With machine learning we can easily take new rules for detections our experts have just identified and apply them to old data, picking up on activities that may have been passed over the first time around.



CONCLUSION

Cyber security is about skills and experience. But the defender's dilemma reminds us that we need to have the resources to be vigilant all the time, whereas our adversaries can strike at will. And with the current shortage of trained experts, the adversaries are winning the game.

What's more, skills and experience alone are not enough. It takes technology to leverage that expertise to monitor an entire organizational network and detect the tiny clues that indicate an attack is underway. And it takes superior technology to do so with precision, without activating unnecessary alerts that consume time and resources security teams should be devoting to actual incidents.

Broad Context Detection™ is a central feature of F-Secure's Rapid Detection & Response solution, which provides companies with the advanced capabilities they need to defend themselves from targeted attacks. With the power of machine learning trained and constantly tuned by elite cyber security experts, Broad Context Detection™ ensures F-Secure's solutions pinpoint only the incidents that matter. It's the perfect combination of man and machine, bringing world-class cyber defense within reach of every organization.

With Broad Context Detection™, your company can detect targeted attacks that were previously undetectable. Now go fight to win.

For a video about Broad Context Detection™ go to www.f-secure.com/RDR

Eating our own dog food

Our system is constantly being fine-tuned and analyzed by our experts. The fact that it's the same engine we use to provide F-Secure's premium managed detection and response service means that it's not just blind data analysis based on learning data. We use the system to provide premium services to customers, and then constantly incorporate learnings from those customer environments back into the solution to provide a much higher quality and more finely tuned solution than the typical vendor who only produces the solution. We ourselves are the principal users of several installations, making our motivation customer satisfaction, but more – the system is also part of our daily operations.

More about F-Secure's premium managed detection and response service at www.f-secure.com/RDS

¹ Enterprise Management Associates. A Day in the Life of a Cyber Security Pro (2017).

² Ovum. Closing the Cybersecurity Gaps in Financial Services (2017).

³ Ponemon Institute. The 2017 State of Endpoint Security Risk Report (2017).

⁴ ISSA/ESG. The Life and Times of Cyber Security Professionals (2017).

⁵ Ponemon Institute for HPE. Cybersecurity Trend Report (2016).

⁶ The MITRE Corporation. Adversarial Tactics, Techniques, and Common Knowledge (ATT&CK™) knowledge base (2018).

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.

