WAVESTONE

FUSION CENTERS: THE FUTURE OF SOCS

Our clients manage operational security through two main organizational structures: the SOC, responsible for the detection, qualification and management of incidents; and the CSIRT , who oversees crisis management, digital forensics, technology watch and threat intelligence.

But what are the key levers to optimizing this dual structure—and what are the most promising avenues to shape it for the future? The concept of the Fusion Center offers the elements of an answer.

The respective roles of the SOC and CSIRT haven't been precisely defined, and some tasks fall where their responsibilities overlap: for example, what is the SOC's role in crisis management? Or, how can the CSIRT help in detection?

In addition, the creation of teams dedicated to new and specific needs (for example, anti-fraud units and business-function SOCs), further complicates the way operational security is organized.

Putting this complexity aside, the record of operational security to date is mixed and recurrent problems remain: a shortage of cybersecurity resources; high rates of false positives generating excessive workloads for analysts; poor quality, or non-existent, data repository; and difficulties in meeting business-function needs within the timescales desired.

AUTHORS

BENOIT MARION
benoit.marion@wavestone.com

JÉRÉMY PAGEAUX
jérémy.pageaux@wavestone.com

This publication has been produced with contributions from Amaury COULOMBAN, Cybersecurity and digital trust consultant

BOOSTING EFFICIENCY THROUGH AUTOMATION

It's clear that, today, analysts are losing valuable time by having to perform manual interventions across a large number of tools: copying information (that has to be copied and pasted between the SIEM, a multitude of repositories, and different ticketing tools) and connecting to resources and security tools, in particular those for remediation. Aside from the lost time, these repetitive, and often non-value-added, tasks create frustration and fatigue for SOC teams.

Automating the incident management process

SOARs (Security Orchestration, Automation & Response)—tools for assisting and automating responses to security incidents—aim to overcome these irritants.

These platforms are designed to act as the single tool used by all those involved in incident management. They allow the analysis and response process to be clearly defined and tailored to each security event. Once a process has been defined, some tasks can then be automated through API-based interactions with IT and security solutions.

During the analysis phase, the tool can automatically enrich the security events by retrieving context information from the IS or Threat Intelligence services. Some solutions even approach complete N1 automation by offering the option of chatbots, for example to check with an administrator (using multi-factor authentication) that he or she has carried out a "legitimate" sensitive action.

This information can be used to automatically close some alerts, as well as prioritize those that analysts need to process, and facilitate their qualification work.

But automation doesn't stop there! Although limited to a few well-established interventions (the simplest being the blocking of URLs found to be malicious on the proxy and the removal of phishing emails), the automation of responses has significant benefits in terms of workload for security teams.

Automating the creation of rules based on known threats

Even though analysis and response processes can be automated, the creation of detection rules currently requires a large amount of "manual" effort, which is often unique to the context: their tailoring to the environment's technologies, their specificities (thresholds), etc. In addition, such rules are often based on standard attack signatures, which are sometimes known—and avoided—by more advanced attackers.

If used well, Threat Intelligence can help surveillance teams address these issues. Threat Intelligence platforms can provide context-specific information about current threats to a SOC in an exploitable format (IOC s that can be recovered using APIs). By interfacing these with the SIEM, these platforms can feed it, both automatically and in real time, with the latest attack signatures to be detected. And this mode of operation is proving its worth: MSSPs already use this approach and report that the majority of their legitimate alerts come from scenarios based on Threat Intelligence feeds.

Threat Intelligence platforms therefore make it possible to automatically create detection rules. In addition, they also play a part in the automation of event analysis, providing (to the SOAR) the information needed to judge the legitimacy of an alert—in real time.

Putting in place advanced detection with Machine Learning

In addition to the complexity involved in deploying them, traditional rules have another flaw: they're based on static analyses and, unless a substantial effort is made to maintain them, they generate a large number of false positives. In addition, a signature-based approach detects only known threats—it cannot detect sophisticated attacks.

At the other end of the scale from a known threat/signature approach is Machine Learning—which makes both behavioral analysis and anomaly detection possible. Tools based on this can detect attacks for which it's impossible to define a signature: where too many cases are involved, or the required degree of correlation is too high. They are therefore better suited to the needs of business functions which SIEM teams find difficult to address, given their technical standpoint: application monitoring (using product-specific log formats) and protection against fraud (based on a very varied range of sources).

What's more, Machine Learning can complement a signature approach: helping to reduce numbers of false positives by automatically adapting thresholds to the context (network volumes, number of users, etc.).

Adapt the organization to this automation

Many of these developments require the roles of SOC team members to be adapted. We see detection and response teams operating on the fringes of the SOAR—and all of them, after all, are working toward a common security goal. Fusion Centers, thus, offer a means of unifying the SOC and the CSIRT. To manage the Machine Learning tools described above, there's also a need for data scientists within the Fusion Center.

These developments are beneficial both for security monitoring and the teams who work on it. Reducing the burden of recurrent tasks and increasing the variety of high-added-value jobs will allow Fusion Center staff to take on new roles—something that will greatly help to limit staff turnover.

ADDING A TOUCH OF AGILITY TO THE SOC AND ITS ACTIVITIES

The transformation to a Fusion Center isn't just about tools. More agile modes of operation must also be considered if the effectiveness of operational security teams is to be improved.

In current SOCs, detection rules are essentially created top down—by working downward from an overall risk analysis to specific supervision scenarios.

Systematic best practices for creating detection rules

While relevant, this methodology doesn't allow 100% of the risks initially identified to be covered, and, very often, the process to be followed when an alert occurs isn't defined. The scenarios deployed remain frozen in their initial state, frequently generating false positives or becoming ineffective over time. However, a number of good practices can be applied to limit these problems.

Too many scenarios, which have been tested in a technical sense, prove not to be functional in real situations—as a result of a problem on the log chain, a different format from the one expected, or poorly defined thresholds. Each scenario therefore needs to be subjected to an end-to-end test—to ensure that the associated alert is raised at the right time.

Such testing ensures the effectiveness of supervision scenarios, but the response procedure must also be well defined, working in collaboration with the business-function project team (client, requester and rule designer).

Once these detection and response rules are in place, they must be adapted and maintained if they are to remain optimal. The most effective way to continuously improve the quality of the rules is to set up a systematic feedback loop. For each alert raised, the players (both analysis and response) debrief the Fusion Center, having implemented the rule. If the alert proves erroneous, the rule (or the automatic SOAR analysis) must be modified to remove these false positives. In all cases, feedback can be used to automate the actions carried out: adding a source of information, automating a remediation, etc.

In addition, a peer review by several analysts during the creation of each scenario will improve the quality of the rules, as well as facilitating the sharing of knowledge and maintenance over time.

Threat Hunting: an effective complement to supervision scenarios

Another way to improve the quality of detection is to supplement supervision scenarios (which are mostly top-down) with a bottom-up approach based on Threat Hunting.

Analysts are assigned a regular slot (for example, a particular day of the week) to go "hunting" for suspicious events. Their findings may be gathered from external sources (for example, confidential information found on the Dark web), or internal ones—such as production resources or security tools (including SIEM logs). The objective, then, is to identify suspicious events or behaviors to be investigated and possibly subject to an alert. By doing this, Threat Hunters will improve existing detection rules, and create new ones. What's more, the method allows analysts to improve their knowledge of their environment, encourages a proactive search for threats, and generates proposals for new, value-added alerts—as opposed to simply responding to SIEM alerts according to defined SLAs.

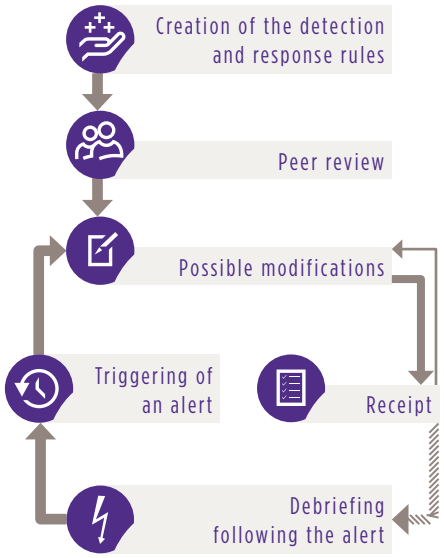

More agility in relations with business-function project teams

All these tasks are complicated and require the involvement of the SOC, which, as a result, finds it difficult to address new needs. Business function and anti-fraud teams therefore tend to create their own dedicated teams in response. But, creating separate entities limits the scope for global correlation and generates redundant effort, something all the more unfortunate given that their information sources are often very similar (infrastructure logs, application logs, user access logs, etc.).

If Fusion Centers are to be placed at the center of supervision activities, including those for the business functions, there's a need to facilitate interaction with business-function project teams. By borrowing some Agile concepts, Fusion Centers aim to involve these business teams more deeply in supervision—both in the design and alert processing phases. To achieve this, new resources ("supervision champions") need to be included in the agile business-function project teams. These teams are then able to set out supervision needs that can be directly implemented by the Fusion Center, and to update them as the business-function project evolves.

In return, when an alert is raised, the Fusion Center analyzes and enhances information on it (automatically if possible!). It then communicates the alert to the business-function project team, which is best placed to qualify—and potentially remedy—an alert that falls within its scope. Finally, feedback is used to continuously improve the process and quality of alerts.

THE FUSION CENTER: AN ENHANCED SOC DEMONSTRATING ITS EFFECTIVENESS!

As a result of automation, Machine Learning, and a more agile structure and mode of operation, Fusion Centers will improve the efficiency of today's SOCs and CSIRTs.

To ensure this improvement is tangible—intangibility being a failing sometimes leveled at SOCs—there needs to be a means of measuring its maturity. There are numerous benefits to doing this: being able to demonstrate results, justify (often considerable) investments in a new structure, or ensure compliance with the regulatory framework.

Standards, limited until this point, have increased in number in recent years. We'll mention two here: the french ANSSI PDIS[1] standard, which is very rigorous and comprehensive, offers a model of the state of the art; as well as the open access SOC-CMM standard[2], which covers all relevant topics and enables SOC self-evaluation using a specific set of questions.

Of all the good practices discussed in relation to the Fusion Center of the future, the one to adopt today is most definitely the regular measurement of maturity which these standards enable.

1 - https://www.ssi.gouv.fr/uploads/2014/12/pdis_referentiel_v2.0.pdf
https://www.ssi.gouv.fr/uploads/2014/12/pdis_referentiel_v1.0_en.pdf
2 - https://soc-cmm.com/

Summary of the new staff, tools and operational methods of the Fusion Center