WAVESTONE

CERT-WAVESTONE NEWSLETTER

N°7

CONTENTS

FEATURE ARTICLE

THE LONGEST NIGHT: THE STORY OF A CRISIS MANAGEMENT SITUATION ARISING FROM A MAJOR CYBERATTACK

This is a true story, based on real facts, but anonymized and adapted to present an insider's view of the management of a carefully targeted attack.

Nearly 15 team members from CERT-Wavestone spent almost two weeks working on this issue:

as members of a corporate crisis management unit, steering the investigation, preparing a defense plan, and acting as the point of contact for the police... our teams played a key part in a broad range of crisis management activities.

This is their story.

Prolog - Incident definition

[Gérôme] It was Thursday night, at six thirty pm, when a consultant came into my office:

- *"I've got a customer who's just called me, it looks like there's been a major incident that has affected their business systems; they want to know if we can help."*

Straight away, I asked Vincent, the technical manager of our incident response team, to call the client and clarify the situation.

[Vincent] Yes, I remember: it was a Thursday night, and I'd just got off the train and was about to take my bus... when I got a call from Gérôme telling me that one of our clients needed help to respond to a data theft incident. At that point, we had very little information.

I quickly found a spot where I could not be overheard, so that I could give the client a call: the CISO for their French operation. They gave me a brief overview of what had happened: they'd discovered malware on their server that could steal clients' data... I needed to get some more detail on this, so I asked:

- *How did you detect the malware and what makes you think it's capable of stealing client data?*
- *Well, we've installed a hardening tool on our servers and it brought up a series of alerts on a file. First, we just deleted the file... but it appeared again, shortly after. So, I sent it to one of our partners to carry out an analysis ... and I've just received the preliminary report: the program creates files, encrypts them, and then collects them via another program. Having said that, we don't know what sort of data is being put into those files...*
- *What kind of server is it exactly?*
- *It's a local server on one of our sites. We've got over a hundred sites in France. We'd like to check that there are no other sites that have been affected...*
- *What type of data does this server deal with?*
- *Principally, it's our clients' personal data...*

At that point, I decided to create a CERT-W intervention ticket and deploy one of our First Responders. These are CERT-W analysts, trained to carry out initial investigatory steps; they go to the affected site to get the analysis underway, acting as the first line of defense for our clients. I remember, I called Ayoub straight away, to brief him on what he needed to do, the background, the initial information we had, and so on.

- *Ayoub, your mission - and you will choose to accept it (!) - is to get over to VictimCorp to assess the scope and extent of the attack. Ideally, start to flesh out*
- *the most urgent remedial actions and collect the information we need to decide quickly what size of the team to put on this; as well as what skillsets we're going to need.*
- *OK, I'm on the case: I'll put my things together and get over there with our toolkit.*

[Gérôme] As soon as Vincent had put me in the picture about how serious things were, and the fact that we'd dispatched a First Responder, we decided, with the agreement of the client, to contact specialist lawyers and a bailiff. Given the nature and the sensitivity of the data that had potentially been compromised, it was vital to bring in legal advice and follow due process.

The First Responder

[Ayoub] Vincent's tone on the call was a bit tense, I could sense the urgency but he was a long way from being alarmed. There are always some applications that are not quite under full control and that are capable of causing occasional data flows that can give rise to false suspicions... Nevertheless, I began getting ready to go to the site: my investigative equipment (hard disks, cable boxes, a change of clothes so that I could last for a day or two - if I needed to, train tickets, hotel reservation, and so on.).

The next morning, I was greeted by the company's CISO who introduced me to the team that was working on the incident: system administrators, network administrators... as well as another forensic analyst. Suddenly, the issue seemed just a little more serious than I had anticipated.

The team explained what it had observed:

- *A suspect executable had been identified by a hardening tool that had been deployed on Windows systems.*
- *The executable persisted despite numerous attempts to stop and delete it.*
- *Files with an unknown extension were discovered on shareson the system.*
- *These shares were renewed on a daily basis.*
- *They hadn't been able to read the files placed there.*

Fearing that it might be due to old Windows 2003 servers, I asked:

- *Which version of Windows does the malware affect?*
- *We've identified it as affecting Windows 2000 so far... that's the business machines that host highly sensitive applications; it was too expensive to migrate them...*

The tension was palpable, and I understood the issues now: if VictimCorp suspicions were well founded, thousands, or even tens of thousands, of pieces of client data might potentially be affected.

We sat in a small room with remote access to an infected machine. This, of course, goes against normal good investigative practice, but we had enough infected machines to allow us to perform a more normal extraction later, if that proved necessary. We had a meeting with the CIO that afternoon and needed to establish the facts quickly; so we decided to take the risk.

It was a fairly trivial matter to identify the malware in question: all binaries in the directory, C:\\Windows\\System\\, dated from 2009, except for the malware, which showed a date of October 2014. And we were in March 2015...

There was a host of questions that remained unanswered, however:

- How many sites in France were affected?
- Was the scope of the issue limited to France only?
- How did the attacker recover the files created?
- Where was the attacker located?

But more importantly, an unspoken question that I had not anticipated had found its way into managers' minds; an issue that had caused numerous complications and that could cause even more damage...

[Gérôme] Back in Paris, I followed the progress of investigations remotely. At 11 am, the first on-site results were fueling suspicions of a major attack. The binary observed was indeed malware designed to collect data. It wasn't a variant known to conventional protection systems; it had been adapted to suit the business and technical environment. We were facing an attack of some complexity but, above all, targeted: deliberately aimed at our clients' systems. We began to work on the basis that the system had been compromised. The task wasn't easy; there were very few logs, and all systems were flat: it was difficult to find analysis points

At midday, other information from a number of different external sources confirmed our suspicions. We couldn't wait any longer, the company's management had to be warned. The CIO had already been informed and had himself initiated the first actions with respect to the executive team.

But given the criticality of data and the number of affected systems, the situation needed to be treated as a crisis.

[Ayoub] We continued analysis to identify the mechanisms of persistence used by the malware: registry keys impacted, services created, etc. It turned out that in terms of technical complexity, it was probably the least sophisticated malware that we could be facing: execution at userland level, RUN registry keys being used to ensure persistence, a common name for all versions of executables, etc. Inspection of additional machines identified three other versions of the malware. These different versions recovered the encrypted files and reinfected other systems if they were ever cleaned: overall, there had to be some 3,000 infected machines.

[Jean] On Friday, at 5pm, just after I'd finished off my last CERTitude task, I took a call from Vincent; I wasn't expecting that this would be my first step to joining the VictimCorp investigation team.

- *Hi Jean, are you free at the moment?*
- *I haven't got a huge amount of time: my graduation ceremony is in an hour and a half at Bercy...*
- *That works fine! I'm going to send you two executables; find out everything you can for me about them.*

I then received an encrypted email containing these executables and began analyzing suspect-file1.exe and suspect-file2.exe files, without arriving at any conclusive results in the time available.

[Ayoub] At 6pm, Vincent called me to get an update:

- *Do you need any help, Ayoub?*
- *Yes, I do...*
- *Roughly how many people?*
- *There's no limit... The management has authorized an unlimited budget for this...*

At 9pm, a member of the management team came to see me with a fairly serious air, despite the sporty outfit and closed the door on us.

- *Ayoub, there are some real issues linked to this incident. VictimCorp has hundreds of thousands of employees. Frankly, I'm wondering about the possibility that there's been inside help on this one.*

I reminded the client that, given what we had observed so far, all scenarios remained a possibility. In anticipation, though, we continued to follow all the rules to allow us to pursue the investigation through the courts, in particular, by drawing on the support of our partner law firm.

- *Oh oui...*
- *Combien de personnes à peu près ?*
- *Illimité... la Direction a annoncé un budget illimité...*

The Crisis

[Gérôme] That Friday night, the first crisis meeting took place, I was there to represent our investigation teams and assist in its management. With such potentially important consequences, representatives from all business functions had been alerted and were present. The CIO and our team gave an initial assessment of the sequence of events that had been established.

Immediately, the question of who had perpetrated the attack arose. The technical information showed malicious code that had been carefully crafted to reflect the environment that the client operated in, over a period of several months. All scenarios were discussed and debated, could there have been insider involvement? Collusion? Without dismissing this possibility, we explained that what we observed could also be the result of an external attacker who had taken the time to understand the client's system.

Then came the question of keeping the incident confidential. The Communications Department was working, in parallel, on an initial position in case the incident became common knowledge. It was quickly agreed that the operational crisis unit should be based in a particular location. IT staff who knew the SI very well were identified as its potential members. We also decided to increase the size of our investigation team: five people were brought in for the weekend.

By Friday night, we were certain that this was a targeted attack affecting the business systems. Some administrator accounts had also been compromised, and there were no guarantees on the integrity of the Active Directory. We moved crisis management over to a parallel email system, in case the attacker had access to messages. But we were still in the dark about how the attacker entered the information system, gaining access to, and extracting data. These uncertainties also fueled the internal-fraud scenario.

The Investigation

[Ayoub] Florent quickly joined me, after I had completed my initial tasks as First Responder. I briefed him on the situation before continuing the analysis; the objectives were to:

- Consolidate the list of indicators of compromise (IOCs) from the different versions of the malware.
- Reconstruct the infection timeline.
- Identify any other impaired assets (user accounts, Unix machines, etc.).

We worked in collaboration with the systems and networks teams to identify malicious behavior in the few Windows logs that we had on each machine.

In parallel, Vincent asked all the teams to enable logs to be created on all devices: routers, firewalls, Windows, etc., to give us maximum visibility on the IS.

After a day of going through the logs, we finally found the much-sought-after holy grail: the attacker's base camp. A server containing tools for network scanning, an unlisted and unknown local user, but, above all, connections at suspicious hours of the day to different infected machines. We were slowly following the trail of the attack back to the source of the breach.

[Gérôme] Saturday was punctuated by two crisis meetings with the senior management team. The law firm that we had involved was, by now, also present at the crisis unit. The various legal scenarios were assessed, both in terms of the company's responsibilities and also its ability to react. It was decided to prosecute as soon as possible, depending on the outcomes of the investigation.

In parallel, we strengthened the teams from our side and assembled investigatory and defense teams, with a manager heading up each: Vincent for the investigation and Baptistin for defense. We also put in place an overall steering structure for the crisis to synchronize our teams and those of the client, decide on priorities and feedback information: I took that role at first before handing it on to Chadi, when I had to leave for Singapore... Slots with the senior management team were confirmed for 11am and 6pm; they became our focal points over the days that followed and helped ensure regular monitoring of the situation.

[Vincent]] I was at a family meal on Saturday night... but my phone never stopped ringing! We were in the process of trying to organize ourselves to bolster the response team while ensuring that there was a rotation of teams to respect the need for normal working hours and to collect information from the lawyers, etc.

Finally, we decided to send Ayoub for a rest and keep Florent in place to ensure continuity in the investigations. Baptistin and I, who had monitored the analyses during the day, began to get our things together to join Florent on site from Sunday morning and start putting both teams in place.

Taking the train at 6am in the morning was a little bit... early. When we arrived on site, we started by meeting the CISO to explain how we would proceed next:

- *We know the purpose of the attack: to steal customer data. And we also know the area targeted: all sites in France. Now, the goal for the day was to recognize the channel that was allowing the attacks to take place on the servers and sites. Given the volume of assets and information to analyze we will need to rely on your staff to perform numerous investigative actions: the collection of data, analysis of specific cases, and so on.*
- *Perfect; there's also the head of SOC who will be here this afternoon, though he's absent at present.*
- *We will bring him up to date when he gets here. In terms of what's next, starting from tomorrow, we'll form two teams: one to conduct investigations, and the other to prepare to defend the information system and the continuity of business activities. We've identified eight people that can join us and will come in to staff these teams. We will need your own experts to help us progress too. Baptistin will now start working on the defense issues while I head up the investigations.*

So, we started work putting in place regular liaison points with both our clients but also with our crisis managers, Gérôme and Chadi. I applied myself first of all to the - now famous - central server. After asking Florent to create an alerting script for us if the attacker connected on this server, I quickly realized that something was wrong: the account the attacker had created was being used right at that very moment!

- *Oh yes, that's me, said one of the internal team.*

Before investigating further, we needed to hold a quick briefing to instruct the internal team how to react to the investigation: discretion, in order not to alert the attacker to the fact that they were under observation, the traceability of investigation actions, and matters of communication and coordination. Faced with a lack of crisis management tools, I asked our analysts, who were due the next day, to bring our MI6 - «Micro Information System Safe» - with them, a platform on which we have all the tools to conduct effective crisis management: messaging, storage space, wiki, ticketing, shared calendars, etc.

With good practice taken on board by the internal staff, we then resumed investigations on the central server, accessing numerous internal documents and browsing-activity on the attacker's account. This consisted of technical architecture documentation, business procedures, descriptions of particular products, etc. But that's not all; we also found traces of navigation on dropbox-style sites. Without being able to confirm the data theft at this point, our suspicions were growing ever stronger. The discovery of network mapping toolsand, more generally, offensive tools allowed us to categorize this server as the attacker's «base camp»: the asset through which the attacker connects systematically before accessing servers on local sites.

The atmosphere was more studious than stressful really. Everyone was quietly progressing their activities. The first pizzas of the day arrived, and we ate them together, using the time to share information, good practice, and lessons learned.

The head of SOC then arrived. Baptistin and I were able to discuss the situation with him in order to work up the defense plan, but also to give direction to some aspects of the investigation. The exchanges were not very fruitful as there were few effective safety systems in place.

And then suddenly everything took off: the alarm system installed by Florent rang.

The attacker was there, within our client's cyberspace. Right there, and up to no good. Stress levels then went up a few notches, but we had to keep a cool head. We must not make the mistake of letting the attacker know that we were on to them, while, nevertheless, laying sufficient traps to be able to analyze their activities and follow the trail of footprints!

The attacker was only active for the space of a few minutes. But that was more or less enough to give us the vital initial information that they were connected to the base camp from within the workplace... as a domain administrator! The AD was very definitely compromised, and we began to integrate the concept of reconstruction into the defense plan...

Analysis continued on the administrator login, but we quickly decided that it was a straight forward bounce… this time from a web server!

- *Hang on… I know that server, said an analyst from VictimCorp's operational security team. We conducted a security audit on it last year, and there were vulnerabilities everywhere, and the middleware was obsolete…*
- *What did you do about it? I asked next.*
- *We carried out some remedial actions, but it was too expensive to do an in-depth upgrade.*

It's such a shame that we have to wait for an incident or crises to occur before we finally put the right security measures in place! Ah … the pentester's curse! (http://www.securityinsider-solucom.fr/2015/03/le-fardeau-du-pentesteur.html), as Arnaud says…

Our efforts then turned to the web server and especially the Apache logs. We quickly identified the attacker accessing a particular file: a webshell that they had filed previously, including exploiting the vulnerabilities identified during the audit… As a result, we started analyzing the webshell code, but it was fairly well protected. So I sent our analysis team to run a full assessment.

After another pizza and building on the various things we had discovered, we decided to put in place some new monitoring mechanisms. VictimCorp didn't use SIEM or a log manager, and I added WATS (Wavestone Attacker Tracking System) to my list of required tools. WATS is our circumstance-based surveillance system, and it uses the open source technologies, ELK and OSSEC. The goal is not to carry out general surveillance as in traditional SIEM, but rather to have the capacity to track very specific indicators (login, IP, file names, etc.) on a broad set of logs.

[Jean] It was Sunday afternoon, and I heard my phone vibrate. It was Vincent again, but he seemed more stressed than the day before:

- *It's me again; sorry to bother you on a Sunday. Are you able to travel out of Paris tonight? It's to do with the malicious strains that you analyzed on Friday. I'll send an email with all the tools we will need; don't' forget to pack them!*

[Vincent] On Monday, ten Wavestone consultants were deployed to the client's offices. The crisis unit had formed; the roles were clear; and the phase of gearing up for a crisis was over.

It was the most complex bit that was left: understanding the entire attack and managing its consequences.

At 6am, VictimCorp's CIO arrived to take me by car to the first crisis meeting of the day. The decision was quickly taken to locate the crisis unit and operational staff at another site and to continue to conduct investigations and make preparations for defense: after the weekend had passed, we could not keep news of the attack confidential, given that it was right in the middle of VictimCorp's openspace…

By the time we had briefed the complementary investigation team, put together so that we could operate in rotation to ensure the maximum effectiveness of analysts, midday had already arrived… along with more pizza!

The pairs of analysts then resumed their activities. The team responsible for investigation on the compromised web-server system provided materials to the other analysis teams working on assets, accounts, IP, or suspicious files.

Analysis of the webshell had unfortunately not yielded much: it was a common webshell with minor modifications in order to facilitate some job execution tasks and protect its own code.

For my part, I was applying myself to understand the attacker's overall approach, their profile, resources, and motivation …

By late afternoon, we had prepared a large inventory of assets and compromised accounts. It is at that point that the attacker chose to reconnect to his webshell, triggering numerous alerts and capturing information, according to the measures we had put in place… and raising stress levels for those who hadn't had much sleep in a long time… I can still remember people shouting: «The attacker's just connected!» «The attacker's just connected!»

There are times when it's important just to keep calm; to do nothing other than observe.

The storm passed, and we could analyze in detail the attacker's every step. Finally, we had the full picture. Silence and calm replaced excitement: that was it… we knew.

It was now time to formalize our entire understanding of the attacker's methods to be able to marshal the evidence for the prosecution that would be pursued, and provide answers to the questions from the police.

Finally, it was most definitely time to sleep… but only after one more pizza!

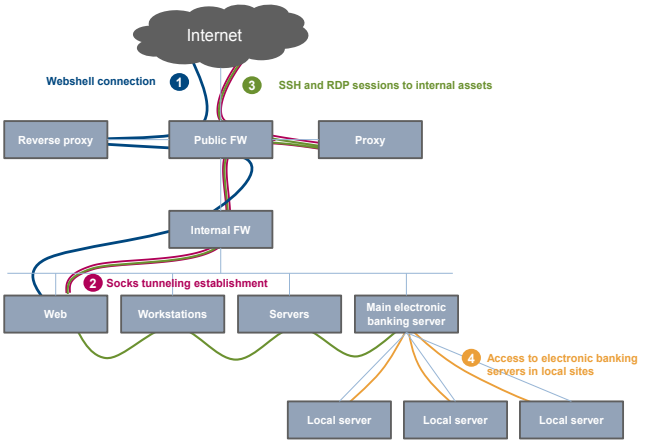

The surveillance activity

[Jean] The attacker logged in every day, at a more or less fixed time, so we thought a more advanced monitoring system would allow us to track their activity, and understand any new areas, accounts, etc. that had been compromised. Our client does not possess a surveillance solution, so we deployed WATS, Wavestone's mini SIEM tool, grouping the ELK and OSSEC stacks, which we largely fed with Domain Controller, firewall, proxy, antivirus, and UNIX logs. A projector was installed in the investigation room, and different dashboards projected on the wall: the ratio of numbers of connections to accounts compromised by the attacker, graphs of their activities, bounce servers, etc. At 6pm, the time when the attacker typically connected, all eyes were on the makeshift screen.

Later in the week, we also had to set up an oversight mechanism for the information stolen by the attacker.

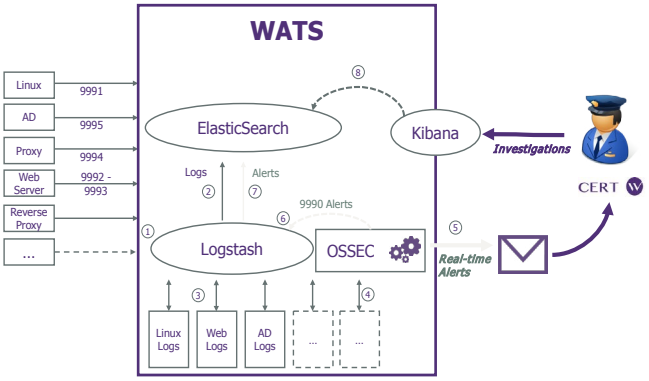

The client did not have a centralized management tool for servers that had been infected, and so we needed to deploy an oversight service for these servers in a rather precarious way. Thus, in the same way that the attacker extracted their information, we recovered our monitoring logs, via a chain of three servers, one of them a DC, so that they could be sent to WATS.

The Defense

[Hélène]] Investigations over the first few days (and nights) had shown that we needed to get a defense plan ready quickly, in order to fend off the attack and be pre-armed against possible recurrences: it was then that Yassir, Fabien, and I, got involved in the defense team, under the guidance of Baptistin.

On Tuesday morning, we joined the Wavestone team already working on the investigations on site. As we arrived on the scene, there was no doubt, that the crisis we were facing was the result of a large-scale attack. Immediately, I met the investigation team on site for a briefing on the initial results of their analyses. At this point, the attack scenario and the attacker's mode of operation were starting to fall into place. We, however, didn't have the complete view of the IS and the possible points of interconnection with providers, partners, international subsidiaries… therefore, Baptistin started the thinking on developing the defense plan. There was a series of immediate questions. Where to start? How the tasks should be divided? We had to consider all scenarios: what if the mode of operation identified so far were not the only one? And what if the attacker had access to other points of entry on the IS? What if recurrences of the attack led to the deterioration or destruction of systems or customer data?

Beyond the measures to be implemented to eliminate the procedure known to the attacker (removing their access to the systems, closing open doors, etc.), we had to get ready for the most critical scenario: that the attack would resume with a vengeance (using a functional Plan B), and other systems were threatened with destruction (with a whiff of Sony in the air…).

To deal with this, we were prepared to isolate part of the IS in order to avoid the attack from spreading.

The idea was that if such a scenario were to occur, we should identify the part of the IS under attack, in order to allow only the most critical flows, to ensure business continuity in the corporate functions.

Based on these different scenarios, we began to list the actions to be taken and the security measures to be implemented; and we further enhanced this plan throughout the period of our involvement.

The defense plan was designed to put a number of threads in place:

- A monitoring device: to act and move forward on the next steps while keeping an eye on the strategic points of the IS.
- A support unit: for dealing with possible side effects of the measures.
- A task force to render the attacker's mode of operation ineffective: ensuring that none of the steps taken by the attacker could be reproduced in the future.
- A controlled isolation process allowing the impacts to be contained if the «new attack» or «degradation» scenarios were to occur.
- Short, medium, and long-term security actions which would allow us to strengthen and sustain SI security levels to reflect the needs of the sector.

A few hours later... the tasks were distributed among the different team members. For my part, I was responsible for the establishment of the support unit and the planning of actions for the potential isolation of the IS... the investigations sprints would now give way to a security marathon!

On Tuesday afternoon, I went off to meet the head of the Support Unit in order to develop a specific escalation process. The goal: to effectively treat the side effects of the actions that would be carried out (incident feedback on potential malfunctions within the business functions and/or IT).

Starting on Tuesday night, a first version of the process was framed, and the main actors at the different levels (L1, L2, and L3) were briefed!

On Wednesday morning, 24 hours after our involvement started, the first actions were taken, and the anticipated scenario unfolded as we had hoped.

The surveillance unit was put in place with Jean's tools; in parallel, I met with people from the business functions to jointly identify with them all the critical flows, and those necessary for the operation of their business units. In short, a piece of work that would take several months, under normal circumstances, to be delivered in under three days! The objective was to identify all the actions needed

in order for a degraded mode operation to be feasible in a case of partial isolation (it should be remembered that, at that point, a worst-case scenario occurring was still a real possibility).

Following the first exchanges, this was far from obvious...

- *Hello Business Functions: in the case where we need to isolate part of your IS network, as a matter of urgency, which flows would you like to see continue as essential?*
- *All of them!*
- *Right... OK... let's try that again...*

The longest night

[Chadi]] Every day since the beginning of the attack, we had been meeting together at 2pm in order to prepare for the senior management's crisis meeting. The objective was clear: to provide a comprehensive and strategic overview to the senior management, in order for them to make decisions. Over the first days, the exchanges very much centered around the investigation, the attacker, and the immediate impacts of the attack. Gradually, the main focus became the defense strategy, both in the short and medium terms. In this regard, the CEO was very clear, «I'll take all the responsibility for past mistakes; but I won't be taking any for future ones.» As part of the prosecution under way, supported by expert lawyers, he asked us to guarantee to him that data was no longer at risk and that he would be able to announce that to the relevant authorities.

The CIO undertook to provide a clear answer on that by the start of business on Thursday at the latest. We still had the afternoon and night to find a solution. The challenge was twofold: first, we had to protect the company's data and that of its customers. On the other hand, we did not want to simply cut off the attacker's access, because it would reveal that they had been discovered, and we did not know at that stage whether there were other points of entry.

Time was short, and each player had a view on what the right strategy was. Some thought it critical to cut off the attacker's access, in order to protect data, others stressed that doing this would not preclude re-access using another door that was not yet identified.

Lack of sleep and fatigue were not helping matters; discussions reached a point of extreme tension, disagreements burst out into the open, and the insults even began to fly... then followed the longest night. Vincent and I then sat down in private to apply some measured thinking to choosing the most rational strategy and then adopting and communicating it to the defense team to be managed and implemented.

At 10:00pm, we had a direct exchange with the CIO. He looked us in the eye, and started counting on his fingers:

- *One, I want a summary from you tonight, clearly describing the situation and the proposed solution. Two, you do not take any action until our CEO has sanctioned it. Three—goodnight—because I'm going home now to tuck my children into bed - having hardly seen them for a week!*

At 1am, the technical teams put an idea that they'd just had to me: a honeypot trap. It was to let the attacker log in, but confine them to an isolated area of the network (a «bubble»), where, in reality, they wouldn't really have access to anything. In this way, we could assure data protection, while giving ourselves the ability to monitor the attacker's actions! Why hadn't we thought of it sooner? Time to validate the gameplan's technical feasibility; I wrote to the CIO to put the solution forward. He approved it immediately.

At around 2am, I asked the most tired team members to take some rest. Unsurprisingly, I was met with some resistance: the excitement of being in the thick of the action meant that there was no way they wanted to stop halfway through.

A relief team worked until dawn to put in place the necessary measures from our plan. So, at 8am, we were able to announce that the measures were in place, and the data protected as far as possible.

Emerging from the crisis

[Chadi]] After this, between Thursday and Monday, we ran two distinct parallel threads.

First, we pursued the legal process. The Police met with us during the crisis, and took a statement, using the evidence put together by our security experts. Their goal was to understand the exact sequence of events, the way the investigation had been carried out, the results obtained, and the way they wanted to pursue the case. So, we provided them with all the material that we possessed.

Second, we put in place very specific monitoring in the attacker's zone of operation. It consisted of a monitor that emitted a visual alert, 24 hours a day, seven days a week, whenever a connection to the "bubble" was detected.

[Jean] The WATS monitoring system in place was by then functional, but wouldn't allow us to trace the attacker using the criteria we already knew about: accounts used and infected servers. We needed to find out if other servers had been compromised by the attacker outside of France, which we already had under surveillance. Vincent then suggested that we used CERTitude, our large-scale research tool for indicators of compromise. Unfortunately, our criteria included the tools having served by the attacker (including PsExec); and we then had a large number of false positives, without discovering any new servers that had been compromised. It was then decided to create a procedure and put a control script in place, then to execute it in each of the relevant countries, under the supervision of our analysts in the investigation unit.

With this step completed and the monitoring and detection tools in place, I was able to withdraw from the investigation team and move on to other work.

[Hélène] Meanwhile, we continued to bring the business functions together in order to map all of their critical flows. The first versions of flow matrices had emerged, while others were still under construction.

We enhanced the defense plan's security measures with the additional feedback gathered by the investigation team, which now had a complete picture of the attacker's mode of operation and the tools they used.

[Chadi] On Friday, as the weekend approached, the attacker had tried to connect twice to carry out their usual activities. Excitement levels reached their peak: the trap had worked! There were several attempts, and the attacker eventually disconnected without taking any additional action.

The weekend was quiet. After being assured that the alert screen would be monitored night and day, the team was able to get some rest for the first time since the beginning of the crisis. Some of VictimCorp's managers, and Vincent, remained on standby during the weekend just in case an alarm came in from the monitor. The beginning of the next week was marked by the absence of any new connections. The attacker had certainly understood things: they had lost their access to the SI.

The state of crisis was then lifted, and a longperiod of overall securing of the IS ensued: this involved all the defense plan measures that were not directly related to the attacker's mode of operation, but would help secure the IS in a comprehensive and sustainable manner. So, of all the scenarios considered by the Defense Team, thankfully, it was the best case that finally unfolded! Phew...

Finally, the actions described here came to underpin the creation of a digital transformation plan across all of VictimCorp's business functions; a project involving an investment of several hundred million Euros...

CAST IN ORDER OF APPEARANCE


Gérôme

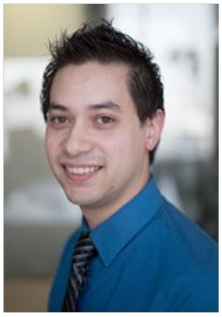
Vincent


Ayoub

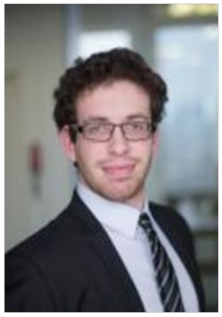
Jean


Florent


Chadi


Baptistin


Arnaud


Hélène

FEATURE

THE NEW TOOLS OF OPERATIONAL SECURITY

In the past few years, there has been a significant increase in the number of cyberattacks, and, in particular, new defense systems have been developed to protect against, detect and react to these attacks.

Let's revisit three types of these new defense systems:

/ «EDR» solutions; next-gen tools for endpoint security.

/ «SIRP» platforms; a panacea for incident response?

/ Machine Learning methods applied to cybersecurity.

EDR: ARE THESE 'NEXT-GEN' TOOLS REALLY EFFECTIVE AGAINST TODAY'S THREATS?

Given the number of cyberattacks affecting businesses every day, and according to the investigators and analysts working on these incidents, the following statement quite accurately reflects the level of protection on information systems:

«The installation of an antivirus software, HIPS, or IDS on any target system presents little or no challenge at all to the attacker.»

While these solutions will force the attacker to take a few standard precautionary measures, they will require common sense, rather than any real expertise:

/ To register new domain names.

/ To reprogram the malware.

/ To avoid system-wide anti-malware scans ...

But current detection and prevention systems in no way prevent the attacker from entering the corporate network, and even less so from broadening their activity to network shares and endpoints of a company's employees.

With this in mind, several market players have brought forward a wave of «innovative» tools, in order to address three major issues identified by companies:

/ To detect an advanced attack.

/ To increase visibility on the endpoints(both workstations and servers)

/ To remotely carry out remedial action on the effects of an attack.

Before looking more closely at new systems such as «Endpoint Detection and Response» (EDR) or «anti-APT endpoint,» let's look at the limitations of more traditional solutions.

LIMITATIONS OF TRADITIONAL TECHNOLOGIE

Antivirus solutions

Antivirus software installed on both terminals and mail servers offers essential security functions and help protect against widespread threats and attacks. However, it is widely recognized that these tools are not equipped to detect, let alone block, any

targeted attack by an unknown malware. In some cases, they would not be able to detect any malware at all (see Hacking Team attack).

Without going into the obscure world of manual obfuscation of code, it is quite easy to bypass the antivirus systems currently available on the market via simple memory-based execution. Indeed, antivirus software automatically scans every file stored on a disk. Therefore, malware that runs exclusively on memory, using one of the many existing techniques (dropper, process injection, DLL injection, etc.) will bypass any protection that has been installed on an endpoint.

Intrusion Detection System

IDS solutions, on the other hand, face an even more challenging problem: signature obsolescence. The level of protection provided by an IDS is only as effective as its signature database. Although it is possible and recommended to implement behavior-correlation rules, via SIEM for instance, few companies take the time to study and implement such rules, either through lack of time, resources, or expertise; or because of the limitations of the tools currently available on the market.

ENDPOINT DETECTION AND RESPONSE: A NEW APPROACH

More and more software developers have been looking at the problems arising at endpoints, in order to offer solutions that can intelligently detect an attack but also, that are able to run investigations remotely and perform remedial actions.

Next, we will discuss in detail each of these three types of functionality in order to shed light on the added value they bring, compared with more traditional antivirus software, but also to highlight their limitations when faced with a real, targeted attack.

Market players

Broadly, there are four major groups of players in the world of EDR:

- Pure players who are exclusively dedicated to the field of EDR and focus their work on its various aspects: SentinelOne, Cylance, Carbon Black, etc.
- Global players who are large software security companies: RSA, Cisco and Palo Alto also offer EDR solutions, which are essentially purchased through pure players, but also have the advantage of being well integrated with their own widely-used systems.
- Players specialized in digital investigation (Digital Forensics - DFIR): some of these players such as Guidance and FireEye have drawn heavily on their experience with actual incidents in order to develop their EDR tools.
- Finally, and surprisingly, developers of antivirus software who complement their own software with a range of functionality taken from the EDR world.

Threat detection

One of the major expectations of EDR is, of course, its ability to detect advanced attacks, for which no IOCs (Indicators of Compromise) are available. We will therefore only report on solutions that offer an intelligent and autonomous detection engine.

Four detection techniques are commonly used by market players.

Detection of the exploitation of a vulnerability

This approach is typically adopted by players such as FireEye, Confer (acquired by Carbon Black) or Palo Alto, and is based on a fairly simple observation.

Whether known or not (zero-day), a vulnerability is exploited by a number of conventional, listable, and, above all, quantifiable techniques: buffer overflow, return to instructions, DLL injection, heap preparation, etc.

These solutions, therefore, target the manifestation of these techniques in the memory of the processes running on the system and raise an alert, if required.

Having said that, in order to have a satisfactory level of visibility on the system, it is necessary to divert a significant portion of the system calls made by the operating system, intervene on the handling of objects by the kernel, etc.; a range of operations that could easily result in denial of service, or even damage to the system.

Consequently, the general approach has been to monitor exclusively the processes that are most often exploited by attackers: zero-day vulnerability on flash would have a high probability of being detected by these tools for example; whereas a zero-day issue on the smb.exe process, a process that manages the sharing of files, would probably go unnoticed.

Detection of malware behavior

Unlike the detection of the exploitation of a vulnerability, which occurs in the early phases of an attack, the detection of malware behavior assumes that the attacker has taken, or is in the process of taking, control over the endpoint.

Thus, some tools such as Carbon Black, RSA or SentinelOne

essentially focus on the chains of actions that are characteristic of an attack. For example, a text editor (notepad) that runs a cmd.exe process and opens a connection on port 443 of a host server located in another continent, might be linked to suspicious activity and will (ideally) be identified as such by these tools.

The major drawback of this method of detection is, of course, the number of false positives that it can generate. A telling example is in Office, where documents containing a slightly more complex macro, are instantly detected as malware by some of these tools. One of the main reasons for this is the huge amount of writing to the temporary folders of the user, %APPDATA% or %TEMP% icons, images, etc.; behavior quite similar to that of malware.

Détection via sandbox

Some developers, mainly antivirus software companies, offer updates of their own software to include «EDR capabilities.» This is certainly attractive from a deployment and operational perspective, but often the update only serves to add an external validation step in the detection process: if a file saved on disk does not match any signature but is still considered as suspicious via a static analysis of this file, it is sent to a sandbox that runs it in a virtual environment to determine its status.

Apart from the detection time inherent to the transmission and the analysis of the file on a remote server, as well as any potential problems with deployment in the cloud of such a sandbox, the detection process is still based on the inspection of files present on the disk only and does not investigate the processes running in the memory of the terminal.

Furthermore, the fact of putting distance between the malware environment (virtual machine - sandbox) and the file environment (endpoint) introduces a bias that can prove fatal. In fact, a number of malwares can take advantage of this separation in order to optimize their malicious potency, thus wrong-footing the sandboxing systems.

Detection using user behavior

Unlike network detection tools, very few EDR players have ventured into the field of threat detection based on learning about user behavior.

The detection techniques previously presented identify clear gaps but these gaps are filed on a «black list» of well-known malicious behaviors. This is a nuance which reflects the principle of signatures, albeit in a more refined form.

It is on this basis that companies like Triumfant have developed quite an interesting approach: they have built a mathematical model that describes the normal state of the activity of a user on a terminal. Any variation from this state: the use of se_debug privilege for the first time, the presence of five active accounts on the same computer instead of one, etc. sends a security alert, which will be validated, or otherwise, by the operator.

Companies such as Guidance also follow this approach but model the behavior of large groups of computers, which effectively detects malware that is spreading on the IS. This approach may fail to detect a targeted attack, affecting a few working endpoints only.

Summary

These detection mechanisms are, of course, complementary to a large degree, and some software developers are happy to draw on several approaches when building their detection engines. However, this does not serve to cover all threats. In fact, almost all solutions incorporate intelligence that has been obtained from software providers' research laboratories, customers, etc.

This intelligence is based on traditional Indicators of Compromise (IOCs), which in the end deliver the benefits but also the limitations of detection systems based on signatures, which have long been part of antivirus software ...

Investigation capacity

One of the major requirements put forward by companies is the possibility of obtaining the status of the terminal and collecting technical information at any given time. This is reflected for example by the recovery, on demand, of the following artifacts:

- The processes and services currently installed on the endpoint.
- A list of files in the %AppData% directory.
- A copy of the RAM.

What once required a manually developed powershell script can now be performed almost instantly from a centralized console, thus offering a comprehensive view of the status of all IT assets.

As well as providing clear efficiency gains, the reliability of results is also improved. In fact, remote search functions natively present on Windows are limited to the user space (ring-4) and cannot detect processes/files/registry keys hidden by advanced malware (ring-0).

This is the major advantage of having software installed at the terminal itself, as long as it is installed in the kernel and carries out low-level functions.

Many software developers go a step further by using APIs

(Application Programming Interfaces), which allows these redundant research actions to be automated. The format of supported IOCs varies from one software developer to the other, according to the strategic choices made: YARA for Guidance which has an investigation background, OpenIOC for CarbonBlack and FireEye, while RSA supports many formats. Nevertheless, almost all software developers are gradually moving toward supporting all formats.

Also, in order to better understand some of the analyses required following an alert, software developers are quite happy to create interfaces with other sandbox tools, either third-party or in-house, in order to carry out more in-depth analyses. This is Cisco and Hexis' approach, for example, but also that of many others.

Remediation options

The third feature of EDR is the capacity to carry out remote remedial actions to endpoint. However, before developing this point in more detail, let's consider the value of such a feature. In fact, deleting a DLL file is a remote process that in no way guarantees the eradication of malware. For example, some malwares, such as Greyfish, will attack the BIOS and can persist even after formatting the disk.

However, daily incidents do not systematically involve malware as complex as Greyfish, and the possibility of removing a registry key for something more typically persistent would make life easier for many analysts.

As a result, some software providers, such as FireEye, propose isolating the infected endpoint only, which offers the guarantee that the threat would be contained to some extent; others, such as Carbon Black or Guidance, propose action on the system itself, performing advanced operations on its kernel in order to try to repair the endpoint.

Finally, other players such as Cisco and RSA propose the remote deletion of some artifacts, but this kind of action deals with the infection at a superficial level only.

Triumfant, a player not very well-known in France, makes it possible to perform surgical-type operations on the memory to correct the malware alterations (IDT tables, SSDT, etc.) and so, is able to process rather complex malware.

A note of caution, though: if the EDR console falls into the wrong hands, it could become a rather powerful weapon of cyberdestruction.

CONCLUSION

Conventional antivirus software such as IDS, etc. clearly offers a limited degree of protection, but it is unrealistic to think that the «new» solutions to be found on the market can provide perfect, infallible options, especially when it comes to detection.

An antivirus software is still required, particularly for protection against basic malware; it cannot be replaced by an EDR.

Finally, the added value of such a tool lies more in the possibilities it offers to collect artifacts, search for IOCs, or even provide the capacity for remote remedial actions, something that can greatly facilitate the work of the response teams in the case of a large-scale attack.

However, such tools need to be particularly well secured in order not to become- themselves-the enabling agents of attacks.

WHAT IS THE ADDED VALUE OF A SIRP?

SIRP, A PLATFORM DEDICATED TO THE MANAGEMENT OF SECURITY INCIDENTS

As suggested by the name (Security Incident Response Platform), a SIRP provides a platform to respond to security incidents. Unlike SIEM (Security Information and Event Management) whose main objective is to correlate logs in order to extract security alerts, a SIRP participates in the management of the alerts raised. As such, a SIRP is closer to a ticket management tool (such as ITSM - Information Technology Service Management tools) and is typically used by IS operations teams on a daily basis.

It also allows the management of alerts and security incidents via some specific features (including workflows for specific processes, IOCs, and threat intelligence, interface with SIEM and other security tools, etc.). The objective of a SIRP is to assist analysts and their managers in their daily tasks; and, more generally, to improve the effectiveness of responses to security incidents.

POSITIONING A SIRP IN THE DETECTION ENVIRONMENT

A SIRP is a platform for responding to security incidents, and as such, it can be considered as an «extension to a SIEM (Security Information and Event Management),» to assist with the analysis, notifications, and organization of the reaction plan.

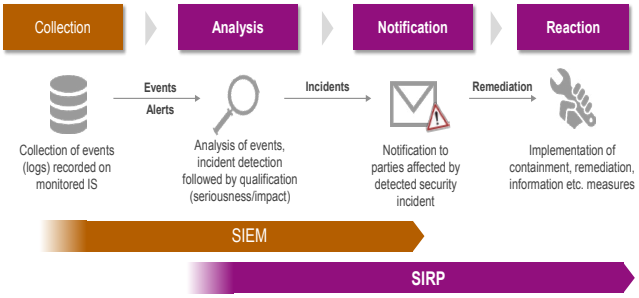

MAIN FUNCTIONALITIES : COLLABORATION, STANDARDISATION AND AUTOMATION

As a central part of the management of security incidents, a SIRP provides functionalities that are specific to its users (analysts, SOC managers, CISOs, etc.). In this regard, in both security and functionality terms, it goes beyond typically used ITSM platforms. Specifically, a comprehensive and mature SIRP platform must be able to provide the following benefits:

/ Increased collaboration through:

- Workflows or playbooks (incident management processes comparable to decision/work trees) as the default option (i.e. supplied with the solution), and that are easily and highly customizable. For example, a process workflow of malware detected on a terminal, with a branch dedicated to managing ransomware.
- An interface with the corporate ITSM (or ITSMs, in the case of highly fragmented IS or organization) in order to communicate transparently (and without changing ways of working – basing things on existing tools) with other IT teams (for example, network teams, endpoints teams, helpdesks, etc.).

/ Improved efficiency in handling incidents through:

- The centralization and accessibility provided by the platform to all analysts with well-defined roles (Level 1, Level 2, Level 3, MSSP manager, SOC, CISO, business unit, etc.), allowing, in particular, the management of priorities (automatically or by assigning incidents) as well as the continuity of service provision by management using rotating teams (for example, a 24 hours-a-day, seven-days-a-week service).
- The automation of analysis, and those actions designed to dispel doubts are usually performed manually by analysts via an interface between the SIRP and other existing solutions: logs managers, IDS, proxies, antiviruses, and threat intelligence solutions. These integrations make it possible for a SIRP to retrieve and enhance any additional information about the incident.

For example, for any alert created in the SIRP via the SIEM and originally generated by several IDS alerts, the SIRP could automatically (or at an analyst's request) consider:

- In terms of traces/markers:
  - » Logs (of the web proxy) if it can find traces of the request and response.
  - » The CMDB or LDAP/AD to retrieve information about the machines affected by the alert (name, type, department, OS, last connected user, etc.).
  - » Antivirus installed on endpoints: alerts and version of the antivirus software.
  - » The email/network sandbox and the last potential threats identified but ignored.
- In terms of contextualization of the threat:
  - » A search for information about the current alert (markers, threats, targets, etc.) within its own database of past incidents (using the concept of capitalization).
  - » An internal threat-intelligence database (for example, MISP instances) or an external one via submissions of markers to external platforms / SaaS (such as Virus Total, IBM Xforce, FireEye, Trend Micro, Palo Alto, etc.)
- For verification purposes:
  - » External applications (for example, sending of a hash or URL toward Virus Total).
  - » In-house solutions, for example, the sending of suspicious files (fed back via a virus or IDS alert) to a sandbox for dynamic analysis (through execution in a controlled environment).
  - » Toward different antivirus solutions currently in place (antivirus on email gateways, endpoint antivirus, etc.) to verify centrally and thoroughly if the threat is known by any of the existing containment solutions.

- The industrialization/automation of the response. Although this is a source of controversy, the automation of the response to an incident (including mitigating actions) is a topic of discussion in many companies. Some SIRP allow, using similar mechanisms (such as interfaces with existing third-party solutions), containment/mitigation

actions to be automated. For example, by providing the analyst with a button to «block the URL and IP on the web proxies» or «generate a signature for the named file on all antivirus software.»

- The large-scale task processing, decisions generated by workflows, and an integrated wiki, are managed by the analysts. For example, an emergency action sheet for the mitigation of a ransomware or a DDOS, a list of contacts for escalation of the issue in any business unit or subsidiary, an example of an incident report, a comprehensive list of lessons learned (action plans for improvement identified after the incident), etc.

Clearly, the interface with third-party solutions is not «natural.» A SIRP generally uses APIs and supports a number of solutions suggested by the software developer. Following this, the capabilities of the third-party solutions, with which interfacing will take place, must be established, either at the application level, using an API , or else at the database level.

CROSS-FUNCTIONAL ADVANTAGES SUCH AS:

/ Advanced reporting capabilities: SIRPs offer metrics that are specific to the management of threats and security incidents (for example, breakdowns by incident type, by impact, SLAs met, etc.), and«reporting role based» metrics: indicators and reports for Level 1 analysts differ from those of Security Operations Center managers and the Head of Information Systems Security for a business unit.

/ Support in meeting regulatory requirements, as a result of increased confidentiality and traceability of information and actions related to the security incidents. Indeed, as a dedicated and securely-designed solution, a SIRP benefits from its own access control and encryption tools, and comprehensively controls all interfaces with third-party solutions (for example, as part of an interface with a corporate ITSM, only the information that is strictly required by the operator is sent to the ITSM). Some solutions can provide traceability and history for all aspects of an incident (for example, the criticality of an incident first created as«high,» later moved down to «medium» on 07/29/2016 by «John Doe Level 2 analysis»).

/ The management (recording, tracking, and capitalization) of «security requests» (for example, the IOC research campaign requested by ANSSI - the French National Agency for IS Security).

/ Support for the organization of war rooms, crisis management, etc.

SIRPS SUPPORT ALL INCIDENT RESPONSE PROCESSES

Finally, as a result of its intrinsic functionalities, and when interfaced with key solutions, a SIRP can enhance the various stages of incident response.

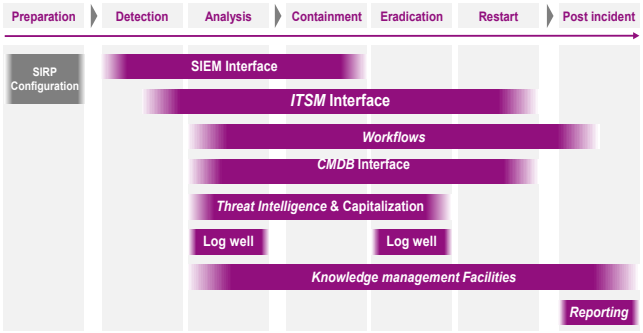

INTEGRATION WITH EXISTING TECHNOLOGY AND ORGANISATIONAL DESIGNS

As a central part of the response to an incident, the SIRP sits at the crossroads between:

/ Technical information: events/logs, alerts, incidents, reports, etc.

/ Management solutions for information systems (CMDB, ITSM, etc.) and security systems (well logs, SIEM, IDS, antivirus, proxies, etc.).

/ The security and IT production teams.

The diagram below provides a view of the functional linkages between the SIRP and other relevant aspects of IS security:

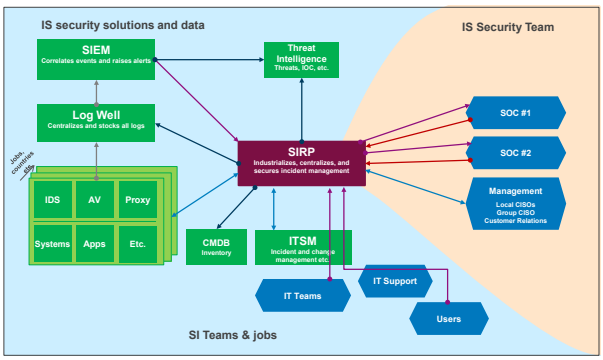

INSTALLING A SIRP

Although this tool may seem very attractive, given its capabilities, there is a range of questions to be considered prior to making a purchase.

PRIORITIZE YOUR NEEDS IN ORDER TO IDENTIFY THE MAIN TARGET

In order to maximize the benefits, it is important to evaluate the organization, in terms of its level of maturity, both in detection

and incident response, and to define the desired target, and then to carry out a gap analysis, particularly with regard to the things that the SIRP can, and must, support.

Here are some questions to ask:

- In terms of current arrangements: what is its maturity of detection and incident response? How are the assets distributed? Which teams currently analyze and manage the incidents? What are the existing or missing processes (handling, escalation, communication, etc.)? What are the main preventative and detection solutions? While tools can help, they cannot replace a dysfunctional or incompetent organization.
- What are the regulatory constraints to be taken into consideration that apply to all, or specific, business functions, and the SI (LPM, PDIS, Basel Accords, etc.)? Can a SIRP play a role in ensuring compliance?
- Finally, what are the expected benefits/gains? For example, is it mainly to strengthen communication between a number of spheres of operation/supervisory teams and incident management? Is it to improve the traceability and capitalization of information? Is it also about standardizing and automating the maximum number of action/response analyses?

Collaborative working is strongly encouraged; input should be gathered from all stakeholders and properly taken into account in order to select a tool that will be best adapted to meeting everyone's needs, optimal implementation, and maximum value.

CHOOSING A SOLUTION ADAPTED TO THE ENVIRONMENT

The needs-identification phase is paramount in a SIRP selection process. There is a great variety of solutions able to meet all or some of users' requirements.

These solutions can be classified according to two main criteria:

- Is the solution open source? That is, is it derived from a project, free, and modifiable, but with, potentially, no technical support; or is it a commercial solution i.e. packaged and ready to use, with technical support, but also paid?
- Was the solution created in order to manage IT security incidents? Or, alternatively, was it developed by the IT production world, and, as such, is it an extension to a ticketing/ITSM tool?

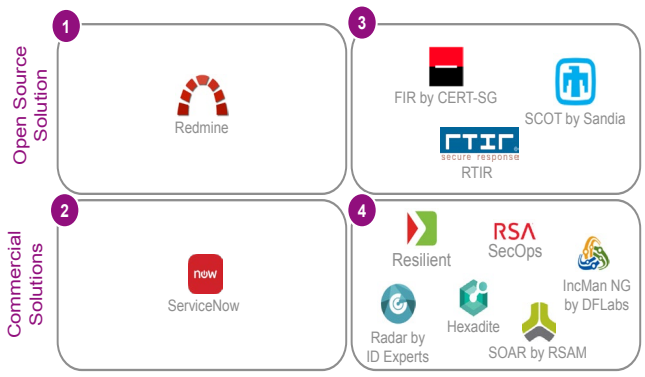

For example:

1. Open source quadrant [3]: the FIR (Fast Incident Response) platform created and maintained by CERT-Société Générale.
2. ITSM quadrant [2]: ServiceNow, the leader in SaaS ITSM which recently added a «Security Operation Management» option supported by three ServiceNow applications: "Security Incident Response,", "Vulnerability Management" and «Threat Intelligence.»
3. Pure-player quadrant [4]: SOAR / PROS:
   - Resilient (acquired by IBM in May 2016) offers a dedicated solution of the same name. Historically, Resilient was undoubtedly the main player in the SIRP market.
   - RSA offers the "SecOps" solution, which is integrated into its «Advanced SOC».
   - The Italian company DFLabs, with their flagship solution «IncMan.»

Typically, if you are a CERT/CSIRT, or a single incident response team with «strong expertise,» a tool developed by another CERT/CSIRT (for example, CERT-Société Générale's FIR or CERT-BDF's TheHive), which you operate and which can be tailored to your needs, is almost certainly the best choice.

On the contrary, if your company widely uses an ITSM solution with a security incident management module, it might be wise to consider carefully the possibility of capitalizing on the ITSM by thoroughly assessing the functionalities which are specific to the security incidents.

If your company has a high level of maturity, and you are definitely looking for advanced features to respond to incidents which would be supported by default (for example, a multi-SIEM plugin, multi-threat intelligence, etc.), it might be wise to consider some of the market's pure players; offerings here are largely driven by the US market and its range of MSSPs (Managed Security Service Providers).

CONCLUSION

Clearly, using a SIRP is not essential, and security incidents can be dealt with without one. However, for some organizations, especially large companies operating across many sites, or with an international reach, a SIRP may be an attractive solution. The use of a SIRP leads to efficiency gains in the handling of security incidents (shorter response times, better response quality, more value from actions, etc.), offers a reporting tool for incident response, and can ensure coordination between different teams.

But, like any new solution installed on an existing information system, its full and successful integration can only be achieved through the involvement and support of all the relevant stakeholders.

MACHINE LEARNING AND CYBERSECURITY

Everyone has heard of «machine learning," «artificial intelligence," «big data» and «analytics." What is there about these concepts, especially about machine learning, that can be applied to cybersecurity?

WHAT IS MACHINE LEARNING?

«Machine learning» can be defined as «the concept of using data and algorithms to allow a machine to learn for itself.»

THE THEORY, THE DISCIPLINES AND THE MODEL ...

Unlike statistical modeling which consists of the formalization of rules between variables in the form of mathematical equations, machine learning is the name given to the concept of an algorithm that has the ability to learn from data without relying on preprogramed rules. In this way, machine learning belongs to the realm of information and artificial intelligence.

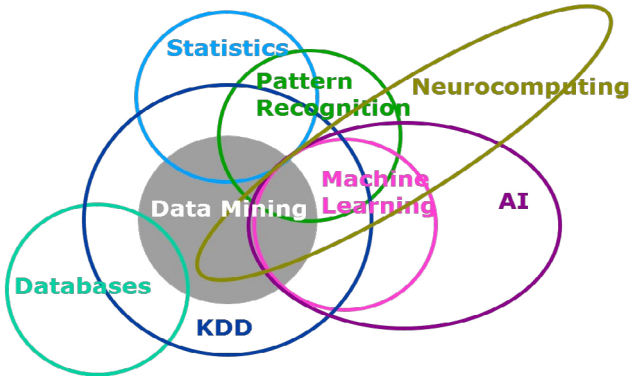

MACHINE LEARNING DRAWS ON RECENT TECHNOLOGICAL ADVANCES

The three main pillars of machine learning are:

/ Data mining: made possible, and justified, by the amount and variety of data produced today, all of which is potentially collectible and available.

/ Pattern recognition - in particular, enabling the creation of links between the data collected in order to highlight patterns.

/ Neurocomputing - as an additional analytical tool, inspired by biological neural networks such as the brain.

These capabilities are made possible by recent developments in technology, as illustrated in the diagram below:

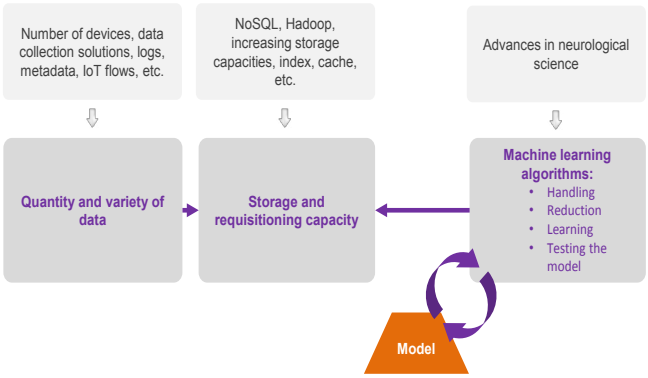

Last, machine learning is the «brain» which allows meaning to be extracted from a data warehouse.

HOW DOES IT WORK?

Giving machines the ability to learn is not self-evident! Here are the key concepts:

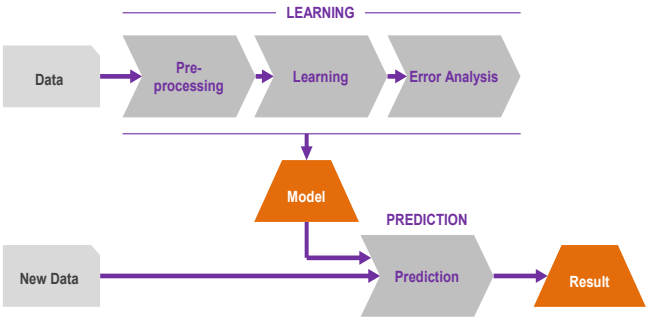

The process can be broken down into several phases:

/ Preprocessing of the data through standardization and data cleansing makes them accessible for processing.

/ Learning can be based on several types of algorithms, including:

- Supervised learning creates a model by building on categorized examples gained from past experience. It requires training data made up of two groups: firstly input data or values (also called «features») and secondly, a category («labels»). For example, phishing websites vs. clean ones. The purpose of creating this model is to predict the classification of data where only the features are available. Specifically, in order for a program to recognize a car, for example, it is «fed» by tens of thousands of images of cars, labeled as such. Once trained, it can recognize cars in new pictures.

- In-depth or unsupervised learning is aimed at leaving the worry (and the work) to the machine (understanding the algorithm) and determining categories by highlighting common patterns and differences. This method differs from supervised learning by the fact that there is no

defined output. This technique is based on a «neural network» in which the results from the layer of neurons are used as the input to the calculation of other layers In March 2016, the alphaGo program, having learned to play the game Go by this method beat the world champion Lee Sedol by 4 games to 1.

/ Error analysis makes up a test phase of the model.

Then, in order to evaluate the classified data, the two groups must be combined.

This phase is undoubtedly more sensitive. Indeed, correlating them involves merging rules related to its functions (with the intended use of machine learning), mixing rules, in particular, those gained through testing and exceptions.

MANY POSSIBLE APPLICATIONS IN A GROWING MARKET

Artificial intelligence and more specifically machine learning is definitely a growth area! With, for example, historical applications such as fraud prevention in the banking sector or the prediction of illnesses and helping with the decisions for the associated care (see Google DeepMind and London Eye Hospital), there are many potential sources (which are not the object of this article).

Attracted by large companies, and especially Google, several announcements about machine learning have been published recently, for example:

/ In May, Google made its «TPU» (Tensor Processing Unit) chips public, specifically designed to be optimized for reduction operations (the basic operation of machine learning), usable with its open source library called «TensorFlow».

/ In June, BrainChip , a producer of microchips dedicated to machine learning, bought the French company «Spikenet Technology,» a provider of multi-application technologies for real-time shape recognition to be used with applications designed for, for example, airport security.

MACHINE LEARNING APPLIED TO CYBERSECURITY

Let's look now at the applications of machine learning in cybersecurity, particularly in the detection of security incidents.

DETECTION ENGINES IN A RAPIDLY CHANGING ENVIRONMENT

Before talking about applications proper, let's set out a summary of current solutions and detection engines:

/ Historical methods (such as antivirus software) are based on signatures: the software looks, particularly in the files, for traces of signatures known to form parts of malwares.

/ The most advanced solutions are based on simulation (i.e. execution of suspected malware in a sandbox environment) to determine whether the malicious behavior results, or not. This involves a combination of behavioral analysis (for example, it is not «normal» for a binary to read and re-write a large number of files placed on storage devices, which could be a sign of ransomware) and of researching Indicators of Compromise – IOCs - (for example, a file sending an HTTP GET request toward a URL/ IP known to be listed as being part of C2)

/ New solutions for large-scale behavioral analyses incorporate large volumes (many Gbps) and varieties of data (flows, metadata, logs, etc.) that look for any discrepancies that could indicate malicious activity.

It is most notably on this latter type of machine learning solution that true added value can be realized.

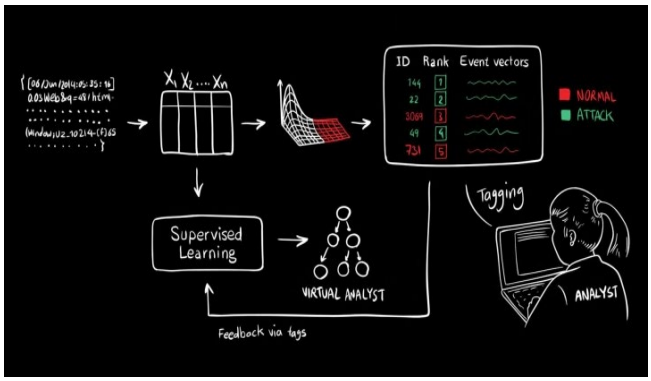

Among the major developers of security solutions, RSA following Security Analytics (the RSA SIEM - Security Information and Event Management system - based on logs and network packets) announced, at its conference in March 2016, the integration of a real-time behavioral analysis engine based on machine learning.

On the "Big Blue" side, in May, the developer and provider of the well-known SIEM platform, «Qradar," announced the release of a cybersecurity application offshoot to its machine learning platform «Watson.»

IBM is planning to feed «Watson for cybersecurity» from its «X-force» Threat Intelligence flux, but also less structured data such as SPAM messages, malware, and research reports helped by the partnership established on this project involving eight North American universities. The publisher expects to process 15,000 documents per month.

Of a more modest size, and closer to us, is the Darktrace company, established in 2013 and based in England. They offer a detection solution called «Enterprise Immune System» that doesn't use IOCs but relies solely on machine learning algorithms. A variation in the field of industrial IT is also available. The developer claims an average learning time (in terms of data capture, manual contextualization, warning system, iterative tuning, etc.) of about two weeks.

The solution, in addition to having a detection engine that differentiates it, offers a particularly visual interface.

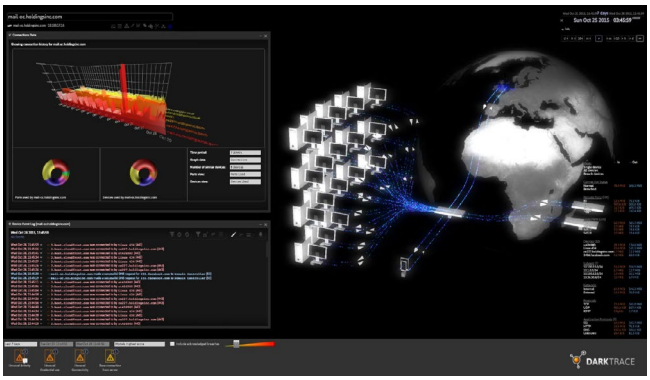

We have recently had the opportunity of seeing it work (in a demo environment) and of testing it as part of an RFP for one of the company's clients. The solution seems particularly promising, and Darktrace raised $65 million in early July 2016 in order to continue its development.

What is particularly interesting about these solutions is that they do not only focus on the detection of «initial infection,» but also, and most importantly, on the detection of «symptoms» (meaning post-infection ones) which offers increased visibility with respect to malware detection.

A SUPPLEMENTARY DETECTION SYSTEM RATHER THAN AN INDEPENDENT ONE

Although machine learning is not new, it has recently proven itself, and there are many possible applications (fraud, health, insurance, etc.) This combination makes machine learning currently very popular and will keep it in the spotlight in coming years.

In the field of cybersecurity and the detection of security incidents, the major players (for example, RSA and IBM) expand their platforms through detection engines based on machine learning and «pure player» companies (for example, Darktrace) are emerging and marketing dedicated solutions. That said, the market still seems quite immature.

Also, there are very few large companies today that are using a security-incident-detection solution that is primarily based on machine learning. Indeed, software security teams are often overwhelmed by the current methods of detection without needing to add another one. A real level of IS security, therefore, seems to be a prerequisite to implementing such a solution.

In all cases, machine learning is a complementary approach as it is not intended to detect the same events: its contribution seems much more significant to malicious behavior detection following a compromise than to detecting the compromise itself.

Finally, machine learning applied to cybersecurity would seem to be a supplementary method and not an end in itself. As far as solutions are concerned, the display capabilities, automation, and reporting seem crucial. On the human side, availability and quality of expertise are needed now more than ever.

In short, a subject area to keep a close eye on!

INTERVIEW

THE THREAT INTELLIGENCE STRATEGY OF AN INTERNATIONAL BANK

Christian Karam is the Cyber Threat Intelligence Director of a large international banking group. Based in Singapore, he oversees the division responsible for providing the business with a constantly updated view of all the current threats, and the operational teams with valuable information to help them identify threat indicators, the techniques, tactics, and processes that allow them to respond to and mitigate threats more efficiently. Christian Karam also conducts recognized research, security and cybercriminality work, in his capacity as an expert who previously worked for Interpol.

Threat Intelligence isn't a very developed field within businesses; why not?

In the majority of companies, it still isn't a very widespread activity. It involves significant investment in terms of human resources, and that's why few companies have started working in this area, to date; especially since the skillset is very rare...it either comes from former police officers who worked on cyber issues or cybersecurity researchers ...

Basically, it means picking up the American Total Information Awareness (TIA) model again. A Threat Intelligence program should serve to revolutionize the capabilities of an organization to detect, classify and identify attackers, and, of course, allow the company to take quick action to protect themselves and thwart cyberattacks.

Nevertheless, banks have realized the importance of developing Threat Intelligence: it is a question of scaring criminals by sending a strong message that «we know who you are, where you come from, and we can arrest you.» The main objective is assigning the attacks: who is attacking, what are their motives, what are they hoping to achieve, what are their capabilities? We work together with the police to arrest the attackers. We have to make it so that cybercriminals consider it «not worth the hassle» to attack the bank: because technically it's too complicated and too risky.

Threat Intelligence also aims to understand how cybercriminals can be better than us in defensive terms, and by doing so, we aim to recast the balance of power.

La Threat Intelligence a également pour objectif d'identifier en quoi les cybercriminels peuvent être meilleurs que nous dans la défense et ainsi de rééquilibrer le rapport de force.

How should Threat Intelligence activities be structured?

First, a process must be implemented: who is going to do what, how, with whom... the roles and responsibilities have to be defined.

Next, there is the whole design phase: what products to use, how to integrate them, which tools will be more effective, and which sources of Threat Intelligence are needed...

In concrete terms, we have established an intelligence center that works in partnership with the SOC, the anti-fraud team, and the cybercrime investigation team.

It is this center that receives all the different threat intelligence flows, handling them on behalf of the other teams, and defining what the threat landscape for the bank looks like, in order to identify the key areas to focus on, depending on the geographical location and the latest cybersecurity events.

The intelligence center was established to address a number of operational and strategic problems:

- On the operational side, the goal is to provide useful information to other teams in order to help them and

guide them in their investigations and in their IT defense activities. Generally speaking, we try to apply the OODA model (Observe, Orient, Decide and Act) for our security operations. The intelligence center analyzes and documents attacks, in order to provide useful information to guide other teams, in particular, the SOC, to help them make good decisions before taking action in the protection arena or defending IT systems.

- From a strategic perspective, it consists of taking a step back from the threat scenarios that threaten the company, and the attacks it has suffered, in order to provide a more comprehensive perspective than merely a view on an individual attack. Each individual case is reviewed to identify metrics and assess the attacker's capabilities. Next, we identify if our defense is at the same level or better than the attacker's by considering the nature of the technology being used to block the attacker.

This evaluation work is done twice a month and allows us to identify areas where we need to make investments.

The center also conducts external research: we are trying to flush out the attackers by identifying their infrastructure, tools, methods, etc.

Currently, we have a dozen analysts, and we have a target of 35 analysts in 3 years. By way of comparison, the SOC today employs around 40 people with a target of 50 within 3 years.

What challenges do you face as you put this team in place?

The main difficulty is the skill requirements. Few people are able to do this type of analysis: you need skills in systems, networks, reverse engineering, cybercriminality, and so on.

Furthermore, we have issues with the Threat Intelligence sources: finding a quality source is very complicated. The Threat Intelligence market is essentially American. As a result, 90% of the indicators give useful information for North America, compared with about 10% for Europe and Asia.

We also face cultural issues:

- The concept of «collecting everything, knowing everything, recording every signal and sound» is not really something that exists in Europe.
- We are confronted by the language barrier in many cases.
- European countries have a tendency to be much more conservative about data protection than others.

Finally, the legislation is not the same everywhere, and we can find ourselves facing legal constraints in the collecting or sharing of information within the same organization, where it is spread over several different geographical regions.

What investments are there to consider?

We have already discussed this a bit, but generally:

- A skilled workforce
- A flow of useful and relevant threat intelligence.
- Basing your strategy only on data flows is bad practice because, over time, there will be a huge number of false positives. You need to start with internal resources in order to prioritize the different areas and enrich the whole through external sources. A mapping exercise on internal intelligence capabilities is required before focusing on the external.
- Be careful, because the more sources there are, the more work there is for the analysts...threat intelligence has to be seen as «searching for a particular needle in a stack of needles—all attacks look similar!" This is all the truer because groups of hackers often seek to imitate each other for camouflage purposes: states want to impersonate groups of cybercriminals in order to remain anonymous, and, conversely, cybercriminals adopt the methods of states in order to appear legitimate.
- A comprehensive toolkit. There are two main types of tools:
  - Threat Intelligence Platforms (TIP): for correlating internal and external information (sources, bulletins, vulnerabilities, etc.).
  - Threat Intelligence Lists: for consolidating indicators and other information collected.

In summary, building a Threat Intelligence strategy consists of:

- Narrowing the scope of what you are working on: it is impossible to work on the entire sphere of a large international group, so it is essential to prioritize.
- Define the department's structure and processes: pay particular attention to analyst skillsets!
- Aligning Threat Intelligence activities (strategic, operational, tactical and/ or technical) don't over-invest in pointless activity if the organization can't track it!
- Build the appropriate tools: TIP, TIR, surveillance tools focusing internally, then externally....neither too much, nor too little!

Our thanks to Christian Karam for agreeing to this interview!

CERT-Wavestone combines a range of technical and business expertise to provide an initial comprehensive response to security incidents. More than 45 experienced experts operate within the CERT-Wavestone structure.

NEWS

The summer has not been "one long holiday" for our analysts: they have dealt with data theft, blackmail, IS destruction, ransomware, an act of revenge by a former employee… our clients have been attacked from all sides! We should remember, however, that the majority of attacks, among the diverse range that clients experience, are perpetrated by amateur cybercriminals, not organized bodies.

You can find our experts' publications at:
www.securityinsider-wavestone.com

SUBSCRIBE : CERT@WAVESTONE.COM

Twitter @secuInsider

To subscribe to the CERT Newsletter:
cert@wavestone.com

REACTION TO ATTACKS OR SUSPICIONS

- Digital Forensic and Incident Response
- • Business function and SI crisis management / Threat Hunting
- Construction of remediation plans

THREAT INTELLIGENCE

- Evaluation of company attractiveness
- Analysis and attack decryption
- Watch & Learn: Cybercriminality surveillance

DEFENSE PREPARATION AND CRISIS MANAGEMENT

- Definition and leading of CERT and SOC processes
- Red team and Purple team
- Crisis scenarios

Publication Director: Pascale Imbert
Editor: Frédéric Goux
Contributors: Gérôme Billois, Baptistin Buchet, Hélène Dutilleul Ayoub Elaassal, Chadi Hantouche, Mathieu Hartheiser, Jean Marsault and Vincent Nguyen.
Photographs: Getty Images - Fotolia
Diagrams: Wavestone
Printer: Axiom Graphics

Wavestone-CERT
CERT manager : Matthieu Garin
cert@wavestone.com

WAVESTONE
www.wavestone.com

Wavestone is a consulting firm, created from the merger of Solucom and Kurt Salmon's European Business (excluding retails and consumer goods outside of France). The firm is counted amongst the lead players in European independent consulting.

Wavestone's mission is to enlighten and guide their clients in their most critical decisions, drawing on functional, sectoral and technological expertise.