A Proactive Approach to Incident Response

A SANS Whitepaper

Written by Jake Williams

August 2015

Sponsored by
Blue Coat Systems

Introduction

Incident response (IR) costs skyrocket every year as the number of breaches increase. Worse, most industry professionals agree that it's not a matter of *if* but *when* a compromise will occur for any organization of suitable size or value, which means IR costs will only continue to rise. Whether it's paying for credit monitoring or credit card reissue, these activities are expensive. They are also brand damaging. Nobody wants to be remembered as the next "largest breach to date."

As became clear in the early days of the U.S. Office of Personnel Management (OPM) hack, not being able to articulate what is taken during an incident can (rightfully) garner distrust among affected parties. Further, investigations are increasingly difficult because attackers, as seen most recently in the Sony[1] and OPM[2] compromises, use anti-forensic and evidence-destroying techniques.

Any incident response has two components that drive overall cost:

1. How long does it take to detect the intrusion after the attackers first gain access?
2. Once detected, how quickly can the incident be remediated?

Finding a solution that addresses both questions with satisfactory answers is the job of any organization that cares about saving costs and protecting data. In today's security environment, though, separating the important signal from the noise is one of the bigger challenges incident responders face.

In testimony about the 2015 OPM breach, the organization claimed to be blocking 10 million "confirmed intrusion attempts" per month.[3] Regardless of whether that number is inflated, it speaks to the multitude of alerts that analysts must investigate. No amount of manpower could possibly investigate that number of alerts, and simply ignoring alerts is negligent. Organizations, therefore, need a way to quickly identify alerts that are worthy of deeper investigation. After those events are identified, they need to be investigated (and quickly). Speed matters but so do accuracy and thoroughness.

Shortening the time to detection and remediation likely is an evolutionary process in most organizations. Maturing the IR process not only saves money, but it can also help reduce the number of records lost and the amount of financial damage perpetrated by intruders. By comparing three typical levels of incident response sophistication, this paper explores why significantly improving IR capabilities is a crucial undertaking for organizations that have sensitive data and explains how to implement it.

[1] "Sony Hack Fits Pattern of Recent Destructive Attacks," www.csmonitor.com/World/Passcode/2014/1204/Sony-hack-fits-pattern-of-recent-destructive-attacks-video

[2] "OPM Hack Probe Hindered Because Digital Trail Has Been Erased, US Official Says," www.abcnews.go.com/US/opm-hack-probe-hindered-digital-trail-erased-us/story?id=31784335

[3] "OPM Chief Says Government Agency Thwarts 10 million Hack Attempts a Month," www.nbcnews.com/tech/security/opm-chief-says-government-agency-thwarts-10-million-hack-attempts-n376476

Moving to Proactive Incident Response

Before discussing the concept of proactive IR, let's consider the alternatives. How do most organizations handle IR today? Three IR maturity models are commonly found in organizations, and their results are definitely not comparable. Where does your organization fall and what might you do to get to the next level? Table 1 provides a quick overview of the three basic IR maturity models. The red text in each IR maturity model indicates less than ideal conditions associated with that model.

Table 1. Three IR Maturity Models

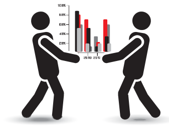 **Manual Forensics**	 **Basic Forensics**	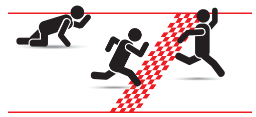 **Proactive Incident Response**
Limited log data	Security information and event management (SIEM) system usually in place to standardize log retention and correlation	Network data available and complete
Logs must be gathered from multiple departments	Some packet data available but not always readily searchable	Networks sensors in appropriate locations
Single-purpose open-source tools dominate	Incorrect sensor placement common	Network traffic decrypted as necessary
Network traffic capture is reactive	Initial vector of compromise can usually be determined	Network data available and automatically enriched with reliable threat data
Results difficult to validate and not easily searchable	Amount and type of exfiltrated data not easily determined	Unknown data automatically investigated
No standardization of retention requirements		Analysts receive only relevant alerts

Manual Forensics

The baseline IR state is manual forensics. In manual forensics, some level of log data is available. These logs must be acquired from multiple people across multiple departments (data center administrators, network teams, and so forth). Forensics professionals know all too well that this is a cumbersome process. Evidence needed by one team may be trampled on by another, and chain of custody is a nightmare with so many departments collecting data. Network traffic is captured only after the incident is discovered. Because network traffic capture is reactive rather than proactive, it is difficult or impossible to discover the original point of compromise. Single-purpose open source tools with little integration capabilities dominate the manual forensics process.

Results from manual forensics investigations are difficult or impossible to validate. Rather than being held in standard repositories, the data used for the investigation is stored in many flat files, and as a result, it is not easily searchable. Retention requirements for different data sources are not standardized. In manual forensics, the investigation focuses on the residue of the attack, and not because examining residue is ideal—far from it. There is simply no other data available. Manually examining collected data is difficult, is time consuming, and dramatically increases the cost of the investigation, ironically resulting in more incomplete results than those obtained in higher maturity models.

Drawbacks to manual forensics:

- Often impossible to ascertain the exact compromise vector.
- Uncertainty surrounding circumstances of the breach. Did lateral movement or privilege escalation occur? How far did attackers penetrate into the network? If manual forensics can answer the question at all, it will be extremely manual and costly.
- Difficult or impossible to determine what and how much data was accessed, viewed or exfiltrated. This vagueness can mean the difference between reporting a compromise of 1 million (uncertain) or 10,000 records (certain).

Basic Forensics

At the basic forensics maturity level, some forethought was given to building a monitoring program. Most organizations at this level will have security information and event management (SIEM) to standardize retention and correlation of logs. Some amount of packet data is available either through rudimentary capture tools or through more sophisticated, security analytics solutions, though probably not the 30–60 days' worth most investigators would prefer. Even if that much network capture data is available, it may not be readily searchable and much of the context for that data is unavailable. Common problems include incorrect sensor placement (resulting in some data not being collected) and the collection of encrypted traffic.

With basic forensics, more data is certainly available for an investigation. Firewall and IDS data may be forwarded to the SIEM and network data is used to enrich these alerts, weeding out false positives. Some threat intelligence data may also be used to provide context to network data. Investigations at this maturity level are much easier and more complete. With sufficient analysis, the initial vector of the compromise can often be determined. Determining the amount and type of exfiltrated data, however, usually is not possible because attackers most often use encryption for command and control (C2) and data exfiltration.

Basic forensics is the maturity level attained by many larger organizations. Unfortunately, several of these organizations believe they have fully matured their IR capabilities and will not benefit from additional service changes. They are wrong—much more can be done to ensure more complete, accurate and cost-effective investigations.

PCAP or it didn't happen.

Some organizations find themselves overwhelmed by the amount of work required to transition from manual forensics to basic forensics. Just because an organization can't immediately transition all technologies to a basic forensics maturity level doesn't mean it can't add value by implementing just a few. The first things that should be on the organization's road map in making the transition from manual to basic forensics are ongoing network packet capture (PCAP) and a SIEM. Most organizations get a SIEM fairly early, but they fail to see the value of packet capture at this stage. That's a huge mistake because network PCAP provides ground-level truth in an investigation. It's so critical to investigations that investigators regularly say "PCAP or it didn't happen"[4] when discussing competing hypotheses in a case.

Proactive IR

The highest IR maturity level is proactive IR. At this maturity level, network data is available and considered complete. Network sensors are placed in appropriate locations to capture every packet of relevant traffic and every available piece of metadata. Network traffic is decrypted as necessary to ensure visibility. The data is indexed and searchable, and high-quality threat and reputation data is automatically applied to the traffic to provide adequate context. This approach differs from the basic forensics level not only because the data is more readily searchable, but also because all network data is available and automatically enriched with reliable threat and reputation data. No network data is lost due to SSL encryption, inappropriate sensor placement or packet loss.

Reputation and threat intelligence data takes care of the knowns, but what about the unknowns?

[4] "QOTD—PCAP Or It Didn't Happen," www.novainfosec.com/2014/05/08/qotd-pcap-or-it-didnt-happen/

At the proactive IR maturity level, unknown data (web pages, PDFs, email attachments, etc.) are also automatically investigated. These files are automatically sent to a sandbox system where they can be executed (or "detonated") for investigation, but we're not talking about your grandma's sandbox system. Proactive IR is all about feeding analysts relevant alerts. The sandbox should be configurable to include at least the organization's "golden image" and even include any other specific configurations used by the organization. Analysts no longer need to waste time on attacks that don't target their OS and app environment or are not effective against the security mechanisms employed at the organization. For those found to be malicious, the sandbox system automatically feeds data back into the network system so that previously unknown network data can be tagged as malicious and blocked at the gateway.

A Sample Incident Response

To underscore the difference an IR maturity level can have on an investigation, we'll walk through the compromise at NoName Bank, a large and fictional regional bank serving the southeast U.S. Like most regional banks, NoName is an attractive target. Its websites are regularly attacked, and it receives phishing emails every day. During the past two years, the sophistication of the attacks has increased, and many more spearphishing attacks have occurred. Employees have even found random USB drives in the parking garage at the NoName building.

In early February, NoName investigates a larger-than-average number of fraud claims. Bank authorities believe a compromise has occurred, but they are not sure of the source of the compromise. At this point, NoName begins its IR process. The following sections examine the IR process for the same incident, using the different IR maturity models.

Manual Forensics

Because NoName does not capture network traffic or historical net flow data, it begins capturing network traffic at the perimeter to try to identify the compromise source. Two IR analysts pore over captured traffic for days, performing spot analyses in an attempt to find indicators of compromise (IOCs). Analysts twice believe they have found IOCs, but later they determine these to be false positives. They have no historical traffic with which to compare current data, so false positives are a virtual certainty.

A week into the investigation, NoName analysts discover three different variants of malware on three different machines. Eliminating possible false positives in network traffic was difficult, and when the machines are finally discovered, analysts take full disk images and perform lengthy analysis to find the specific malware. Some false positive network anomalies are conclusively excluded only after taking disk images and performing known hash exclusions, conducting autorun analyses and manually checking the reputation of suspect binaries. All infected machines are in administrative areas and do not have specific access to financial data. All malware samples are commodity. Because they lack a centralized log repository, NoName analysts must work with different departments to obtain log data. They coordinate with the workstation team for event logs, the server team for the domain controller logs and the network team for any firewall logs that might be available. After a lengthy process of analyzing log data from the individual machines, NoName analysts determine the infections, while significant, are not the source of the fraud observed.

Weeks into the investigation, NoName analysts finally discover that three machines in the credit underwriting department are regularly communicating with a strange web server using HTTPS. Because the traffic is encrypted, analysts cannot determine specifically what data is being exfiltrated. Analysts again coordinate to obtain logs to investigate the three machines. This time analysts discover the same malware on all three machines. The event logs on the machines are incomplete and, in any case, do not cover the time of the initial compromise, which filesystem timestamp analyses indicate occurred just before the Christmas holiday weekend.

NoName analysts remove the active malware and rebuild the machines in question. However, because they cannot determine which data the malware accessed, they send breach notifications and reissue credit cards for all 2.2 million customers and provide them with credit monitoring for a year. Notifying customers of a breach caused substantial brand damage and was a nightmare for the PR team.

Basic Forensics

NoName analysts begin the incident response. NoName uses a SIEM to centrally manage logs and manually forwards unknown files delivered via web or email to a malware sandbox. This approach allows NoName to detect the three infections that misled them in the manual forensics example and remediate them before the attackers can move laterally, saving valuable time during the investigation. NoName has been capturing packet data using a security analytics tool and has 14 days available online for the investigation. Analysts begin the investigation by examining a single day of network traffic. They easily discover several suspicious traffic patterns. However, unlike with manual forensics, analysts quickly identify them as false positives because they have historical packet and net flow data available.

Using manual forensics, it took analysts weeks to identify the compromised machines in the underwriting department due to delays in acquiring and analyzing data. At the basic forensics maturity model, however, analysts discover the compromised machines with only two days of investigation.

After the infected machines are discovered, the picture is much different than what was seen with manual forensics. Using stored net flow and packet data, analysts are able to determine the activity patterns of the attackers. Analysts can also estimate the amount of data that was exfiltrated. Whereas analysis may have suggested that the attackers had access to data for all 2.2 million customers, it's clear from the relatively small amount of data exfiltrated that a wholesale database dump was not performed. Net flow analysis confirms the first communication with the attacker's domain was on Dec. 21.

Unfortunately, however, the analysts have packet data for only 14 days and cannot see the source of initial compromise. Given a date to begin investigating, NoName queries the SIEM for event log data from the domain controllers to find suspicious logon patterns to the compromised machines. They find a department manager made Remote Desktop Protocol (RDP) connection logons to each of these machines on the infection date after hours and from a VPN connection. Using net flow data queried from the SIEM, NoName analysts still see unusual traffic patterns between the VPN user and the three compromised machines. When they examine packet data to gain additional context, though, they discover that sensor placement captures VPN traffic while it is still encrypted. NoName attempts to analyze the C2 and exfiltration traffic in available packet data, but efforts are hampered because the communications are encrypted.

What a morale blow to the NoName analysts. They don't have enough packet data to fully support the investigation due to storage limitations, and they can't use much of what they do have because of poor sensor placement and encryption.

NoName analysts query the SIEM to determine which customer records the attackers accessed. Although auditing is insufficient to show specific customer records accessed, the level of detail stored at the SIEM sufficiently indicates that some databases were never accessed at all. This discovery reduces the number of accounts potentially breached from 2.2 million to 300,000. Even better, NoName analysts complete the investigation within days rather than weeks. Because the data was stored in standard repositories (packet capture), the investigation is repeatable and chain of custody is easy to establish. More important, the investigation is more complete. The analysts are able to determine the machine originally compromised and follow up with more in-depth analysis. Using manual forensics techniques, this machine (still controlled by the attackers) was not discovered. Like many organizations, NoName failed to successfully remediate the incident on the first attempt using manual forensics.

Proactive IR

NoName begins the investigation by querying network data for established communications with destinations that have low or absent reputation data. Assuming a breach, they sort by the total amount of traffic sent outbound during the past 30 days. The attacker's server is the sixth destination they examine. Using proactive IR, NoName analysts identify the compromise three hours into the investigation.

NoName analysts again identify the RDP connections, but in this scenario, NoName has full packet data for 60 days. The sensor locations were carefully chosen to decrypt network data in transit from VPN users. After acquiring the private keys from the compromised machines, NoName is able to use RDP Replay[5] to see exactly what the attackers saw and did when they compromised the machines in the underwriting department. Some attackers use RDP, VPN, and webmail once they have sufficient access to the network. In these cases, the ability to quickly target anomalous behavior is paramount. Finding this specific anomalous data of interest was possible only because the packet data is indexed and can be easily searched. In some cases, artifacts can be viewed as they are captured, which saves a step and allows for faster review and analysis.

NoName analysts don't need to query the SIEM for incomplete database audit data; they can say with certainty that only 25,000 customer records were compromised.

Packet data is decrypted, so analysts have full visibility of all HTTPS traffic, including the malware C2 and data exfiltration. NoName analysts have a complete, plaintext record of every command issued to the malware and every byte of data exfiltrated. There's no need to query the SIEM for incomplete database audit data; NoName analysts can say with certainty that only 25,000 customer records were compromised.

The source of the RDP connections, a laptop that frequently travels outside of the organization, has been repeatedly seen attempting to communicate with low-reputation destinations. When the reputation is known at the time of the communication, the attempts are blocked. In other cases, alerts were generated when past communication with malicious destinations was discovered. Multiple malware sandbox alerts were also generated from this laptop. Bottom line: NoName analysts are able to see a pattern where the laptop owner, the department manager, was repeatedly targeted until compromise was finally successful. These attackers put the "persistent" in advanced persistent threat (APT).

[5] www.contextis.com/resources/blog/rdp-replay/

With full knowledge of the compromise scope, NoName analysts turn their attention to the initial compromise vector. Using packet data, they discover the user downloaded a malicious Java applet from a landing page on Dec. 20, just one day before the machines in the underwriting department were compromised. The HTTPS referrer indicates that the user clicked a link from a web-based email solution. It appears the user may have clicked a link in their webmail, but NoName authorities want to understand what enticed the user so they can tailor new phishing training. Fortunately, NoName can view the specific email that targeted the user because they implement SSL decryption and the email was viewed while the machine was on NoName's network. NoName is also able to download the specific exploit used, even though it was delivered over HTTPS, and submit it to their sandbox vendor for additional tuning.

In the proactive IR scenario, NoName analysts are able to fully assess, analyze and mitigate the compromise in hours, not days or weeks.

In the proactive IR scenario, NoName analysts are able to fully assess, analyze and mitigate the compromise in hours, not days or weeks. Unlike the other scenarios, NoName doesn't guess at the root cause of the compromise—they know what it was. They can implement corrective training and tune their malware sandboxes to determine why detection originally failed for this exploit.

In all three of these IR maturity levels, the investigation began at the same point, but proactive IR alerted the investigators to signs of a compromise earlier by offering fewer, higher fidelity, actionable alerts. With proactive IR, attacks may be detected and blocked in real time, eliminating the cost and time to remediate a full-blown intrusion had the attackers been allowed to fly under the radar of countless low-quality alerts.

Organizations do not simply move from the manual forensics maturity level directly to proactive IR. Proactive IR requires tools, processes, and training. Organizations often examine tools and try to buy best-in-breed options for each capability. With proactive IR, however, best in breed isn't really best if it doesn't integrate with tools already owned and operated by the organization. With proactive IR, integration is just as important as capability (if not more so).

Buying the right tools is not enough. Proactive IR requires getting the right personnel on the team as well. Although proactive IR makes investigations easier and more complete, that only happens with highly trained and experienced staff. Organizations all too often fail to budget for staffing and training. Having more visibility in a network is always a good thing, but someone has to investigate the alerts generated by monitoring tools. There are indications that the IRS received[6] alerts before they subsequently identified that 104,000 taxpayer accounts were compromised. Whether these other alerts were miscategorized or simply ignored is unknown, but either way a personnel (rather than technology) failure seems to have contributed to the breach. From what is publicly known so far, it is unlikely the IRS had a proactive incident response. Maybe the IRS would have been better positioned to deal with a lower volume of alerts if it had implemented a proactive IR program.

Proactive IR Ingredients

Organizations can't simply will a proactive IR program into existence. They must plan a balance of skill and capabilities to be ready for proactive IR. Key components include the following and are illustrated in Figure 1.

Adequate capacity planning. Nowhere is capacity more challenging than with packet data, which offers ground truth in an investigation, whereas other solutions simply support log-based investigations. Given today's complex investigations, not capturing packet data is absolutely negligent.

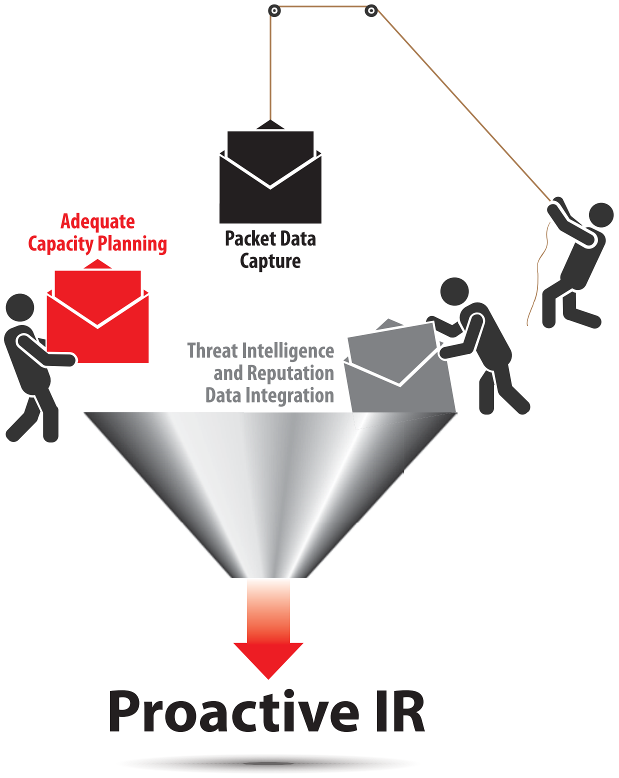
Adequate Capacity Planning
Packet Data Capture
Threat Intelligence and Reputation Data Integration
Proactive IR

Figure 1. Successful implementation of proactive IR requires several ingredients.

[6] "IRS Investigating $39 Million in 'Suspect' Tax Refunds After Data Hack," www.washingtonexaminer.com/irs-investigating-39-million-in-suspect-tax-refunds-after-data-hack/article/2565442

Packet data capture. Capturing "some" packet data is not sufficient. At a minimum, organizations should capture 30 days' of packet data. Sixty days' worth is even better. Network bandwidth usage almost always increases over time, and usage spikes are often observed during incidents. There is no telling when the packet data will be the difference between breach notification letters going to 2.2 million or 25,000 customers. Capturing packets before an attack can be beneficial and done fairly simply using widely available tools.

There is no telling when the packet data will be the difference between breach notification letters going to 2.2 million or 25,000 customers.

Threat intelligence and reputation data integration. This integration is a critical component of proactive IR, but all threat intelligence is not created equally. Reputation data will sort a destination into one of three classes: known good, known bad or unknown (or uncharacterized). With most threat intelligence feeds, the last category is the most common. If a threat intelligence feed has a limited customer base, it is unlikely to classify large numbers of destinations. All things being equal, the more customers a given feed has, the more likely it is to provide additional (and valuable) context to an investigation. Carefully evaluate your threat intelligence feeds before purchase or contract renewal.

Conclusion

Proactive IR is clearly the better choice for organizations that want to save money and better protect data. While reaching a state of fully proactive IR is indeed a journey, it starts by shortening the time to detection. After a compromise is detected, proactive IR ensures containment and remediation are not only faster than traditional approaches, but they are also more complete. As shown in this whitepaper, any investments in building a proactive IR program are easily offset by being able to quickly and accurately determine the scope of a compromise, something that can pay real dividends when it's the difference between assuming the worst and knowing what actually happened.

After a compromise is detected, proactive IR ensures containment and remediation are not only faster than traditional approaches, but they are also more complete.

A key component of a successful proactive IR program is network packet capture enriched with high-quality threat intelligence and reputation data. Organizations that do not currently have network packet capture should look to adopt this technology. Those that already have some packet capture should ensure their coverage of their environment is complete and begin decrypting SSL.

Proactive IR doesn't just happen. Only through careful planning, training and integration can the benefits of proactive IR be fully realized.

About the Author

Jake Williams is a SANS analyst, certified SANS instructor, course author and designer of several NetWars challenges for use in SANS' popular, "gamified" information security training suite. Jake spent more than a decade in information security roles at several government agencies, developing specialties in offensive forensics, malware development and digital counterespionage. Jake is the founder of Rendition InfoSec, which provides penetration testing, digital forensics and incident response, expertise in cloud data exfiltration, and the tools and guidance to secure client data against sophisticated, persistent attack on-premises and in the cloud.

Sponsor

SANS would like to thank this paper's sponsor:

BLUE COAT®