cybereason

THE SEVEN STRUGGLES OF DETECTION AND RESPONSE_

Detection and response is still a challenge for organizations even though they have a variety of security tools at their disposal. So what's keeping companies from detecting and responding to breaches in days or even hours instead of months and years? To answer that question, Cybereason reviewed major security incidents to see if a detection and response issue played a role in how the incident was handled and reached out to security leaders to get their take on the struggles of detection and response.

Even with a raft of security tools readily available, detecting and responding to security incidents remains challenging for organizations. U.S. companies took an average of 191 days to detect a breach in 2017, according to the Ponemon Institute. That's a slight improvement from 2016, when the average was 206 days.

For more evidence, look at how long some companies took to detect data breaches. Take Yahoo, which is now owned by Verizon. In 2016, the former Internet giant disclosed that attackers breached a database and gained access to the encrypted passwords, email addresses and security question answers of 500 million users. While the incident shouldn't have come as a surprise -- no company is immune to an attack -- security experts wondered why Yahoo took two years to detect the breach, which occured in 2014.

Then there's credit reporting agency Equifax, which reported in September 2017 that a breach exposed data on up to 143 million U.S. customers. According to media reports, the attack began in March 2017 and Equifax detected the breach in July, giving adversaries a few months to move deeper into the company's network and cause more damage.

Quickly detecting and responding to a breach provides defenders with an obvious advantage. The sooner analysts spot a compromised machine and remediate it, the less damage there is. But there's another factor that could motivate organizations to accelerate detection and response: the General Data Protection Regulation. GDPR requires any organization that handles the personal data of E.U. citizens to provide extensive information about a breach, including what users were impacted and a remediation strategy, within 72 hours of the incident being discovered. And while GDPR is a E.U. law, it applies to any organization that handles E.U. citizen's personal data, regardless of where it's based. GDPR (and the other regulations it's likely to spawn on the handling of personal data) may encourage companies to look at technologies that will enable them to accelerate detection and response.

So what's keeping companies from detecting and responding to breaches in days or even hours instead of months and years? To answer that question, Cybereason reviewed major security incidents to see if a detection and response issue played a role in how the incident was handled and reached out to security leaders to get their take on the struggles of detection and response.

01. LACK OF ENDPOINT VISIBILITY

Attacks take place on endpoints and, in most cases, serve as the infiltration point into an organization's network. Plus endpoints are the attacker's ultimate target since they hold valuable data. In theory, knowing what activity is occurring on a company's computers and servers should be at the top of a security team's priority list.

But companies often struggle with endpoint visibility. Organizations operate incredibly complex IT environments. There are Windows, Linux and Mac environments. In some companies the IT department completely controls endpoints while others have a BYOD policy. There are remote workers who infrequently connect to the company's network. And there are the tools that security professionals are using to gain endpoint visibility, many of which weren't designed for that task.

Take RTI Surgical, which manufactures medical devices and biological and synthetic implants. The company used antivirus software and a firewall, but neither one provided it with endpoint visibility and couldn't monitor endpoints that weren't connected to RTI's network.

"We had nothing beyond the standard tools, which give you limited visibility, especially if you're performing incident response. I need 100 percent visibility into my endpoints. I need to know that when employees travel I can get a report on what their machines are doing or not doing," said Jeff Wright, the company's security manager.

02. DEALING WITH THE ENDPOINT DATA DELUGE

Even if a company is using a tool designed to provide endpoint visibility, such as an EDR (endpoint detection and response) platform, it may just flood the security team with data instead of providing them with insight on incidents.

That was the case at a Fortune 500 bank with more than 60 million customers across the globe. The EDR tool it was using collected reams of endpoint data from the bank's hundreds of thousands of servers and computers. However, it didn't provide security analysts with any context on this information. Instead, they had to undertake the time-consuming process of manually querying the data to make sense of it.

"Our previous EDR tool just provided lots of data. That's not helpful when you have such a large infrastructure. You just get overwhelmed with data," said the bank's deputy CISO, who wished to remain anonymous.

03. ALERT FATIGUE

Data isn't the only thing overwhelming security analysts. The security tools they're using inundate them with alerts, leading to alert fatigue. When analysts can't distinguish the important alerts from false positives, the entire organization can be impacted.

Look at Target, which suffered a data breach that affected 70 million customers, cost it $202 million and resulted in the CEO losing his job. While this incident occurred in 2013, it clearly demonstrates that alert fatigue can have ramifications that extend beyond information security, a lesson that's still relevant today.

Target's security team was notified about suspicious activity in its IT environment in late November 2013 but didn't action until a few weeks later in mid-December. Alert fatigue lead to the delayed response, according to news stories, which said that Target's analysts likely received hundreds of alerts every day. This rash of alerts made separating the legitimate ones from false positives a challenge.

Judging from two separate studies, false positives, the main contributor to alert fatigue, is a major problem for many enterprises. The Cisco 2017 Security Capabilities Benchmark Study found that, due to various constraints, organizations can investigate only 56 percent of the security alerts they receive on a given day. Of the alerts that are investigated, only 28 percent are legitimate. The rest are false positives, meaning that a substantial amount of time is wasted spent looking into bogus alerts. And a Ponemon Institute survey quantified false positives, claiming that organizations spend $1.37 million each year responding to false positives and waste 425 hours every week investigating them.

04. STRUGGLING TO LOOK BEYOND IOCS

The typical approach for detecting attacks entails looking for indicators of compromise (IOCs). Common IOCs include virus signatures, malignant IP addresses, MD5 hashes of malware files and URLs or domain names linked to botnet command-and-control servers. If any of these are observed on either a network or operating system, a breach has most likely occurred.

But as the deputy CISO at a Fortune 500 bank noted, IOCs are very easy for attackers to change and companies that rely on them to detect threats have limited visibility into all network behavior. "Indicators are an aging thing in security. You have to move beyond them. You have to detect techniques and tools, which are much harder for adversaries to change," he said.

For example, changing hash values and signatures is only a matter of either rebuilding or obfuscating malicious code, which is done automatically in some commodity programs like the Angler exploit kit. Maintaining a supply of IP addresses is easy to accomplish by using botnets, hacked servers, anonymous hosting or a domain generation algorithm (DGA) mechanism, which provides a constant supply of domain names that can host malware and are never reused. As for host artifacts, their URLs can be randomized and payloads can be fully randomized, which is a feature commonly found in malware.

05. FOCUSING ON MALWARE-BASED ATTACKS

Adversaries are increasingly using fileless malware attacks, which essentially turn Windows against itself by using legitimate tools like PowerShell and WMI for malicious activity. Fileless malware attacks don't result in code being installed on a computer so there's no signature for antivirus programs to detect. "Attackers have moved on to fileless attacks. They're always trying to get by your defenses and using Windows tools allows them to do that since those are trusted," said the deputy CISO at a Fortune 500 bank.

While some advanced security tools address fileless malware attacks, the methods they use to block this threat are less than ideal. For instance, some tools block all PowerShell activity and force users to whitelist approved PowerShell functions, a time-consuming process.

06. A LACK OF QUALIFIED SECURITY TALENT

Finding skilled security talent is a perennial challenge for companies. According to a report from Frost and Sullivan, the global security workforce will have more than 1.5 million unfilled positions by 2020. Meanwhile, ISACA predicts there will be a global shortage of 2 million security professionals by 2019.

Making the hiring situation even more dire, security workers are coveted by every business. If these employees aren't getting hired by hot technology companies like Facebook, Apple and Tesla, they're being recruited by security vendors. There aren't enough people to meet the demand.

"One of the things that concerns me in the cybersecurity field is the incredible demand for talent, for people to do cybersecurity, security engineering," said Phil Mazzocco, CSO at Peraton, a Herndon, Virginia, company that provides security services to the U.S. government.

While technology can empower the security teams companies already have and decrease the need to hire more analysts, even the best security tools are ineffective without people who know how to act on the information they provide. Even in an era of machine learning and automation, organizations still need analysts who understand what alerts mean and know how to triage an incident, scope a breach and generate a full picture of the attack.

07. THE GAP BETWEEN DETECTION AND RESPONSE

Figuring out if an organization is under attack is typically a time-consuming, labor-intensive affair. Analysts usually gather data from multiple security tools and then analyze it, a process that can take hours, days or weeks. For example, the security analysts at law firm Cadwalader, Wickersham & Taft had to comb through firewall logs and conduct packet analysis to build an attack picture, said Dimitri Josh, a member of Cadwalader's security team.

"You'd spend two or three hours looking at what's being blocked or what the source addresses are and you'd have to run antivirus and scans," he said.

The plethora of security tools that analysts use is partly responsible for hindering the incident response process, according to a study from Enterprise Strategy Group. Of the 203 security and IT professionals polled, 27 percent said that the amount of time required to conduct incident response increased when their company used security management and incident detection technologies. New tools require staff to learn how to use them, custom configurations and integration with tools that the enterprise already uses. These processes all require time and resources and add complexity to understanding and reacting to security incidents.

But time is at a premium when responding to incidents. The longer adversaries linger in an IT environment, the more machines they can move to and the more data they can access. "The ability to protect yourself in seconds makes a difference," said a security analyst who works at a healthcare revenue cycle management company.

PREVAILING AGAINST THE STRUGGLES OF DETECTION AND RESPONSE WITH CYBEREASON

The struggles of detection and response aren't impossible to overcome. Here's how the Cybereason EDR platform can help.

- **Increased endpoint visibility:** Cybereason offers visibility into all endpoints, including those not connected to a company's network. Analysts can see what processes are running on any machines that run Cybereason.

- **Automated data collection and analysis:** There's no manually querying endpoint data with Cybereason. Instead, the platform automatically collects data from every endpoint, correlates it and uses behavioral analysis to link together seemingly unrelated incidents and spot malicious activity.

- **Complete attack story:** Cybereason doesn't issue alerts on every threat, lowering the prospects for alert fatigue. Instead, it provides analysts with a malop. Malops show the full attack story and present information like the initial penetration vector, what machines were compromised and if the attackers moved laterally to other machines.

- **Behavioral detection**: By looking for attacker behavior instead of IOCs, Cybereason can detect more advanced attacks, like those that use PowerShell to conduct malicious activity, as well as more basic but lethal threats like ransomware.

- **Empower analysts of all abilities**: By presenting a full attack story, analysts are given all the information they need to take action on an incident without using multiple tools. This means that even junior analysts can use Cybereason to tackle more advanced threats.

- **Shortening the time between detection and response:** Cybereason's graph database detects incidents in minutes instead of hours or weeks, allowing analysts to immediately investigate a threat.