PREVENTING THE INEVITABLE: THE NEED FOR RAPID DETECTION AND RESPONSE

You've fortified your defences. You follow industry best practices. You've purchased the latest and greatest technology. Yet attackers still penetrate your defences. In a world where it's expected that attackers will successfully breach your perimeter, what chance can you possibly have to protect your business?

BJØRNAR PRESTAASEN
Head of Security Operations Center
mnemonic

Attackers will breach your defences. This is not a hypothesis, exaggeration or a fear mongering statement – it is a simple, undisputed fact.

Regardless of the preventative measures we put in place, a determined attacker with the right motivation, financial backing and skillset will evade these measures. In some cases, the attackers do not even need to evade your preventative measures – your employees take care of that task for them.

It should be no surprise that there is a direct correlation between the number of users in an organisation and the number of confirmed security incidents the organisation experiences each year. From our 15 years' experience, we find that for every user an organisation should expect to see 0.2 to 0.3 confirmed security incidents annually. That means that for a company with 1000 users, there are an expected 200 to 300 confirmed incidents each year. This does not speak to the severity of the incidents, but serves as an indicator to the immensity of the task security teams are faced with.

By 2020, 60% of enterprise information security budgets will be allocated for rapid detection and response approaches, which is an increase from less than 30% in 2016.

Special Report: Cybersecurity at the Speed of Digital Business, Gartner

PREVENTION EXAMPLES INCLUDE:

- Vulnerability management
- Password management
- Access control
- Inline security products with blocking capabilities (e.g. firewalls, web/email proxy, anti-virus, endpoint protection)
- User awareness training

"IT risk and security leaders must move from trying to prevent every threat and acknowledge that perfect protection is not achievable. Organizations need to detect and respond to malicious behaviours and incidents, because even the best preventative controls will not prevent all incidents.

Special Report: Cybersecurity at the Speed of Digital Business, Gartner

THE ROLE OF PREVENTION

This is not to say that prevention is not an important component of a well-rounded security strategy – quite the contrary. Prevention is a critical capability that fortifies your cyber defences, and represents best practice for protecting your organisation.

Prevention shapes the attackers path, and makes it more difficult for them to infiltrate a network, move laterally, escalate privileges and steal data. Attackers are only human, and are most likely to pursue the path of least resistance to achieve their goals. If it is dangled in front of their face, attackers will go for the low-hanging fruit. However if you cut the low-hanging fruit, hungry attackers will bring a ladder to reach the higher branches, or a chain-saw to simply cut the tree down.

In opportunistic attacks, where the target is arbitrary, deterring an attacker with enough preventative measures may be enough to cause them to simply move on to another target with lower hanging fruit. Or it might not be. This will vary depending on factors such as the attacker's ambitions, skill level, motivation, and because they are human, their mood.

> **"**
>
> The harder we make life for an attacker, the more likely they are to knock on different doors, generate noise, and trip an alarm.

On the other hand, a determined attacker with a clear target and goal will relentlessly raise the sophistication of their attack to match the security maturity of their target. However, the harder we make life for an attacker, the more likely they are to knock on different doors, generate noise, and trip an alarm that allows defenders to detect and respond to their presence – provided there is someone listening for the alarm.

IF AN ALARM GOES OFF AND NO ONE IS THERE TO HEAR IT – DOES IT MAKE A SOUND?

There are a wide-range of techniques that can be used to detect suspicious activity in your environment. You can inspect network traffic, install agents on endpoints, collect logs in a SIEM, execute files in a sandbox, monitor user account activity, uncover anomalies in user behaviour, identify spikes in bandwidth usage, amongst a plethora of other techniques. A vast amount of technologies exist supporting these techniques, exclusively designed to detect an attacker's presence.

The challenge is that the security solutions designed to alert you to suspicious and unwanted behaviour do exactly that: they alert you - a lot. In our own security operations center, we see that an organisation with 2,000 employees will, on average, generate over 400,000 security events every day. It is simply not feasible for any organisation, regardless of how well funded they are, to manually assess this volume of alerts.

This challenge is compounded by the fact that the alerts themselves are often low-value and commonly considered unreliable. According to a study by the Ponemon Institute, 81% of malware alerts are considered unreliable by security professionals. It is really no surprise then that only 4% of these alerts are ever investigated.

> “
>
> An organisation with 2,000 employees will, on average, generate over 400,000 security events every day.

But let's assume that an alert is investigated – then what? These automatically generated alerts provide marginal value as they have limited information, little-to-no context, and often represent an isolated, technical event from an individual point in a network.

These automated alerts serve the function of bringing the existence of a potential security incident to an operator's attention. It is up to a security team to make sense of these alerts, validate the security incident, assess if they represent a threat to the organisation, and take appropriate action to respond to the potential threat.

> Without a plan and capability to respond to a detected threat, you are not much better off than having not known about the threat in the first place.

DETECTION MEANS NOTHING WITHOUT THE APPROPRIATE RESPONSE

The ability to detect malicious and unwanted behaviour in a network only represents half of the battle. Without a plan and capability to respond to a detected threat, you are not much better off than having not known about the threat in the first place.

Responding to a security incident requires that an organisation understands more than just the technical details. An effective response requires that an organisation observes the threat not just as a single malware alarm that has been triggered on some client, but as a security incident that has the potential to (further) impact key business processes if not handled appropriately.

STAGES OF DETECTING AND RESPONDING TO SECURITY INCIDENTS

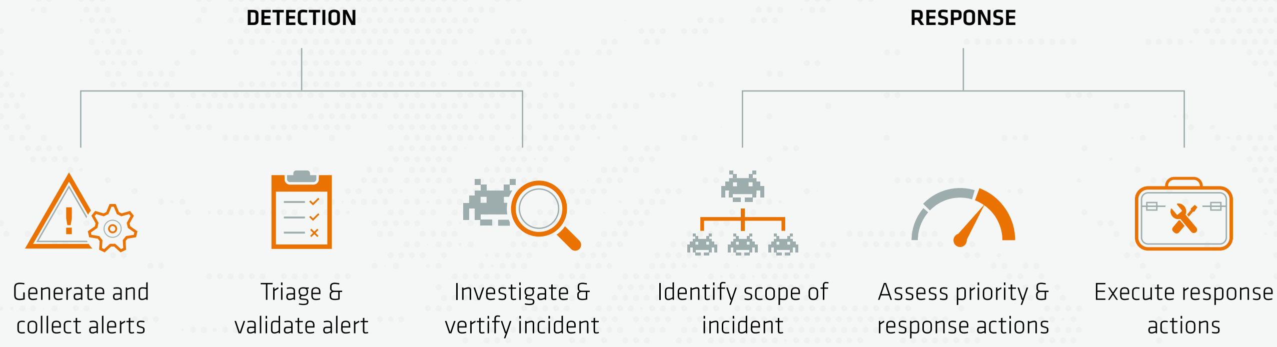

DETECTION

Generate and collect alerts

Triage & validate alert

Investigate & vertify incident

RESPONSE

Identify scope of incident

Assess priority & response actions

Execute response actions

Every organisation will need to go through various stages when detecting and responding to security incidents. These stages are the same, regardless of the threat, industry or technology a company may have.

The detection phase identifies and validates potential threats against your organisation. Much of this phase can be automated and driven by technology, however people will need to be involved at some point to validate the threats.

The response phase is focused on understanding the threat in the context of your business, and taking appropriate actions to remediate the incident. What is the significance of the assets, data and users involved? What services are impacted? How will this affect core business functions? Evaluating the incident in this context enables a response that reflects the severity of the incident based on what it means for your business, rather than solely on the threat itself.

IT'S ALL ABOUT THAT CONTEXT

Context is an important aspect of all decision-making processes. The more information we have surrounding the circumstances of a situation, the more likely we are to make an informed decision on how to proceed. This applies not only to incident response, but every decision we make.

Consider Halloween night as an example. On any other night of the year, if you saw a bloodied person with an axe in their head, you would quickly assume they are severely injured and likely call an ambulance. However the information that it is Halloween night provides context to the situation and will influence how you assess the situation and ultimately, the decision you make. While the decision you make may not change, the extra information adds context and supports a more informed decision to be made.

The same concept applies when responding to security incidents. The more information we understand surrounding the incident, the more informed of a decision we are positioned to make.

Some examples of context that can assist in the decision making process includes:

USER CONTEXT:

Which users are involved? What is their role? Where are they located? What systems and information do they have access to? What influence do they have in the organisation should they be impersonated by an attacker?

INCIDENT CONTEXT:

What systems are affected? Were more than one system involved? Was the threat blocked? Are there alerts from multiple systems? What is the technical scope of the incident?

THREAT CONTEXT:

What type of attack is it? How sophisticated is the attack? Is this a targeted or opportunistic attack? Can the attack be attributed to any individual threat actor? Where is the attack originating from? Have we seen similar attacks in the past? Are other organisations in my industry or region being attacked?

BUSINESS CONTEXT:

Are any of the involved systems or users connected to any critical business processes? What is the potential impact towards these business processes?

How a security incident impacts a business is driven not only by the type of threat, sophistication or attack vector, but the business under attack itself. Capabilities aside, each organisation will have a different set of key business functions and respective priorities for these functions.

For an online retailer, this may mean prioritising that the webstore is available and able to process sales. Meanwhile a law firm may be far less concerned with downtime on their website as they would in protecting their clients' personal and confidential data.

Putting a security incident into the context of a business' core functions enables a response that proportionately reflects the severity of the incident as it relates to how the business is impacted, rather than on the threat itself.

In a world where minutes can be critical,
months are an eternity.

THE IMPORTANCE OF RAPID DETECTION AND RESPONSE

There are an array of regional and global reports that provide insight into how breaches happen, how long attackers go undetected, and how prepared organisations are to respond. One report puts the average time from compromise to detection at around 2.5 months, while another has it at just under 5 months.

The exact figure is less important than the reality of these numbers – we are measuring our detection time in *months.* And bear in mind this is just the time it takes to detect the attacker, not the time it takes to respond. Imagine what you have done in the past 2.5 months, let alone 5 months. In a world where minutes can be critical, months are an eternity.

Despite the statistics in each report varying, the main takeaways are the same:

Takeaway 1: The longer it takes an organisation to detect and respond to a compromise, the more costly it will be.

The direct costs for responding increase as an attacker moves laterally throughout a network and widens the scope of the incident, and increase the chance of direct financial loses as the attacker has more time to execute their actions. Similarly, indirect costs rise as the probability of intellectual property theft (and in all likelihood, the magnitude of this theft) increases the longer an attacker has to operate unimpeded.

Takeaway 2: Security incidents that are discovered by an external party take far longer to respond to than those detected internally.

It makes sense that incidents detected by external parties (e.g. law enforcement, regulatory bodies, customers, etc.) will take longer to respond to simply from an operational perspective to alert the right person in the organisation with the necessary information. External detection is also commonly based on the magnitude or residual effect of a compromise rather than the compromise itself.

However this expected delay does not account for the fact that organisations take 7 to 28 times longer to respond to externally notified incidents than internally notified incidents (internal includes those detected by Managed Security Service Providers). One plausible explanation is that organisations without the capability to detect incidents internally are also less likely to have the plans, processes and routines required to rapidly respond to security incidents in general.

Takeaway 3: The probability for data being stolen rapidly increases when an attacker's presence in an environment moves from days to weeks.

This should not come as any surprise. The longer an attacker has in your network, the more likely they are to succeed in their goals.

Takeaway 4: The longer an attacker is present in a network, the lower an organisation's confidence becomes in their understanding of the full scope of the incident and ability to completely extinguish all effects of the breach.

Time is an invaluable resource for your adversaries. The more time an attacker has in your network, the more complex the incident is likely to become, and the more difficult it will be to understand the full scope of the incident. This leaves an organisation with avoidable gaps in their knowledge, forces them to make less-informed decisions, and a general sense of uncertainty.

Takeaway 5: Technology alone is not enough.

Incident response is a decision-making process. Technology is a tool to support this process, however it is exactly that – a tool. Assessing a security incident, understanding the potential impact it has on a business, and determining how to respond involve a series of complex decisions that only people can make. If you are an online retailer, and you suspect your payment systems have been compromised the day before Black Friday, technology cannot decide whether to take your website offline during the biggest shopping day of the year - only people can make that type of business decision.

OPERATIONALISING INCIDENT RESPONSE

Rapid detection and response operationalises incident response. Handling security incidents become a part of the daily routine and an organisation enters a state of continuous response. Security incidents are no longer seen as an anomaly but an expected occurrence on any given day. From an operational viewpoint, the handling of security incidents is seen no differently than traditional operational activities, such as server maintenance or user management.

The effect of operationalising incident response is that the efficiency of the process dramatically increases. These efficiencies mean that an organisation moves from a purely reactive state of defence to a (near) real-time level of incident response. The resources required to facilitate the response decrease, the elapsed time between initial compromise and recovery shortens, and the process as a whole becomes predictable and measurable.

In the end, this translates to cost savings, an ability to confidently report on the state of cybersecurity and reduces the potential impact a security incident may have on the organisation.

A FINAL NOTE

Cyberattacks will happen. Some of them will be blocked, others will be successful. Some will be caused by employees, and others will elude even the most fortified defences. While it is not possible to prevent all attacks, it is not realistic or feasible to completely abandon prevention and rely solely on detection and response.

It is not a choice between either prevention or detection and response, but how to best use them in combination.