Improving Effectiveness of Incident Response

In the information security profession, already considered to be very challenging, incident response (IR) is a complicated process to do well. There are a number of factors that make this so, many of which are familiar to the security industry and are exacerbated in IR.

This document offers a look into the operational aspects of IR, analyzes the strengths and weaknesses of current IR practices, and highlights solutions that can both increase security analyst productivity and significantly strengthen an organization's overall security preparedness.

Why Incident Response is So Difficult

Skills Shortage

The security industry is constantly in a state of shortage for skilled professionals. Good security people typically have a mix of skills and experiences, often a combination of general technical skills, system administration, networking, development, and a smattering of others. Developing a person that is this well-rounded takes time and resources, thus the general shortage. Often IR requires even more specialized skills like packet analysis or malware reverse-engineering.

Detection and Sources of Actionable Data

Two conjoined issues are the detection of unwanted or unauthorized activity and acquiring sources of actionable data.

Detecting activity that the IR team might be interested in is often a difficult proposition. A solid threat detection mechanism is often built on a foundation of complex and widespread security instrumentation to watch systems and networks at a high level of detail. Such instrumentation, in and of itself, can be difficult to architect and expensive to implement. Putting it in place does not immediately gain us anything, and the information output is not immediately actionable.

The second of these issues is arriving at actionable data. Actionable data is found at the intersection of good data and good intelligence. We may be instrumented to the furthest extent possible, but if we have no context for the data, we can't do much with it. A list of IP addresses communicating does nothing for us if we don't know the source or identities of the senders or receivers. It is only at the point where we can say one of them is a known malware IP or a Tor exit node that the information becomes actionable.

False Negatives and Positives

Security tools, in general, often have problems in correctly detecting and classifying attacks. We can commonly see this in the firewall and Intrusion Detection Systems (IDSs) that comprise the core of data fed to SIEMs. Such devices primarily tend to fail in one of two ways: either they don't notice an attack (false negative) or incorrectly classify benign traffic as an attack (false positive). It's worth noting here that Cyphort's Anti-SIEM has been certified by ICSA Labs to accurately identify 100% of advanced threats during their extensive testing process.

While we might assume that false negatives would be the worse of the two situations, false positives can be just as bad or worse. False positives produce alerts which we must spend resources to investigate. If our systems produce bad alerts often enough, our IR team will begin to ignore alerts from the system in general—good or bad.

Manual Data Analysis

Much of the work done by incident responders is a manual process. During an incident investigation, IR teams may need to analyze network traffic, sift through a wide variety of forensic data collected from multiple systems, review email messages or attachments, and perform many similar tasks.

While we might like to think that this type of analysis happens in some automated fashion, it does not. This kind of analysis happens on spreadsheets and notes typed into incident reporting documents. This is, of course, slow and inefficient, and can lead to mistakes and inaccuracies—all of this is highly undesirable in the area of security and most likely to venture into legal proceedings.

Mitigation

Mitigation is, perhaps, one of the most difficult problems that incident response teams face. Mitigation in real-time is not often possible with the standard set of security technologies at our disposal, with the possible exception of Intrusion Prevention Systems (IPSs). Whenever we have a security technology that requires a human decision maker in the loop, our reaction speed is far too slow. By the time a person sees an alert, checks to see what the problem is exactly, makes a decision, and takes action, the attacker may already have what they need. For mitigation to be successful, it needs to happen at compute speed, not human speed.

Help With Improving IR Effectiveness

There are several technologies that we can use in conjunction with our existing security tools to make IR more effective and minimize some of the issues that we just covered.

Instrumentation

One of the primary foundations of good incident response is good instrumentation. If we aren't monitoring our environment adequately, we won't find the bad guys. This type of monitoring often involves placing sensors at the choke points between differing zones of security on our networks, as well as the boundaries around our most sensitive systems. As part of this strategy, it's obviously important to have continuous visibility into web and email traffic, the two top attack vectors used by cyber criminals.

Presuming a solid security architecture is in place, an attacker will need to move into and out of an area covered by our instrumentation, and we will, hopefully, be able to classify their traffic as being unusual in sufficient time to act on it. This is, however, one of the tricky parts.

UEBA

User and Entity Behavior Analytics (UEBA) solutions use complex algorithms and statistical analyses to detect anomalous patterns of user behavior in our environments, across data, applications, systems, and networks. UEBA is singly focused on user behaviors (which can render it somewhat myopic).

While UEBA is somewhat useful, it provides only circumstantial evidence of a threat and cannot on its own deliver a conclusive threat verdict. This is because anomalous behavior does not always equate to malicious or unauthorized activity. When we are collecting information from log data only, further investigation and manual intervention are required to differentiate an anomalous threat signal from the noise.

The five virtues of the Cyphort Anti-SIEM

- **Direct detection of threats.** Directly identifies evidence of advanced targeted attacks and malicious exploits.
- **Exploit to exfil.** Visibility into threat progression across the complete Cyber Kill Chain. Does not rely on circumstantial anomaly evidence.
- **Scope of attack.** Visibility into the scope of attack accelerates understanding of the devices affected, how active the threat is, and where it's going.
- **Risk-based prioritization.** Contextual understanding into which critically important assets (host and user) IR teams need to focus on first
- **One-touch mitigation.** Isolates affected assets and blocks threat movements. Reduces the time to mitigate threats through efficient integration with other tools throughout the security infrastructure.

Improperly tuned or tempered UEBA analyses can result in very high volumes of false positives and subsequent bad alerts—which can mean lots of wasted time. This, as we have discussed, may result in productivity impacts for our IR team and increase the potential for alert fatigue.

In contrast, when we can collect raw data directly from endpoints or network traffic and thoroughly analyze it, we can produce actual direct evidence of threats which provides us clear and actionable results for mitigation.

Machine Learning

Machine learning is a branch of artificial intelligence that enables computer systems to learn patterns or carry out tasks without explicitly being programmed to do a particular task. Machine learning systems learn and change as they are exposed to new inputs. Machine learning can be used in the detection process to discover new attacks that we have not previously seen based on uncovering covert deception behaviors and patterns. Machine learning can also automate some of the tasks that our Incident Responders normally must handle manually, such as poring through logs or analyzing data to find actionable information.

Mitigation

As we covered earlier, one of the large problems for us, as defenders, is responding to attacks in time to do something. Very often IR investigations are based on what has already happened. This is partially due to our inability to react swiftly enough to stop the actual attack as it was happening.

One solution to this is to automate some of the steps associated with mitigation. While we do see some small measure of this, on a limited basis, it typically only occurs in firewall and IPS systems. Automated threat mitigation needs to be implemented across our entire security infrastructure to be truly effective at stopping attacks in real-time.

While these technical components are good and useful by themselves, they do not present a complete solution when acting in isolation. To arrive at a workable solution, we need something to tie them together and make them work smoothly in concert.

Is a SIEM the Right Security Tool?

Given our need to effectively tie all of these technologies together, the obvious solution would seem to require the deployment of a SIEM. There are, however, a few issues with traditional SIEM technologies.

Deployment

SIEMs are notoriously difficult to deploy. Yes, we can roll the equipment out and ship logs to it, but getting anything meaningful back out of it can be a challenge without time-consuming customization and configuration, often involving costly, third-party professional services.

Once data is flowing to the SIEM, which can be challenging in and of itself, parsing the log format can often be problematic. While most SIEMs can cope with standard log formats such as Syslog, or output from Microsoft servers, others can be an exercise in deciphering the format and teaching the SIEM how to interpret it. More esoteric formats, such as those from custom applications can present even more of a problem.

Once we have the data in a usable format, writing useful rules for alerting can be even more difficult. While we can easily do simple alerting, based on statistics such as a count of failed logins, correlating log data across multiple incompatible systems requires a great deal of familiarity with both the data and the intricacies of the SIEM itself. This often requires either hiring dedicated resources to develop rules or the purchase of professional services from the vendor.

Maintenance

Once we have our SIEM installed, keeping it up to date and functioning properly typically requires dedicated resources.

The foundation of the data and alerting produced by the SIEM is the set of inputs from the systems feeding logs to it. Unfortunately, the task of keeping these systems up is a cumbersome one, particularly when dealing with large numbers of sources. The breadth of systems that we have forwarding logs will often require several methods of picking up logs—from passively receiving data to actively authenticating to the systems on the other end to actively retrieve the desired data. This, of course, makes for a very change-sensitive set of tasks requiring updating whenever the systems on the other end of the exchange are modified.

As with maintaining the connections that provide us with data, we also need to maintain the set of rules that correlate the data when it arrives. Here we need to both manage the addition or loss of systems providing data gracefully. Also, we need to have a method of coping with changes to the format of the data, or the secondary sets of data that we must maintain for correlation purposes, such as account names or IP addresses.

Perhaps the worse of the maintenance issues are coping with the body of data that we maintain in the SIEM. This poses several challenges, including the retention of the data over an extended period, the infrastructure that backs the massive volume of data that we are storing, and the licensing, often sensitive to the quantity of data being processed or events per second arriving at the SIEM.

SIEM Calibration

In addition to the issues mentioned above, coping with the output of the SIEM is a matter of its own. As we discussed above, SIEMs require regular tuning to produce meaningful data. The need for such tuning is often discovered through the arrival of a vast number of false positives. Unfortunately, each of these false positives requires manual review by an analyst to ensure that it is not an actual event. Also, tuning a SIEM is often a trial-and-error event, sometimes taking several attempts to get a particular alert firing in a way that is useful to us.

In short, building and maintaining a SIEM is not for the faint of heart and not a trivial endeavor. SIEMs can benefit from third-party functionalities that augment their effectiveness, and these include advanced threat detection, advanced threat analytics, and one-touch threat mitigation.

How the Anti-SIEM Architecture Augments IR Effectiveness

The Cyphort Anti-SIEM solution is a distributed software platform that combines advanced threat detection, advanced threat analytics, and one-touch threat mitigation to ingest, store, analyze, and visualize threat data to enhance mitigation. Cyphort's architecture relies exclusively on direct threat evidence collected from across the network, continuously ingesting web, email, and lateral spread traffic from a range of devices (including in-line devices and endpoints) to unmistakably determine malware exploits and command and control threats.

Ingest

Cyphort's detection technology, based on machine learning and behavioral analysis technologies, ingests raw data and log data from Cyphort collectors as well as other information sources such as firewalls, intrusion detection systems, and endpoint detection and response tools. This data is stored for analysis and future visualization in a timeline view.

Store

Cyphort's integrated storage architecture is easily scalable based on the requirements of the individual customer, scaling from a few months of historical data, up to several years and beyond—all on commodity hardware. This storage flexibility enables the timeline view of security incidents to be extended to weeks, months, or more based on the historical data stored. Licenses are based on the number of users, not data storage volume or events, enabling customers to store as much data as is needed.

Analyze

The SmartCore multi-stage data correlation and threat analysis engine leverages machine learning in combination with advanced behavioral analysis. The SmartCore analytics engine and native detection engines are focused on making interactive investigations productive and efficient for analysts and incident responders. In addition to making use of raw data from its collectors, SmartCore ingests events from other detection and identity sources. It then employs correlation algorithms that enable data reduction, risk-based prioritization, and automation of tedious manual processes.

Visualize

After the analytics engine correlates threat detection data with all related security events generated from other devices in the network, relevant data is consolidated with host/user identity information and presented to security analysts as a timeline view of a complete security incident relating to a named user. These visualization capabilities reveal threat progression through the Cyber Kill Chain, as well as the scope and impact of a threat against specific hosts and users.

Based on this information, analysts can immediately determine the nature of the threat and take steps to mitigate it directly or escalate it, in the case of advanced threats.

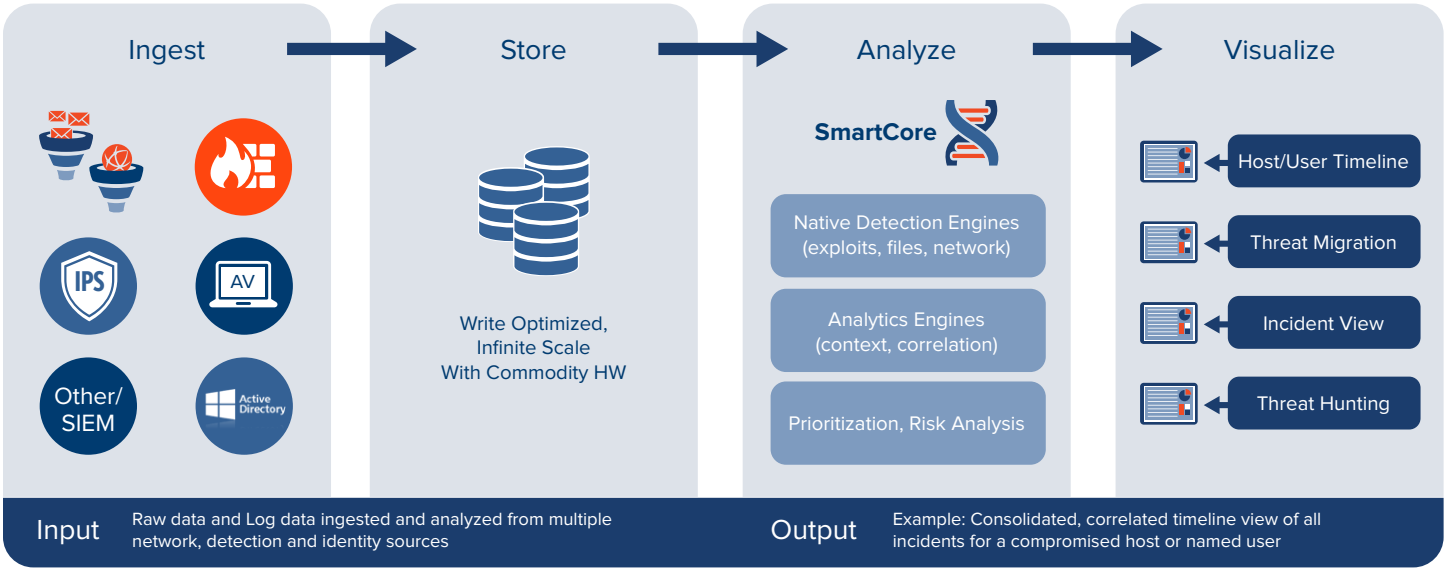

The Anti-SIEM Reference Architecture

Conclusion

Operating an effective and efficient Incident Response team is a difficult proposition. We are often told that we need to reduce costs, do more things with fewer resources, catch the bad guys more often, and respond more quickly. To do all of these things, we need the support of automation backed by smarter, better, and faster technologies that start with accurate proactive detection of advanced threats.

Using security technologies augmented by solid instrumentation, machine learning, and behavioral analysis, we can indeed begin to perform more efficiently and at a smaller cost. We can free our talented analysts and responders from the repetitive tasks that waste their time and talents and can apply them where they can be of the most use.

The future of the security industry is in enabling our foundational security tools by combining them with advanced threat detection, intelligence, analytics, and automated response.

About Cyphort

Cyphort, Inc. is a security software company providing mid- and large-size enterprise customers with innovative security analytics for advanced threat defense. The solution is built with an open architecture that integrates with existing security tools to discover and contain the advanced threats that bypass the first line of security defense in an organization. Based in Santa Clara, California, the company was founded in 2011, is privately-held, and distributes its software through direct sales and channel partners across North America and international markets. Learn more at www.cyphort.com.

CYPHORT

5451 Great America Pkwy, Suite 225, Santa Clara, CA 95054 | P 1.408.841.4665 | F 1.408.540.1299 | info@cyphort.com

Sales/Customer Support P 1.855.862.5927 | F 1.408.540.1299 | support@cyphort.com