CYBER STASH

How To Tell A Good Post-Breach Cyber Incident Response Plan From A Bad One

eclipse

The Problem

The standard cybersecurity incident response process that we have all come to know must be better utilised to minimise business impact following a breach. That process was originally derived from frameworks developed two decades ago but now remains largely unfit for responding to modern-day breaches.

The incident response process hasn't radically changed over the years even though the tools available to responders have matured and the attackers have shifted and improved their methodology during an incident. If we're going to confront the realities of having to respond to the sophisticated nature of today's attacks that breach our systems and impact business, we must accept that we can't protect every system 100% of the time, and there's an unknown level of vulnerability to manage when a breach occurs.

The Challenge Ahead

One of the biggest challenges organisations face is in attempting to manage post-breach response activity using an incident response plan that presumes it's possible to ascertain the full impact of incidents. This presumption leaves gaps in the level of response that organisations believe they require and presents an opportunity for the incidents to reoccur – and they do.

Most major breaches are identified after business impact.

To believe that a typical security team and the tools and processes they utilise can detect, contain, and eradicate all sophisticated attacks within a period of time that would avoid impacting business is, at best, a framework-aligned, fingers-crossed, wishful-thinking, proposition.

The motivation behind the methodologies used today to detect threats is the incident response process itself, which requires that organisations first and foremost plan how they will detect an incident. This has, unfortunately, compelled many security practices to focus heavily on detecting and responding to events that, in most cases, do not lead to a breach, and therefore do not eventuate in business impact. The ROI for a threat detection practice that follows this methodology is difficult to justify and rarely attainable. Most organisations rely on pre-breach methodologies for detecting incidents. That is, they use detection engines that attempt to prevent or detect threats in real-time.

Collecting millions of events every day in the hope that you will detect the threat in real-time and will have enough time to respond before business impact takes place, is an event-driven and reactive methodology that has proven itself to be severely ineffective and one that's not optimised for risk and resources. The focus should rather be on how an organisation can quickly detect breaches, not threats, and how they can become proactive in doing so to avoid business impact.

Real-time detection engines such as those used in Endpoint Protection Platforms (EPP) and Endpoint Detection and Response (EDR) products, don't have the luxury of examining every possible event down the kill-chain because they are running in-line with the attack itself. They are all forced to sample between 30 to 150 samples per minute, and when they sample, they leave gaps,and where there are gaps, breaches occur.

Defenders however now have these two advantages that enable them to quickly establish the facts and respond to breaches:

A rapidly accumulating knowledge of how breaches can be forensically detected and validated

The use of recently developed tools that meticulously audit forensic artifacts and analyse systems at scale and with speed.

These advantages allow organisations to establish a periodic breach detection and incident response practice within a time that is meaningful for avoiding or limiting business impact. This should encourage organisations to review their existing incident response plans and consider how these can be enhanced for getting ahead of breaches after they've occurred

Accepting Reality

It's important for security defenders to understand how adversaries operate, and they can do this by reverse-engineering the adversary's steps and reviewing case studies to draw lessons from failed threat detection methodologies. There is a need to accept and adjust to the new reality that a large percentage of reported incidents that impact businesses are detected externally and not by real-time threat detection platforms, and many continue to remain undiscovered.

In most cases, by the time incidents are first discovered, business impact has already occurred. To exacerbate the matter, before anyone begins thinking that they need to respond to a breach rapidly, or by some pre-determined SLA, remember that an adversary has, on average:

Maintained persistence on the company's network for 2 months

Already discovered its high-value information assets

Already exfiltrate data, all while remaining mostly undetected

It should not come as a surprise that 2 out of 3 companies that experience a major incident, experience another one within the next 12 months. Maintaining prolonged persistence on a company's network gives ample opportunity for adversaries to implant several different backdoors that allow them to re-establish their access to the company's networks should their primary means of access be discovered and cleaned up.

1. The reality is, technologies that leverage real-time detection engines, depend largely on having prior knowledge of an attack before they can block or detect it.

2. Many detection engines depend on knowledge of how the malware behaves and what the Indicators of Compromises (IOCs) are.

3. When security advisories are released, the associated TTPs and IOCs are well documented and then used to enhance the defensive capability of real-time technologies.

4. The problem for defenders is that they are always 3 steps behind where they need to be: detecting compromise without having prior knowledge of a particular attack or the IOCs.

Refocusing

Companies should begin by developing a new incident response process that focuses on minimising business impact, one that does not heavily depend on a methodology that attempts to detect threats in real time but rather on a methodology for quickly detecting breaches forensically. A post-breach detection mindset is radically different from a pre-breach one whereby the focus is not on detecting the attack on its way in. How a piece of malicious code makes its way into memory or onto a system should not be of any concern to an effective post-breach detection methodology. Whether sensitive data has been exfiltrated and held for ransom, or a critical service has been disrupted, the goal of the organisation must be to minimise business impact. It's time to put on our business hats and determine what must be accomplished and in what priority to avoid or minimise losses.

To help guide the response plan at the very highest level, companies must pivot their new process so that it answers these three questions:

Can the incident reoccur? The incident could be the same – that is, the same attack exploiting the very same vulnerability or a different vulnerability and causing the same or similar business impact.

What's the impact if the incident occurs again? The impact could be associated with the same information asset or a different one.

How can the level of impact amplify? This could be a direct result of the attacker's actions, actions taken or not taken by the organisation itself, action taken by a regulatory body or actions taken by stakeholders (investors, users, customers, etc.)

When business impact occurs, typically, the security team, whether internal or outsourced, is asked to investigate. A computer forensic analyst is engaged and called in to reverse-engineer the attack in attempts to chain together the sequence of events leading to the incident and ultimately to find the root cause (vulnerability) so that it can be fixed (remediated). If a computer forensic investigator can perform this deep level of analysis, why can't we use the same methodology to detect and respond to breaches quickly before irreversible business impact takes place? The focus of any investigation should not be too narrow, either with the scope limited to the assets or the locations impacted. That's how unknown vulnerableness can resurface. The scope of the investigation must instead include the entire IT environment to re-establish a higher level of trust from business stakeholders.

Scenario

Let's assume that an incident has taken place whereby sensitive data has been breached from an organisation. Should the investigative lens be limited to the systems involved and the data that was exfiltrated? Should the response team start by using the indicators of compromise available to them, working backwards, and following the evidence in an attempt to discover how the attack was successful and to uncover the root cause?

If you are inclined to say yes to these questions, you would be limiting the scope of the investigation, and in doing so, making the following costly assumptions:

Costly Assumption 1

The adversary gained access to the company's environment through a single attack. By tracing back through the line of evidence, we will discover the full story and all kill-chains

Costly Assumption 2

We have found the vulnerability and have mitigated it. The attacker is unable to re-establish access back to our environment. We will detect if there are any other occurrences of the breach because we now know how to discover it in our environment.

Costly Assumption 3

The only data exfiltrated from the environment is that which we know of. The adversary did not exfiltrate any other data and is no longer in a position to cause further damage

A less optimal response doesn't consider the wider potential implications of the incident because it makes the above costly assumptions which leave open opportunities for the adversary to attack again.

In contrast, a well-informed response plan with a focus on avoiding or limiting business impact:

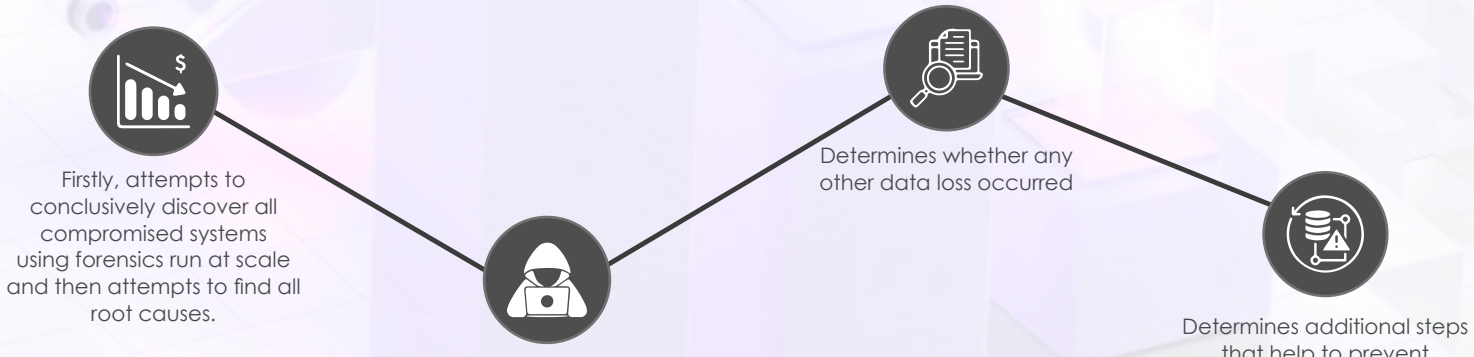

Closing Remarks

Responding to cyber breaches is stressful and challenging for those involved. What helps ease tensions, minimise business impact, and optimise risk and resources, is an independent process that firstly establishes confidence by minimising what's unknown. Remember that you cannot rely on the same tools and processes that allowed the breach to occur in the first place to then validate facts. Specifically, by ascertaining the current compromise state of all endpoints within an environment following a breach, a better sense is formed about the level of response necessary to contain and manage an incident. Then, using this independent post-breach forensic-depth detection methodology in a periodic manner, companies can establish proactive security practice that reduces the likelihood of business impact. This must be achieved proactively through periodic threat hunting as opposed to reactively, whereby detection depends on an event being triggered and having prior knowledge of that event or the behaviour.

At CyberStash, we conduct periodic memory analysis of every endpoint within an enterprise environment, that's how we're able to establish a higher level of trust and maintain it without depending on prior knowledge of the attack. We forensically examine every endpoint proactively to determine its compromise state and respond to breaches before irreversible business impact takes place. The forensic leads we discover are then reverse-engineered through dynamic analysis, code re-use analysis, threat intelligence and human analysis to then map the piece of code back to its origins and to derive its level of risk to the organisation

Explore
eclipse.xdr

CYBER STASH

Become a Reseller
info@cyberstash.com
1300 893 802

Inquiries
info@cyberstash.com
www.cyberstash.com/contact/sales-inquiries