The 13 Deadly Sins of APT Incident Response — Part 1

blogs.blackberry.com/en/2022/07/the-13-deadly-sins-of-apt-incident-response-part-1

Mark Stevens, Rocky De Wiest, The BlackBerry Incident Response Team

When you become aware of an attacker within your network, it can feel like the sky is falling.

Incident response (IR) puts your security team, operations team, and executive team under extreme pressure, as the response process involves many elements of crisis management. This is especially true when fighting an Advanced Persistent Threat (APT) because multiple attackers may be actively battling against you to complete their objective.

The biggest difference when dealing with an APT instead of just another malware is that *you are typically facing a group of humans* that can react to your response. By definition, an APT can employ continuous and sophisticated techniques, implemented by highly skilled attackers. This makes it crucial to avoid what we call “the 13 deadly sins of APT incident response.” In this blog, we’ll discuss just the first four on our list of 13 fatal mistakes.

Real-World Incident Response Lessons

We created this list based on real incident response work conducted by the BlackBerry Incident Response team.

Many in-house security teams have limited opportunities to learn from real-life incidents. It's often the case that one big event causes those internal team members to seemingly "go from zero to 100MPH" overnight, in terms of their exposure to a major incident. When this occurs, organizations frequently reach out to us to help guide and augment their efforts. With more than 100 years of combined IR experience, the BlackBerry IR team encounters major incidents on a regular and recurring basis, giving us a significant advantage in dealing with an extremely broad array of threats and threat actors. Because we act as an extension of the in-house team protecting each client organization, we are often in a position to observe and coach less experienced teams and individuals. This provides insights into frequent "opportunities for improvement," many dealing specifically with APTs, which we have distilled into this list.

Any IR training course will tell you what steps to follow, but not always the "why" and the "what could go wrong?" This article explores these deeper questions through specific, real-world challenges, and suggestions on how to resolve them. Lessons learned "the hard way" are valuable, but it's often preferable to learn from other people's mistakes, when possible, especially when the stakes can be so high.

This blog is the first in a 3-part series that will cover important aspects of preparation, active response, and recovery and aftermath. In this first installment, we'll focus on the first four deadly sins, all of which relate to preparing for an APT incident.

Preparation

Whether you prefer NIST's incident handling guidelines, or the SANS Institute's PICERL model, both frameworks state that "lessons learned" is an important aspect of any incident response effort. But have we really learned our lessons after incidents?

Security vendors—BlackBerry included—love to talk about how they discovered the latest malware sample, but we rarely hear about opportunities for learning when things do not go according to plan.

As full-time incident responders, we leverage our combined experience to plan and execute what we believe is the best strategy for any given situation. But these situations are dynamic, and no two incidents are ever quite the same. Sometimes we are called late to the party, inheriting scenarios that are less than ideal, and while the vast majority of our clients follow our advice, there are some that may have to take other factors into account, causing us to deal with a few curveballs.

Cyber Threat Actor Categories

Currently, we see various types of actors within the threat landscape, ranging from inexperienced “script kiddies” to professional ransomware crime syndicates, to well-funded and organized nation-state actors. None of these adversaries should be dismissed. As we saw from the recent leak of data from the Conti cybercrime group, successful ransomware actors generate significant profits that they often re-invest in their operations to make them more robust. For the purposes of our list of sins, we’ll focus on sophisticated threat actors conducting targeted, human-operated attacks, where defenders may face reactive countermeasures in response to any actions they take.

The List: Deadly Sins of APT Incident Response

As we stated previously, the first four “sins” on our list of 13 are all about preparation. These are related to things that all companies should do before they have any indication of a cyberattack, so that when they do, they already have a plan ready to put into action, with the necessary information at their fingertips. Building IR capabilities on-the-fly is time-consuming, and stressful. You must do this beforehand because the most valuable commodity during an incident is time! That’s why failure to adequately prepare in any of the following areas is just plain sinful.

The First Deadly Sin of Incident Response: Lack of Visibility and Historical Logs

While cyber awareness is resonating more with both SMBs and enterprises, cyber IR preparation unfortunately still ranks low on many organizations’ agendas. This is especially likely if a company has yet to experience a significant incident.

However, when that first major incident suddenly occurs, the first thing responding teams typically reach for are logs that can provide crucial visibility to understand the extent of the threat, and help them ascertain what may have already been compromised. So it is critical to enable proper logging—and centralized storage for those logs—when preparing for APT attacks. One reason is that these threat actors often “dwell” in a network for several years prior to eventual discovery. In fact, the longest dwell time our team has encountered is seven years. Imagine possibilities of what an adversary could do in your environment during 364 weeks of unimpeded access.

In addition to siphoning away your data for an extended period of time, threat actors could also modify your data to suit their needs. Adding an entry to a database, perhaps to provide unauthorized access to a secure facility, is just one example. The type and amount of mayhem a successful APT actor can commit is limited only by their imagination.

Maintaining good logs isn’t glamorous, or even very interesting—until you find an APT in your network. Security operation center (SOC) teams are often inundated with a never-ending slew of alerts to be resolved, and can sometimes forget the importance of historical

logs. But these records are what allow an incident response team to perform forensic analysis on not only present, but also past events.

While storing seven years of detailed logs is not viable for most organizations, the span of one to three years should be considered, depending on your organization's requirements and retention policies.

As consultants, we quite often come into a customer environment and find that all system logs are capped at two or three days, and are stored on the endpoints themselves. This creates a significant problem in historical data and can raise questions about the integrity of the log entries themselves. Threat actors can manipulate logs if they want, or just simply delete them.

One of the frequently overlooked logs, which nearly every network device can create, is "NetFlow." This log can assist during an incident, answering critical questions such as, "Was data exfiltrated?" or, "Which systems are connecting to the APT's command-and-control (C2) server?"

Other logs that can be centralized and indexed include PowerShell, Windows Event, Network, and AV logs. Attention should also be given to what is logged. This is often configurable, and your choices can have a considerable impact on storage costs, among other factors. While any IR team would be grateful for every logging option under the sun to be switched on, it may make more sense to have enhanced levels of logging only on more critical devices. Cloud-based EDR/XDR and managed XDR solutions are available that automatically move events to a secure, centralized source that can still be accessed reliably and analyzed, even if the source data is tampered with or destroyed.

Questions Asked During Incident Response

During an incident, clients often ask questions such as "How and when did the attacker get in?" "What was stolen?" and, "Was any data altered?" While these questions often get progressively more difficult to answer, it is significantly easier with evidence that has integrity and is centrally located. Think of it this way, if the crime was a physical theft, the case could often be solved by having CCTV footage of the incident, from start to finish. Think of cybercrimes in the same way: Your logs are key evidence of what happened and when. If there are no logs, you have a blind spot, and you can't go back in time to fix that omission after the crime has occurred.

When making choices about what to log, and how long to keep those records, ask yourself, "What is the most precious information in our environment, and where are those 'crown jewels' stored?" These locations, and any related systems, should be covered by multiple "cameras." In our case, this means supporting multiple logs, along with toolsets like SIEM and EDR, with rules and alarms configured. This approach also requires having a "security guard" on duty, as all the alarms in the world will not help you if no one is there to hear and

respond to them; for us, this means having either a 24/7 SOC or a managed security service on alert at all times. Unfortunately, our team has seen many easily manageable incidents cause severe damage to organizations, simply because no one was on guard to investigate the alarms.

How APTs Attempt to Hide from Network Defenders

Forensic investigations can often provide a good understanding of what happened during a cyberattack. However, the more a specific machine was used by the attackers, and the more time that has passed since the activity took place, the slimmer the chances of recovering valuable artifacts.

More advanced threat actors are often extremely careful about leaving evidence behind—by deleting logs, time stomping evidence, or using other covert techniques. They leave smaller footprints behind to follow. They practice stealth to prolong their access for as long as possible. This is what allows them to spend years in an environment before being discovered.

In the unfortunate event that a particular threat actor has been lurking in your network for years, having centralized logs for an extensive period can potentially answer your investigatory questions. When dealing with ransomware cases, and the now common double extortion techniques where data exfiltration is commonplace, having certain forensic data (such as logs) moved and securely stored off of the endpoint can make all the difference. This is because when devices are encrypted by attackers, logs and forensic artifacts are often encrypted in the process.

We have worked cases where the on-premises SIEM solution, which was used to store system logs remotely, was encrypted along with the rest of the servers. This means securing the SIEM solution itself is also part of the preparedness challenge.

Organizations can improve their defenses by ensuring standard security practices, such as having a disaster recovery plan worked out, segregating the SIEM solution to strictly allow administrative access only from the security team with multi-factor authentication (MFA) enabled, and ideally, using a cloud-based Linux® solution.

Benchmarking Incident Response Preparedness

Matt Swann's Incident Response Hierarchy of Needs is an excellent reference when trying to determine if you are indeed fully prepared for a major incident. The odds are best if you can confidently answer the questions shown below. (Pro tip: Start at the bottom and work your way up.)

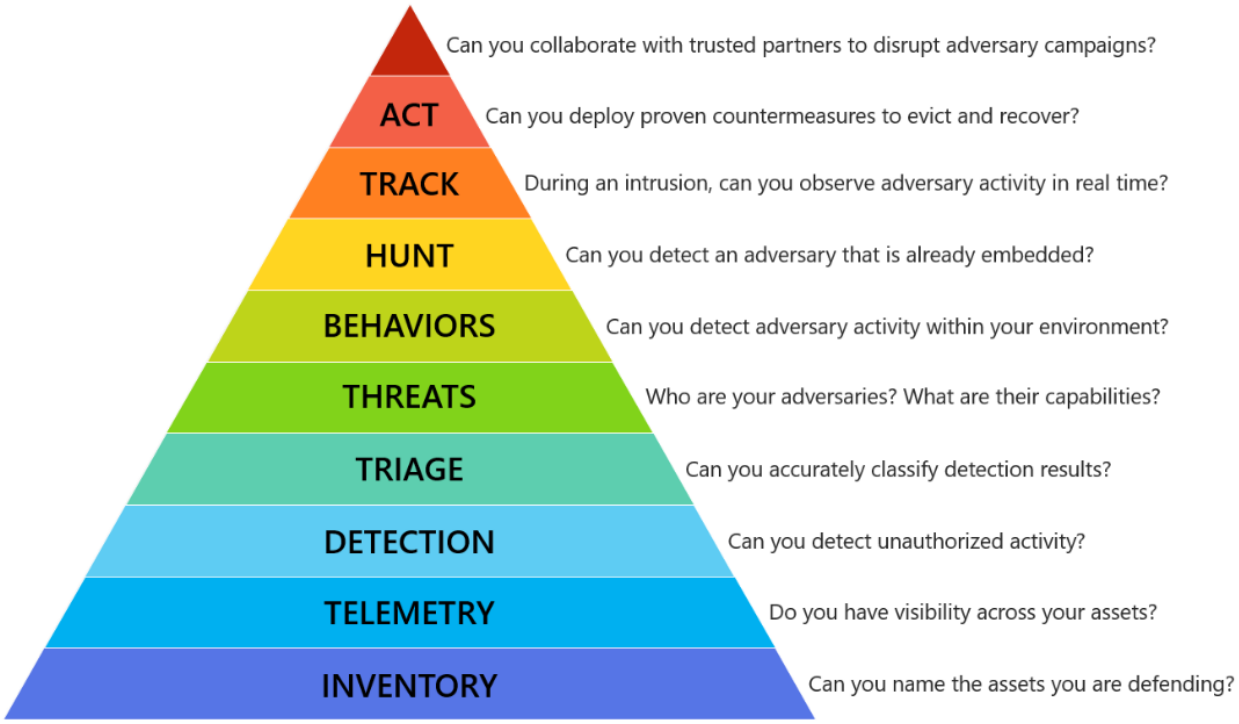

The Incident Response Hierarchy of Needs.

The diagram is based on the well-known principles of Maslow's hierarchy of needs, borrowed from the field of psychology. An organization should start at the base of the pyramid to find out if their primary security needs are covered. Having an asset inventory or a network diagram is the first step in prepping for APT-style attacks. In fact, eight of these 10 questions should already be covered if you have properly implemented EDR and SIEM systems (along with the appropriate staff and processes to operate and maintain them, of course).

Increasing Visibility to Detect APT Attack Vectors

One of the first goals of investigating an APT is to increase visibility. There are multiple reasons for beginning with an increase in visibility:

- Unless you catch it very quickly, whatever you, your security team, or an external third party initially discover will probably just be the tip of the iceberg.
- Fully scoping the incident requires visibility of, and the ability to hunt for, indicators of compromise (IoCs) across the entire enterprise.
- Some APT groups rely on multiple "footholds," or mechanisms to maintain access, to persist in a victim's environment.
- It's important to identify secondary or tertiary backdoors deployed with stealthy installations. These are often left dormant and used to re-establish the threat actor's presence at some point in the future, if the primary access point is lost.

The Second Deadly Sin of Incident Response: Not Having the Right People, Processes, and Technology

Military troops are trained so that no matter how stressful the situation is, they will follow procedure, which ensures that responses are appropriate, strategic, and consistent. The same logic applies when responding to incidents; however, procedures and policies are quite often written after a company's first major incident, when it is already too late. When such plans are in place before an incident occurs, we may find that they have never actually been tested. In too many cases, we see inexperienced but well-meaning staff often will tamper with an infected system and try to self-resolve the incident, before taking a forensically sound image. They may even format and rebuild the affected system in an effort to restore to "business as usual" as soon as possible, and in the process, destroy valuable evidence and information about the attack.

While reimaging is often a best practice for resolving commodity malware affecting a single host, when it comes to more serious incidents like APTs, that evidence is key to understanding the impact of the attack. You often don't know what is critical until the incident is fully scoped and analyzed. If you track an attack back to what is likely to be the initial entry point (also known as "patient zero") only to find that it was reimaged by an overzealous IT security team, it can have serious repercussions on your ability to understand the attack, and how to prevent it from happening again.

In one example, a system admin with limited security experience attempted to resolve a previous incident by reformatting the affected system, without first assessing the possible impact of the attack. The system in question was a development server that was exposed to the internet. Three months after the initial incident, the attacker came back using stolen domain credentials obtained from that previous server. This breach resulted in the organization's entire domain being ransomed, and a large amount of information being stolen.

In similar APT cases, this same technique has caused what we often call a game of "whack-a-mole," where backdoors are found and "resolved" by re-imaging the machines, resulting in what companies often believe to be multiple repeat incidents. However, a thorough investigation often reveals that they were all part of one large incident where the attacker was never successfully removed.

The key to dealing with security incidents is catching them early, scoping them fully, and having a plan in place to resolve them quickly and appropriately. Implementing these best practices requires the ability to differentiate between rudimentary malware and more complicated human-operated attacks. Having an IR process in place for both types of attacks will help mitigate the damage.

Another challenge is that the experience necessary to identify, investigate and respond appropriately to an APT-based attack comes at a cost. It's hard to keep current on the laundry list of threats active and "in the wild" on any given day. Analysts gain their experience through dealing with real-world incidents, and their skills can get stale when not dealing with incidents on a regular basis.

This increase in visibility, as well as the later remediation phases, often relies on the ability of organizations to deploy software, configuration changes, and scripts across their entire estate. Ensuring you have this capability ahead of time can make all the difference.

Even if you purchase IR services, they will require you to install agents or run scripts across the whole enterprise to start feeding the external IR team the data they need to perform their initial analysis. There is nothing more frustrating for a CISO, CIO, or security manager than having to keep an external team of IR consultants waiting for data before they can start to investigate and understand what happened. The clock is ticking, response is delayed, and in the meantime, potential damage can spread.

The Third Deadly Sin of Incident Response: Analyst Burnout and Alert Fatigue

Alert fatigue happens when analysts are exposed to a large number of alerts over an extended period of time, and it's made worse when these alerts include frequent false positives or low-priority alarms. In these cases, analysts can become desensitized over time. Whenever SOC members are not looking at only high-fidelity alerts, and not focusing on continuous tuning, it can delay the identification of critical events when they do occur.

For example, on multiple occasions, BlackBerry IR consultants found that a target organization's antivirus or EDR tools had triggered multiple, highly suspicious alerts for weeks before the incident became critical. However, no importance was given to these alerts because the files were quarantined, or activity was blocked.

Understandably, it can become very tiresome if every single alert must be investigated. This is why Florian Roth created a very helpful cheat sheet to judge how relevant each detection is. While this resource provides an excellent starting point for recognizing legitimate alerts, you will also need to develop your own tuning and prioritization process, along with playbooks appropriate for your organization.

Again, experience plays an important role. When put in front of an experienced threat hunter or IR consultant, many of the warnings would raise immediate concerns and trigger an investigation to look for the root cause. But a SOC analyst who has never experienced a major incident may not have the ability to identify these in a timely manner.

The Fourth Deadly Sin of Incident Response: Not Having an Incident Retainer

Another step organizations should take ahead of a data breach is purchasing an incident response retainer (IRR) with a trusted security advisor. When an incident occurs, you don't want to waste time scoping out which company is best placed to help resolve your incident, and then spend valuable time collecting quotes, or executing contracts and purchase orders to start the response and remediation process. With the current demand for incident response services, vendors might be fully occupied and give priority to their existing retainer clients. An added benefit of having a retainer is that it usually comes with a discounted price tag.

A retainer can speed things up and provide access to world-class experts, intelligence, and tools, when and where you need them most. Retainers often include access to the vendor's products for the duration of the incident. Even if some of the skillsets within a retainer are available in-house, it's advisable to have expert help available when you need it.

Assessing and Absolving Your Sins

When we classify these sins as potentially "deadly," it may not be sheer exaggeration, especially for smaller and mid-size organizations. BlackBerry's recent 2022 Threat Report stated that more than 70% of SMBs have faced cyberattacks, and of those attacked, many cannot recover from a serious attack and ultimately close their doors. An APT attack could be existential for some organizations, while for others, it may be survivable but still catastrophic. The main point is that there are measures you can take now to minimize your risk and to not do so would indeed be a sin.

We hope this information helps you reflect on whether you or your organization may have committed any of the transgressions noted above, and we at BlackBerry are always available to help you prepare and benchmark your incident response planning, for the long-term health and success of your organization.

In Part 2, we will share more "deadly sins," based on incident response lessons learned when actually experiencing an APT attack.

About BlackBerry Incident Response Services

BlackBerry® provides incident response that can help you mitigate the impact of any breach, ensure your recovery follows best practices, and secure your IT environment for the future. Our cybersecurity experts provide answers to your questions so you can protect your IT environment during the current attack and defend against future cyberattacks.

For similar articles and news delivered straight to your inbox, ***subscribe to the BlackBerry blog****.*

About Mark Stevens

Mark Stevens was Technical Director of Incident Response at BlackBerry.

About Rocky De Wiest

Rocky De Wiest is a Principal Incident Response Consultant.

About The BlackBerry Incident Response Team

The BlackBerry Incident Response team is made up of world-class consultants dedicated to handling response and containment services for a wide range of incidents, including ransomware and Advanced Persistent Threat (APT) cases. If you're battling malware or a related threat, you've come to the right place, regardless of your existing BlackBerry relationship. We have a global consulting team standing by to assist you, providing around-the-clock support where required, as well as local assistance. Please contact us here.

Back