Whitepaper

The Forgotten Factor: Timing Incident Remediation

WITH secure

Contents

The Forgotten Factor: Timing Incident Remediation 3
Day-to-Day vs Major Incidents..... 4
Risks of Poorly Timed and Executed Remediation..... 5
Introduction to Stages of Remediation 7
The Theoretical Knowledge: Points of Consideration 9
Practical Implementation: The Flow 16
Conclusion 18

The Forgotten Factor: Timing Incident Remediation

Every day, security teams must decide whether to remediate an incident immediately or pursue further investigation. Our research indicates that the concept of the "Containment Striking Zone" was initially introduced by Mandiant in their 2010 M-Trends report. This idea was further expanded upon and updated in the 2012 third edition of Incident Response & Computer Forensics by Jason T. Luttgens, Matthew Pepe, and Kevin Mandia. However, this information was targeted for nation state actors, and very little has been published about the strategy of remediating incidents. Moreover, these publications precede many of the major technological advancements in the cybersecurity field, including the widespread acceptance and use of endpoint, identity, and cloud detection and response technologies. All and all, the timing of remediation activities remains one of the most undocumented aspects of incidents.

As seasoned incident response professionals, we almost have an internal clock that influences when we act. However, this skill is rare and difficult to teach to new responders. In this whitepaper, we aim to clarify the concept of the “Containment Strike Zone.” We also provide updated guidance on remediation timing, including insights into the key factors that seasoned incident responders consider in the current landscape when deciding on remediation timing.

This paper does not dive into remediation planning and investigations, but rather focuses on a particular area of incident management that is timing the remediation.

Day-to-Day vs Major Incidents

A vital component of choosing a remediation strategy hinges on your ability to identify major incidents. Many of the considerations we will discuss are pertinent to major incidents, but applying them to routine, day-to-day incidents is inefficient and could result in adverse outcomes.

An incident is deemed to be a major incident if:

- the number of affected endpoints is greater than 1, 3, or 5 (based on the maturity of security operations)
- critical assets are compromised (“Crown Jewels”, i.e., Domain Controller, exchange server (PII))
- the investigation requires full disk forensics
- the infection vector originates from an endpoint or identity about which the organization has no live visibility
- on-site assistance is necessary
- full coordinated incident management is needed, including if:
  - incident resolution will take longer than five hours
  - the incident is too big for a single analyst to manage
  - Remediation requires coordination with IT.

The number of affected endpoints highlighted in the first criteria depends on maturity of your security operations. A reputable managed detection and response team would typically be proficient in remediating an incident through EDR with minimal IT support for up to 5 endpoints. In contrast, for many internal teams, this number is more likely to be 1, and 3 for those with more experience. This decision should be made by security operations leaders. Ask: "How many endpoints can my team handle without requiring my direct intervention?".

Any incident falling outside of these criteria, regardless of the severity, is a day-to-day incident. Some examples are:

- A VIP email account was compromised. The incident is rated high severity due to the potential impact. An investigation is on-going, but the containment action was simply resetting the user account password.
- A user endpoint was infected with malware. It is unclear what the malware has done and investigation is required. The device was isolated and user passwords were reset by the security team.
- A public-facing application was compromised. The threat actor succeeded in lateral movement, but all actions taken have been timelined within a 5-hour timeframe. IT is currently deploying a patch, the malicious files have been removed, malicious IP addresses have been blocked, and credentials have been reset.

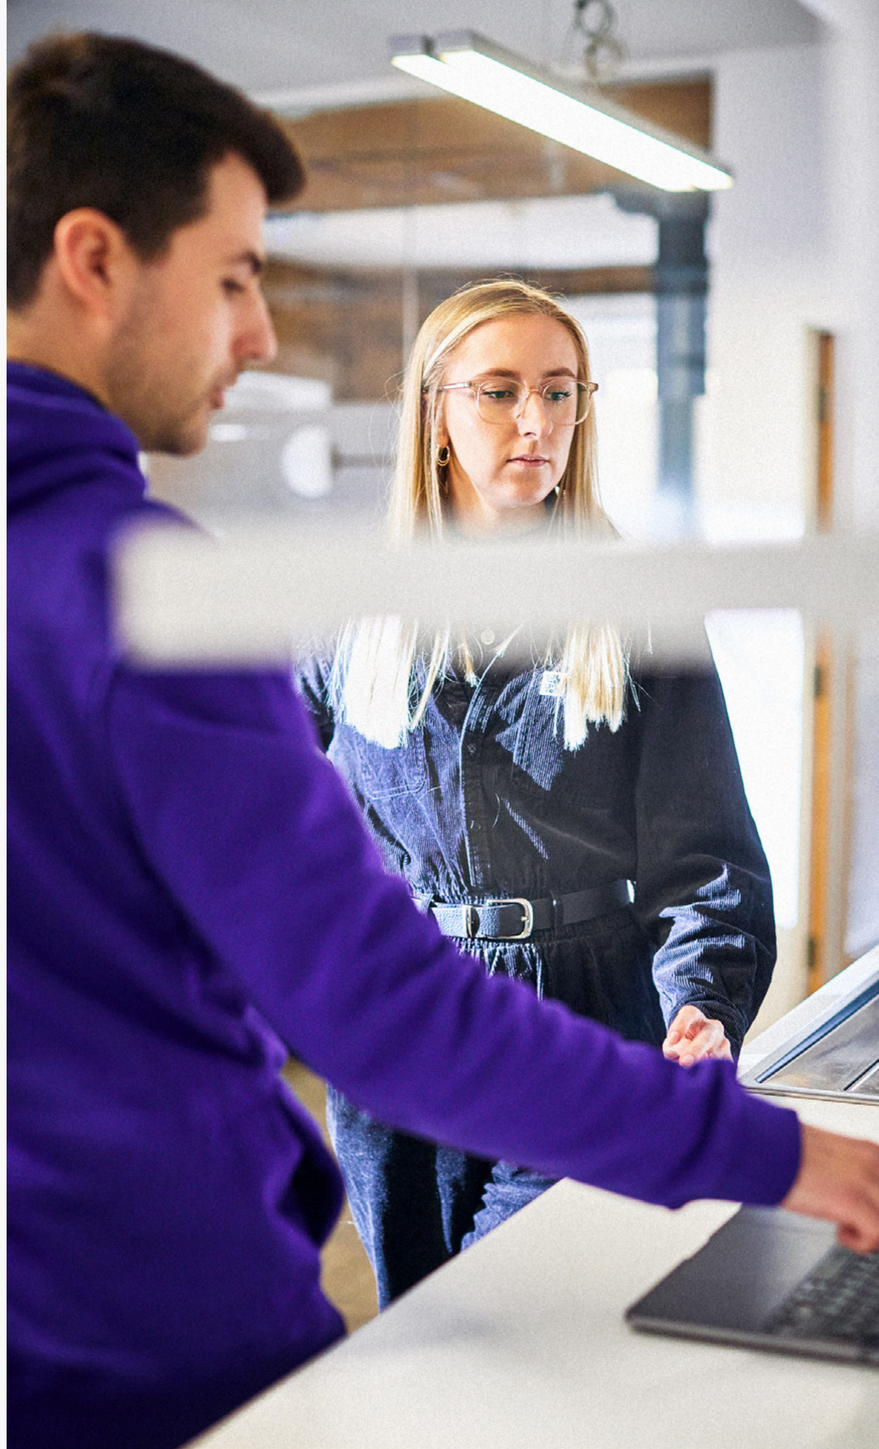

Risks of Poorly Timed and Executed Remediation

The greatest concern regarding poorly executed and timed remediation is inadvertently alerting the threat actor. Historically, it's been agreed that this could lead to five potential outcomes. In this section, we cover these outcomes/risks and update their relevance based on the current landscape.

Although these risks may still be pertinent in a unique, small subset of scenarios, they often do not apply to most routine cyber operations. Furthermore, while these risks are still valid, they can be substantially minimised with live monitoring and detection on your environment; i.e., endpoint or identity detection and response tooling.

Opportunity to Observe the Threat Actor

The present legal and regulatory landscape largely prevents organizations from postponing remediation efforts in favour of gathering intelligence on the threat actor. Adopting such a strategy could lead to significant backlash from regulators and customers alike. Additionally, with the surge in threat intelligence and evidence available to responders, the need for this intelligence-gathering activity has decreased significantly. With numerous sources available, organizations can usually understand a threat actor's objectives without allowing them to advance in their attack trajectory.

Threat Actor Becoming Destructive

The risk of threat actors becoming destructive following an initial compromise was previously rare unless provoked with a failed mis-scope remediation. Threat actors preferred going undetected and unnoticed, aiming to cause as little operational impact as possible. However, because of the emergence of ransomware as a predominant cyber threat and an increasing number of these actors working systematically to maximise profits through causing impact, the risk created by delaying remediation is increased. Giving threat actors a few more days in the environment while you complete the investigation is substantially riskier now than it was historically.

As a result of advancements in the detection and monitoring sector, many organizations can now implement best-effort strategies to mitigate impacts and confidently lean on live monitoring to handle any lingering access as the investigation continues.

Threat Actor Becoming Dormant

Over the years, many vendors have highlighted the risk of threat actors going dormant upon detection or remediation attempts. However, there's a significant lack of evidence to suggest that most threat actors behave this way.

The WithSecure™ Incident Response team has observed a few unique cases where politically motivated threat actors, such as those backed by nation-states, went dormant after the deployment of EDR or during best-effort containment. These observations make up less than 1% of our engagements.

Maintaining comprehensive visibility across your estate is a sufficient countermeasure against this risk.

Threat Actor Changing TTPs

Although politically motivated threat actors have been noted to change TTPs as their campaign evolved[1], our research has concluded there have been less than a handful of occasions where a threat actor has changed all their TTPs within a short period in a single compromise.[2] This requires significant skill and resources.

In contrast to nation state attacks, most attacks are performed by financially motivated groups who aim to expedite their attack cycles each year. Their goal is to quickly achieve their objectives across the many targets they access during extensive campaigns. This strategy often leads them to employ a larger number of operators, typically with a lower skill set. These operators rely on a series of playbooks to execute attacks, which inherently restricts their ability to change TTPs rapidly.

Additionally, each year we continue to observe more and more threat actors performing attacks solely utilising live-off-the-land binaries, open-source tools, and valid credentials. Based on the 2023 Global Threat Report by CrowdStrike, malware-free attacks now account for 71%.[3] It's fair to say threat actors have learned this is the most cost-effective way to avoid detection following initial access. But this in turn have hampered their ability to switch TTPs as their toolboxes have become more limited and publicly well documented through projects like LOLBAS.[4]

Modern detection and response tooling also focuses on a wide range of TTPs, rather than just indicators of compromise, so the chance that new TTPs offer better stealth have been seriously reduced. This means that the risk of threat actors changing TTPs is significantly reduced if you have good estate visibility.

Threat Actor Attempting to Overwhelm or Distract the Organization

As threat actors became more structured and organised, we have increasingly seen that they follow a certain optimised pattern to achieve their goals. This trend led to fewer actors employing tailored attack strategies, such as using distractions or attempting to overwhelm an organization to obscure their true intentions.

Also, while this may have worked against a traditional forensics-led response, this does not work against organizations with live monitoring capabilities. Most organizations today have scalable cyber operations/tools and can maintain better overall situational awareness.

1. DeadRinger: Exposing Chinese Threat Actors Targeting Major Telcos (cybereason.com)
2. Diving Deep into UNC4841 Operations Following Barracuda ESG Zero-Day Remediation (CVE-2023-2868) | Mandiant
3. 2023 Global Threat Report | CrowdStrike
4. LOLBAS (lolbas-project.github.io)

Introduction to Stages of Remediation

The terms containment, eradication, recovery, and remediation are frequently found in incident response literature. However, we found that these terms often lack clear definitions and are not even listed in the glossaries of well-known frameworks like NIST or organizations such as SANS. Furthermore, there were notable inconsistencies in the explanations provided across different resources. To mitigate any confusion, we will offer concise definitions for these terms in this section.

Containment

Containment refers to the actions taken to prevent the spread of a security breach or incident within an organization's network. The aim is to limit the damage by isolating affected systems and networks to prevent further compromise. This may include isolating affected resources (credentials and systems), ensuring your backups are safe, and taking further backups of data if necessary.

Eradication

This term refers to the process of eliminating all changes made by the threat actor, and potentially the root cause of the security incident, from the network or system. The goal is to remove the threat so it no longer poses a risk.

Recovery

This phase involves restoring affected systems and processes back to their normal operation. It might involve rebuilding systems, restoring data from backups, and reinstalling applications. This step is critical to returning to “business as usual” and ensuring that systems are hardened against similar threats in the future. This stage may not always be required if the incident is dealt with in a timely manner.

Remediation

This term refers to steps taken to manage, contain, and mitigate the damage caused by a cybersecurity incident or breach. Put simply, this term encompasses three critical stages: containment, eradication, and recovery.

Although remediation traditionally consists of three critical stages, the landscape of IT networks, cyber-attacks, industry expertise, and defensive tooling has significantly evolved since these terms were introduced. As a result, incident responders often find themselves capable and required to execute all three stages at the same time.

Approaches to Stages of Remediation

Before delving into the key considerations to determine the right approach, it is worth noting that the timing of the incident remediation typically follows one of three approaches: immediate, best-effort, and timed. While the “immediate” approach is most suitable for day-to-day incidents, the “best-effort” and “timed” approaches are relevant for major incidents.

Regardless of the chosen approach, a remediation plan should be drafted as promptly as possible. For each remediation action, the remediation team should determine how the action will be performed, by whom, and how long it will take so that all remediation actions can be executed in coordination. Reducing the execution window of the remediation ensures the threat actor has less time to react if they identify you are attempting to remove their access. This process is often referred to as posturing.

The three approaches to timing remediation can be applied holistically or incrementally across all three phases. For example, you might initiate immediate remediation, discover the incident is still ongoing, and then execute a series of best-effort containment measures followed by a timed remediation. Alternatively, you might opt for best-effort containment and eradication first, finishing with a timed recovery. The chosen approach will depend on the specific circumstances and the desired outcomes.

Immediate
This strategy is usually employed for day-to-day incidents when the attack is detected early at the initial access stage. An immediate response could involve disabling compromised user accounts, blocking malicious IP addresses, or isolating affected systems from the network.

An appropriate case for this approach is a security analyst detecting a user has been successfully phished and malicious attachment had executed. However, so far, no hands-on threat actor activity has been observed or activities have been limited to reconnaissance. This approach may have a small impact on business operations at team or individual level and can be carried out by the security team without coordination between multiple teams.

Best-effort
This approach is typically adopted when inaction could lead to severe consequences for the organization, such as a total network loss due to ransomware or allowing threat actors to exfiltrate data. This strategy might require multiple waves of actions, each with distinct objectives. The initial wave could be aimed at halting data exfiltration or restricting the threat actor's access. It could also be an attempt to eliminate the threat actor while maintaining awareness and vigilance for their potential return.

Best-effort remediation might be restricted to a specific segment of your network or directed towards your entire network, with remediation actions that could exceed the immediate scope of the incident. It often encompasses the entirety of the network; i.e., resetting all user passwords, limiting internet connectivity, removing known malicious files and persistence, or a full domain take-back exercise.

Remedial action taken in the first 12 hours of an incident is generally regarded as a best-effort action. The duration to execute a best-effort remediation can vary based on the scope of actions. It could range from mere minutes for simpler tasks like isolating a host to prevent data exfiltration, to longer durations such as 12 hours for more complex actions like a domain take-back exercise.

Timed
This approach, also known as “delayed action,” is adopted if the response team needs time to prepare. Reparations might involve determining the potential impact of the recommended changes, coordinating actions with individuals responsible for making network adjustments, or simply managing a global workforce. This approach might also be chosen if crucial investigative paths, essential for a comprehensive remediation plan, are nearing completion. For instance, this could involve determining the threat actor's initial access point, means of lateral movement, or persistence. With the right visibility and a seasoned incident response team, you can execute timed remediation within 72 hours of the detection of the incident.

The Theoretical Knowledge: Points of Consideration

At high level, we have found most incident responders' remediation timing strategies are based on four considerations. They are:

1. Estate visibility: your reaction time to threat actors' actions.
2. Scope of compromise: likelihood of successful containment and eradication.
3. Threat actor progress to action on objective: critical triggers against impact.
4. Threat actor objective impact: determines remediation strategies.

These are listed in order of importance. The ranking is based on the critical dependencies of the consideration on each other as shown in the diagram below.

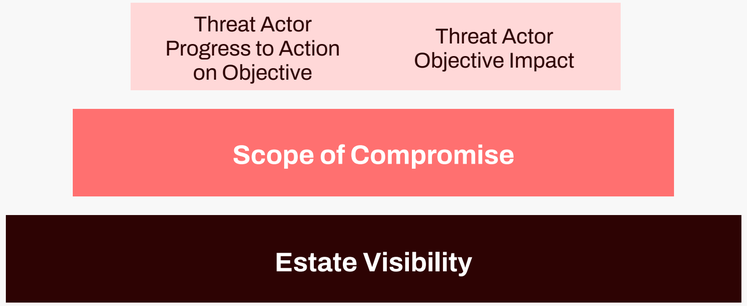

Although we hoped to provide more concreate relationships between these considerations, they are absolutely intertwined and further insight in one area may lead conclusions in another. These considerations are the theoretical knowledge required to follow the practical process flow diagram shared at the end of this whitepaper.

Estate visibility

This can be defined as an organization's telemetry coverage. Estate visibility is the most critical item when it comes to remediating a threat actor because it can enable you to:

- Establish the scope of compromise faster and more accurately.
- Respond quickly to hands-on keyboard attackers who may propagate on your network in an unpredicted manner.
- Confidently monitor your environment if the threat actor attempts to use any remnant access.

Most organizations' estate visibility capabilities can be categorised into three types:

- **Bare**: No central SIEM/logging. Logging is stored on each device. Reliance on Anti-Virus dashboard.
- **Traditional**: Centralised SIEM or logging, including system, antivirus, firewall, and I(D)PS logs.
- **Baseline**: Centralised SIEM or logging, including process creation, network connection, file/registry creation, persistence creation, system, antivirus, firewall, and I(D)Ps logs.

Estate visibility across different parts of the network should be rated separately. Different remediation strategies may need to be applied based on coverage.

Although day-to-day incidents can be managed with forensics and the limited visibility provided by the **bare** and **traditional** types, these deployments severely risk failing to remediate and manage incidents involving hands-on keyboard threat actors, such as ransomware and nation state actors.

Even seasoned investigation teams relying solely on forensics would lag behind the threat actor by at least four hours. This four hour window encompasses several steps that assumes efficient coordination and execution. First, the investigation team identifies who has access to the asset. Next, they provide the administrator with instructions and tools to gather and send the evidence. Finally, someone from the investigation team processes, analyses, and presents their conclusions to the wider investigation team.

While this is possible for most seasoned teams, on average most investigations may take up to eight hours or even longer.

This window is ample time for the attacker to obtain more credentials and compromise additional hosts. In the 2022 Falcon Overwatch Threat Hunting report, it was noted that threat actors took an average of 1 hour and 24 minutes to move laterally, with about 30% of these actions completed in less than 30 minutes.[5] This aligns with our observations from the WithSecure™ Countercept service, where we've seen remnant access exploited to achieve lateral movement in under 30 minutes.

To remediate an active threat actor in your network, you ideally need to be 5–15 minutes behind at maximum. This is possible if you have a baseline visibility. This number is based on our collective experience as team dealing with active threat actors who may attempt to use remnant access or action on objective while the investigation has just started.

Having better visibility with a good detection capability should also help you to spot incidents early on and collect considerable amounts of information on techniques utilised by the threat actor up to the point of detection. This will allow you to establish the scope of compromise rapidly and enforce enhanced monitoring for the techniques used by the threat actor.

Recommendation

For organizations with bare and traditional estate visibility, we strongly advise establishing baseline estate visibility, at least during major cyber incidents. This may involve:

- implementing Sysmon, central logging and monitoring,
- establishing a retainer contract with an incident response vendor that can deploy with a cloud based EDR or xDR on demand,
- setting up an open source EDR solution centrally to be used by your security team if required.

5. https://www.crowdstrike.com/resources/reports/2022-overwatch-threat-hunting-report/

Scope of Compromise

Your understanding of the scope of compromise can be directly correlated to the success of the remediation effort; missing a single host or a credential may mean that the threat actor can re-compromise your network.

The following stages can be used to describe the understanding of the scope of compromise:

Beginning
a. Alert is a true positive.
b. Based on the initial findings, including knowledge of whether the threat actor is using specific malware families, as well as the initial set of credentials used. This also includes other assets that maybe involved due to specific deployment/application; i.e., other assets the AWS credential has access to, web servers linked to the SQL server from which an SQL injection was generated.

Estimated
a. Includes beginning.
b. Critical TTPs employed by the threat actor have been identified. This includes the valid credentials, lateral movement techniques, and persistence mechanisms that have been used by the threat actor.
c. No new TTPs have been discovered in recently identified hosts.
d. An experienced investigation team should be able to reach this stage within 72 hours maximum.

Complete
a. Includes estimated.
b. All identities and endpoints have been triaged OR a complete remediation plan is at hand, including specific actions covering valid credentials, root cause, malware, and network remediation actions.

In all investigations, the investigation team should strive for complete understanding of the scope of compromise. However, based on the pre-existing estate visibility, this may take several days, and the investigation team may not have time as the threat actor will have moved on to the actions on objectives stage.

If you have baseline levels of estate visibility, reaching an estimated stage can be trivial and it is generally a good stage to execute remediation. It is important to keep communicating with the key stakeholders as you approach this stage so they can perform the drafted remediation actions and start posturing at short notice if required.

While this paper does not go into remediation planning, it is worth noting that the worse your understanding of the scope of compromise, the more drastic remediation actions may be recommended or needed. Depending on the estate size and the nature of the incident, this may mean operational impact. If your understanding of the scope of compromise is at the beginning state, you may need to isolate a wider range of hosts and reset all the passwords in your network. This is in comparison to resetting credentials for a small set of credentials or simply removing malware and its associated persistence. **If you had to perform a best-effort containment, it is highly advisable to actively monitor the estate at a baseline level to account for the risk that the threat actor may have remnant access.**

Root cause

Root cause is generally in the forefront of everyone's mind when an incident occurs. We often explain that the root cause of an incident can be remediated before the incident itself is remediated. Unlike a backdoor which may be beaconing back to the command and control server, most initial footholds can involve exploitation or a user action that is not trivial for threat actors to monitor. Equally, a best-effort remediation can be performed before the root cause is addressed if you have established active monitoring of the estate at baseline level. Ideally, both will happen around the same time, but this is not a must have.

Identifying the root cause of a compromise may not be always possible. Evidence may be destroyed or expired, or the incident may have started years ago. We can divide root causes in four categories:

- **Unknown**: you do not know the root cause.
- **Known**: you have identified the root cause but cannot remediate it due to business requirements, or you are aiming for timed remediation.
- **Remediated**: the root cause has been identified and remediated.
- **Timeworn**: we have a hypothesis about the root cause, but we do not have the evidence needed to prove it.

Ideally, the root cause will be either known or remediated, but often organizations that do not have the baseline estate visibility may find themselves in the timeworn or unknown state.

Threat actor progress to actions on objectives: critical trigger points

Although determining the scope of compromise is an extensive task, often a few indicators are enough to determine how far a threat actor got in achieving their objective. This is crucial because it establishes few critical points where you may need to attempt to contain or hamper the threat actor's progress to reduce the impact of the incident. Furthermore, if you believe your scope of compromise is complete, you do not need to consider the threat actor progression: you should attempt containment immediately.

Several frameworks provide insight into the intrusion phases of an incident, including Lockheed Martin's *Cyber Kill Chain*® and MITRE's ATT&CK™ frameworks. The *Unified Kill Chain* (UKC)[6] aligns perfectly with the perspective of containment. According to the most recent release from UKC, an intrusion comprises three phases: in, through, and out. Although we agree with all the other updates made, for intrusion stages we prefer to use terminology from the previous version (pre v1.2), which defines these stages as initial foothold, network propagation, and action on objectives. We favour this terminology because it is more accessible to non-native English speakers, it aligns better with established industry terminology, and it is self-explanatory.

6. https://www.unifiedkillchain.com/assets/The-Unified-Kill-Chain.pdf

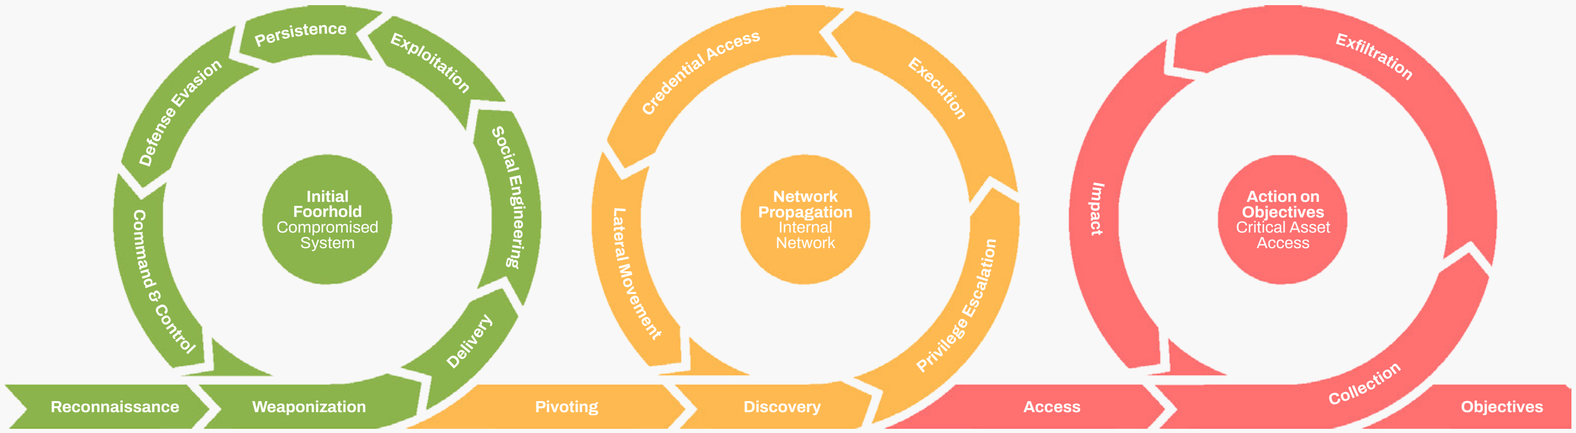

The UKC provides a good definition for all three stages, however for the purpose of this paper these stages can be simply explained as:

- **Initial Foothold**: Threat actor has gained access to a single credential or asset and their access is persistent. If you think that the threat actor's access is limited to a single host, the objective should be to contain the threat actor before they propagate further into the network.
- **Network Propagation**: Threat actor has used their initial access and has propagated their access to further assets or credentials. At this stage, the investigation team should attempt to increase estate visibility to baseline level and identify the complete scope of compromise as quickly as possible, which would trigger remediation.
- **Action on Objective**: Threat actor has gained sufficient access to execute their desired objective. The investigation team should quickly attempt a best-effort containment, accept further best-effort containment may be required, and aim to perform a timed remediation later. The objective is to slow down the threat actor as much as possible. You should also be actively monitoring the environment while completing the investigation.

Although you can break down threat actor progress to finer stages, from the perspective of incident remediation the additional value that would be gained from this would be minimal versus the effort invested.

Guidance on Determining Threat Actor Progress

Based on estate visibility levels and threat actor objective (if determined), organizations can use different indicators to identify the threat actor's progress towards their objective. Some of these will provide critical trigger points where the drafted remediation plan may need to be executed as soon as possible in a best-effort manner to avoid further impact to your business. These indicators are:

1. **Credential Access**: Some of the most crucial and easily available data during incidents are login/log out or credential audit trails, which may be VPN logs or Security/audit logs. Even a simple antivirus detection may give you a hint about what credential the threat actor is using and thus the level of access gained. This is a great indicator to use particularly for financially motivated actors and on-premises networks compared to SaaS or Cloud environments. There are, generally, three categories of credentials:

   a. Standard user: an account issued to a standard user with no administrative privileges.

   b. Administrator: an account with administrative privileges on an asset or multiple assets in the network. This account could also be a local account or a service account.

   c. Domain/enterprise administrator: an account with control over the network's authentication mechanisms.

2. **Data Access**: Politically motivated threat actors do not necessarily require domain or enterprise-level privileges to fulfil their objectives. Often, they seek specific data, which can potentially be accessed with a limited set of credentials. It can be helpful to visualise the threat actor's progress based on their data type objectives. Some examples of this are:

   a. Targeted Data: The threat actor would be considered in the network propagation stage if they have been found to be targeting certain strings; i.e., relating to specific Personal Identifying Information or intellectual property, but the systems containing the information are not compromised. Equally, if these systems are compromised, it suggests that the actor might be on the brink of, or already in the process of, the action on objectives stage.

   b. Recovery Data: If your identified threat actor has attempted to access or achieved access to your back-ups or operational resilience systems, this can also help you differentiate between network propagation and action on objective.

   c. Authentication Data: If the threat actor has administrative control over the authentication services, it would indicate that the threat actor has required privileges to action their objectives.

Threat actor motivation: containment strategy

Containing a politically motivated threat actor is different to containing a financially motivated one. Knowing the type of threat actor you are dealing with will help you determine key indicators of the threat actor entering the action on objective state and should inform your overall remediation strategy. Threat actors and the related containment strategies can be divided into three simple categories:

- **Politically motivated**: continuous threat
- **Financially motivated**: immediate threat
- **Emotionally motivated**: unpredictable threat

Determining the objectives of a threat actor can be challenging. For organizations with limited investigation capabilities that are not part of critical infrastructure, we recommend treating incidents as financially motivated until proven otherwise.

However, if your organization operates in a sector that is listed as part of the critical infrastructure and does not have a strong incident response team, we highly recommend contracting with an accredited cyber security company to support you during cyber incidents.

Politically motivated

Politically motivated threat actors are resourced to target specific organizations, so they are likely to attack repeatedly until they are successful. Once you are a target, a breach is only a matter of time. Political agendas of countries change approximately every five to ten years, so you should expect the threat actor to be interested in you for an extended period.

If you are targeted by a politically motivated threat actor and do not have baseline estate visibility, you should upgrade visibility immediately. You should also contact an accredited cyber security company to support you.

When dealing with political threat actors, dwelling time is also a factor when considering remediation timing. If they have been in your network for over three months, you should weigh the impact of giving the investigation team more time to build a comprehensive remediation plan. If the threat actor is idle, giving the investigation team few more days to perform comprehensive timed remediation may not have any further impact.

With politically motivated threat actors, further controls may need to be implemented off the back of the incident in crucial parts of your network. This will ensure that threat actors are faced with continuously updated impediments that will be difficult to bypass and provide you opportunities for detection.

Financially motivated

Threat actors motivated by financial gain aim to move as rapidly as possible, which means you or your security provider may need to move quickly to contain them. These actors are often easy to link to commodity campaigns published online and simple OSINT based on few indicators of compromise.

You should aim to contain a financially motivated threat actor as you reach an estimated scope of compromise state, or immediately if you think they are about to establish actions on objectives. If you have baseline estate visibility, a good benchmark to follow is reaching an estimated scope of compromise state in a maximum of 72 hours.

As most financially motivated threat actors move to the actions on objectives stage during non-business hours, leaving an uncontained threat actor in your network over the weekend may lead to your business being taken down. If you are facing a financially motivated attacker who is ready to carry out their action on objective and your understanding of the compromise is in the beginning phase, we highly recommend reaching out to an accredited cyber security company. This is especially true if the investigation or monitoring of the situation may be limited due to staff absence, for example not working over the weekend. They can assess the risk and offer urgent advice before the weekend.

Emotionally motivated

Emotionally motivated threat actors, such as hacktivists and insiders, are harder to predict. You may be dealing with an insider who is selling information to your competitor, or an upset administrator whose sole purpose is to embarrass their colleagues.

With this type of threat actor, depending on if you wish to follow legal proceeding or not, you may need to take different approaches. We generally recommend that you aim to contain the threat actor either after gaining complete understanding of scope of compromise, or when you believe the threat actor is about to perform an action that would have unprecedented impact for your business.

Unless you suspect that the threat actor is an insider, you should assume they are financially motivated.

Practical Implementation: The Flow

Although it is impossible to cover every scenario, baseline estate visibility provides opportunities for your organization to quickly establish an understanding of the scope of compromise, profile the threat actor objective impact, track threat actor progress to actions on objectives, and finally provide assurance that containment was successful. But how should you navigate from one consideration to another as the incident is developing?

NOTE: It is essential that you have reviewed the theoretical background provided above to fully understand this process and make informed decisions at each stage. Consider also that this flow relates to decision making with regards to the timing of remediation, not remediation planning, overall incident management, or investigation. It is a subset of the decision making involved when managing incidents.

As with all incidents, the process begins with validation and triage. Assuming it is a legitimate incident, a remediation plan will likely be drafted. Here is where the theoretical knowledge we have discussed comes into play. Does this incident meet the criteria for a major incident?

If it does not, you should perform immediate remediation to eliminate the risk.

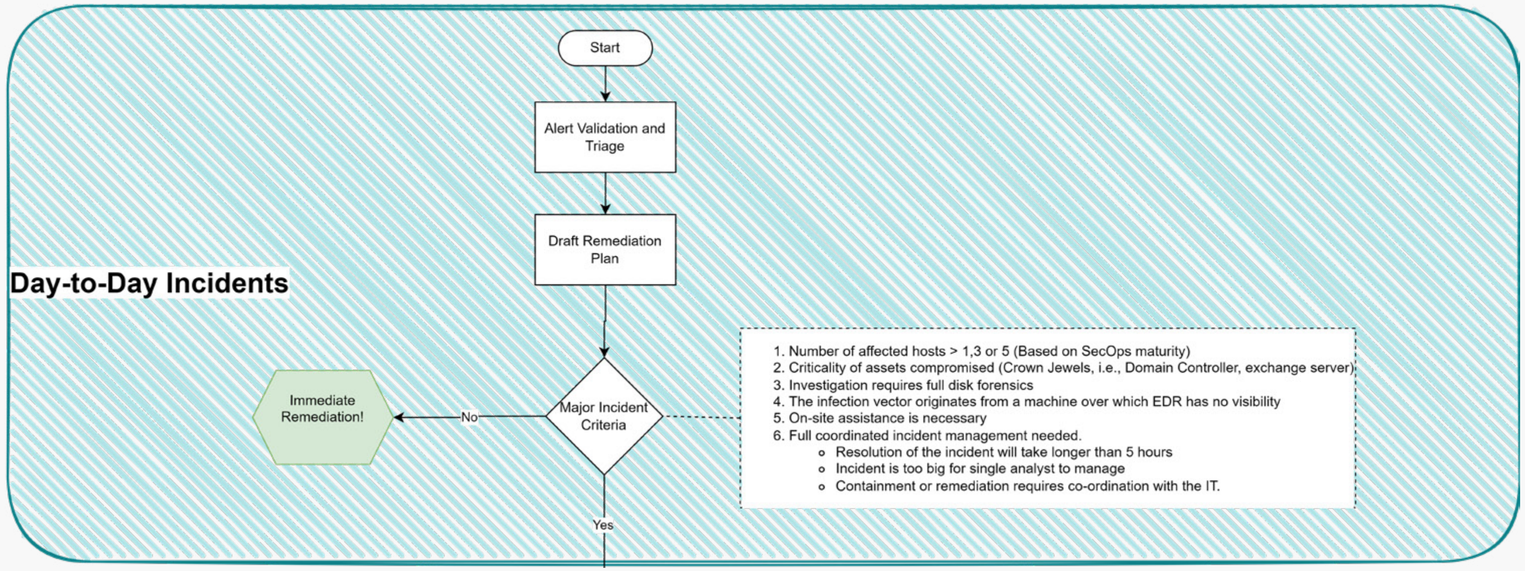

If we have hit the major incident criteria, you have likely started the process of incident management, engaged the relevant teams to swiftly investigate and remediate the risk.

Within the first incident management call, the key question to be answered is whether the threat actor is believed to be at or near the actions on objectives stage. Consequently, are they position to exfiltrate data or cause significant impact to your network? If the answer is yes, you will want to take a best-effort action with the awareness the threat actor may not be fully removed from your network. As mentioned, if you are attempting remediation within first 12 hours, you may have to keep the scope of remediation actions wider than anticipated.

If the scope of the compromise is at the estimated stage, you might be able to remediate the attack now. However, if it has not yet reached that stage, be prepared to take additional best-effort measures.

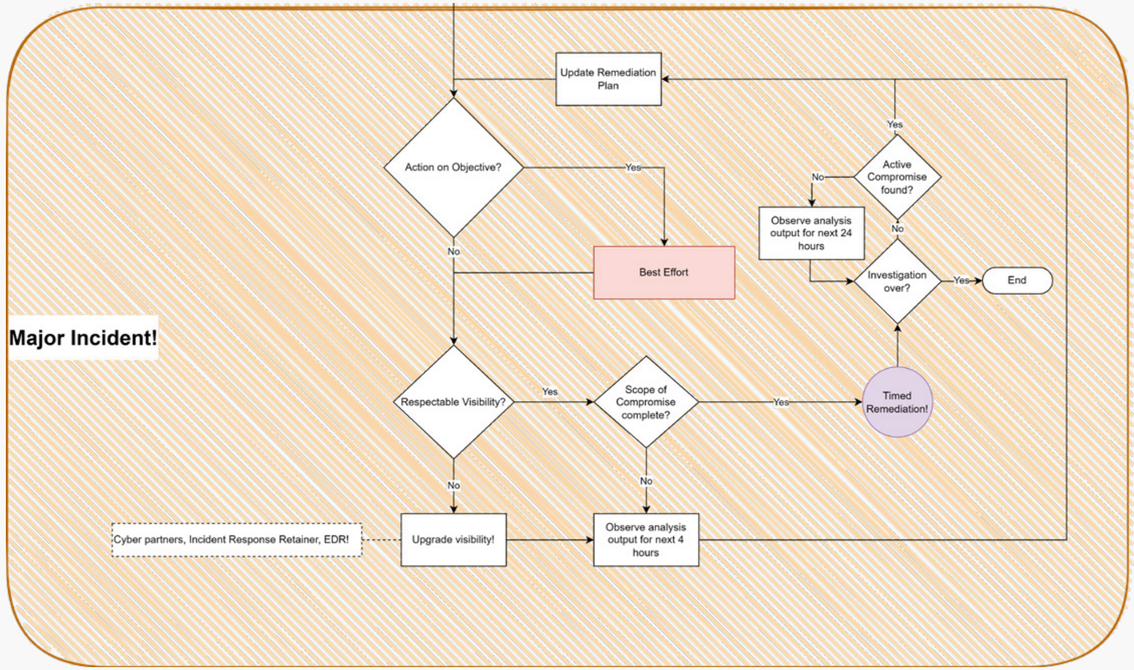

If you have determined there is no immediate threat, you might have chosen not to take best-effort action. However, the next priority should be enhancing estate visibility. Do you have real-time monitoring of the attack? This is crucial to boost your visibility to balance the risk of impact against the need for ongoing investigation. Today, numerous vendors offer cost-effective solutions for monitoring or under-attack scenarios, even covering your cloud and SaaS platforms. Given the current landscape of threats and threat actors moving from initial access to objectives in an efficient manner, attempting to contain live threat actors with forensics is simply not viable.

If you have live monitoring over the incident and you believe you have completed your scope of compromise, congratulations! You have reached the ideal scenario that is possible for matured security operations. You can now confidently move towards timed remediation at the earliest opportunity and continue your investigation if necessary. Should the investigation reveal further risks or remnant access, follow the flow to determine the best strategy for addressing those risks. If the scope of compromise is not complete, it is advised to re-evaluate the situation within the next four hours, update the remediation plan accordingly and go through the flow again.

Major incidents do not take place every day and although they may have a lot of commonalities, organizational differences may apply. However, this flow should be able to guide you through 90% of the major incidents take place. Major incidents are coordination and communication challenges; it is important to be aware of the risks raised by the people involved and adjust your approach where appropriate for your organization.

Conclusion

When we started this work, we hoped to create a calculator that gave you a score, indicating what remediation approach should be taking and the level or risk. We wanted something that tells you are X% into the “Containment Strike” zone. However, as we dove into the subject on timing remediation, we found how little has been documented and how hard it is to take the human out of the equation.

As we tested our prototype calculator, we noticed how key data that would inform the calculator or help our industry progress incident management does not exist, simply because we have not been tracking this regardless of the incident management platform. The timing of the containment has truly been forgotten by both our products and our industry.

We may have failed to create a “Containment Strike” zone calculator, but we decided that there is still huge value to start a conversation on this subject, challenging ourselves and the industry to be better.

In this paper, we have shared our knowledge, experience, and best practices as responders. We hope that you find the content both informative and thought-provoking. We openly invite readers to challenge our views, enlighten us on any oversights or inaccuracies, and contribute to the broader industry discussion on how we can enhance our practices.

Who We Are

WithSecure™, formerly F-Secure Business, is cyber security's reliable partner. IT service providers, MSSPs and businesses – along with the largest financial institutions, manufacturers, and thousands of the world's most advanced communications and technology providers – trust us for outcome-based cyber security that protects and enables their operations. Our AI-driven protection secures endpoints and cloud collaboration, and our intelligent detection and response are powered by experts who identify business risks by proactively hunting for threats and confronting live attacks. Our consultants partner with enterprises and tech challengers to build resilience through evidence-based security advice. With more than 30 years of experience in building technology that meets business objectives, we've built our portfolio to grow with our partners through flexible commercial models.

WithSecure™ Corporation was founded in 1988, and is listed on NASDAQ OMX Helsinki Ltd.

WITH secure