Remediating Targeted-threat Intrusions

Jim Aldridge

Introduction

Successfully remediating a targeted intrusion by a persistent adversary requires a different approach from that applied to non-targeted intrusions. Targeted threats are different for three main reasons:

- Experience has shown that such threats will continue to target victim organizations regardless of the countermeasures instituted.
- They are adept at avoiding detection and take steps to avoid being eradicated from an environment.
- Responders must recognize that the adversary may analyze and react to each action they execute.

To manage the risk posed by these types of threats effectively, organizations must change the way they think about intrusions, understand how targeted threats operate and re-evaluate their security priorities. The approach described in this paper is necessary because traditional remediation approaches often fail when applied to targeted intrusions. Though it has proven successful in the field, organizations must customize the details to fit their particular situations.

Two key principles support this approach. One is that incident responders must recognize that each defensive action may prompt the adversary to react: organizations should delay implementing actions that will directly disrupt the attacker until they are ready to eradicate the threat completely. The second is that while the investigation is proceeding, organizations should prepare for a remediation event specifically designed to contain and eradicate the threat within a short time span. An effective remediation plan should:

- Enhance defenders' visibility to detect indicators of compromise
- Enhance responders' abilities to respond rapidly and effectively to intrusions
- Inhibit attackers' activities post-remediation

Implementing a different approach to remediation (no more "whack-a-mole")

Traditional incident response doctrine focused on each infected system in isolation from others rather than as related pieces of a larger puzzle. Responders identified and then immediately contained infected[1] systems. This approach was appropriate when faced with automated threats. Organizations that follow this approach when faced with a targeted intrusion will immediately block known attacker command-and-control (C2) domains and IP addresses, reset known affected user accounts' passwords and rebuild known compromised systems as responders identify them. If the organization lacks visibility across their environment, which is common in organizations experiencing their first targeted intrusion, this approach has the following effects:

- Responders remove all known compromised systems from the network, leading to a sense of accomplishment.
- The responders "tip their hand" to the attacker. Based on the systems contained, the attacker will know that responders are aware of particular malware variants, utilities and C2 channels.
- Using the backdoors implanted in other compromised systems, of which the responders are not yet aware, the attacker will take steps to ensure continued access to the environment. The

[1] The term "infected" in this context refers to systems on which an attacker has installed malware, e.g. backdoors, sniffers or proxies. Mandiant differentiates "infected" systems, which should generally be rebuilt as part of the remediation process, from "accessed" systems, which may not need to be rebuilt. The term "accessed" refers to systems that an attacker has accessed, but where there is no evidence that the attacker has installed tools. For example, this includes systems from which the attacker has dumped passwords or copied files. "Compromised" systems include "infected" and "accessed" systems.

attacker will avoid using malware, utilities or C2 channels previously discovered by the responders.
- The attacker will continue his work.
- The responders will continue to be blind, and unaware; typically, this lasts until an outside party, e.g. law enforcement, notifies the organization again that they are compromised.

This "whack a mole" approach leads to a cycle of continuously investigating and remediating without ever fully eradicating the attacker from the environment.

Targeted, persistent attackers leverage multiple avenues of access into a compromised environment. Some attackers demonstrate mature operational capabilities, for example using distinct C2 domains and IP addresses for each piece of deployed malware. This helps the attacker ensure that the discovery and removal of one family of malware does not hinder the attacker from accessing the target environment. If the incident responder does not detect and eradicate all the attacker's avenues of access, the attacker will immediately regain access to the environment and continue their campaign.

It is important to point out that the "whack-a-mole" approach may be a reasonable approach in certain circumstances. For example, if criminals are stealing cash from a financial organization in near real time, it is not advisable to spend weeks investigating the compromise prior to taking action. In this case, the victim organization should immediately contain the attack and preserve evidence where possible to assist in the subsequent investigation. The immediate containment approach is appropriate in this situation because the impact of the attackers' continued access to the environment is more damaging than the attacker knowing that a remediation effort is underway. In essence, incident responders confront the attacker and force him divert efforts away from the theft in progress.

Understanding the targeted attack lifecycle

Effective incident response requires an understanding of the attacker's tools, tactics and procedures (TTPs). Targeted threats operate differently from commodity threats: they are sentient, human adversaries. The adversary uses malware and other tools as a means to an end, but it is the adversary, not the tools, that are the actual threat.

Targeted threats typically follow a predictable sequence of events when attacking an organization. Mandiant uses the "targeted attack lifecycle" model to explain this sequence of events. While not all targeted attacks follow the exact flow of this model, it is a useful visual representation for both understanding the flow of attack as designing effective countermeasures.

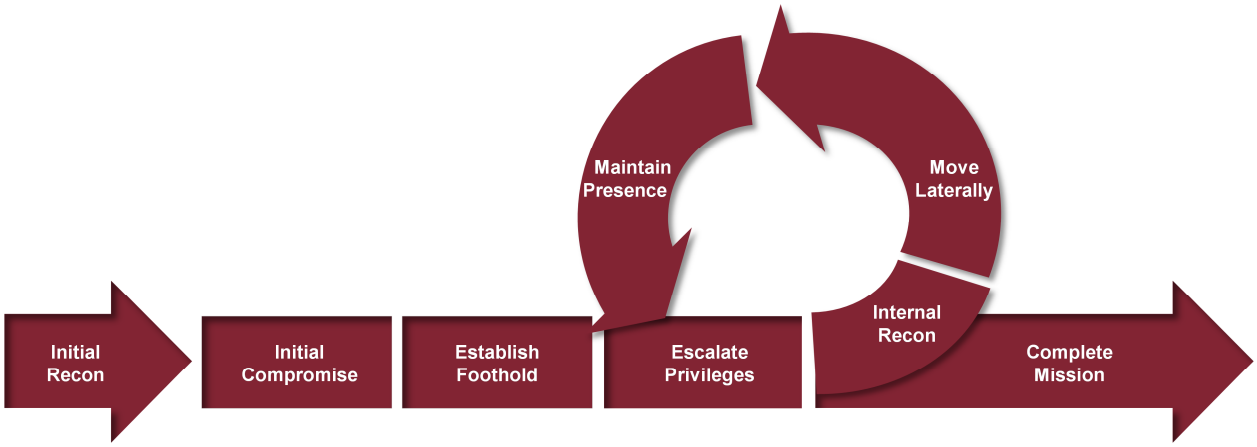

- **Initial Reconnaissance**: The attacker conducts research on a target. The attacker identifies targets (both systems and people) and determines their attack methodology. The attacker

MANDIANT
ADDRESS 2318 Mill Road, Suite 500, Alexandria, VA 22314
PHONE 800.647.7020
WEBSITE www.mandiant.com

may look for Internet-facing services or individuals to exploit. The attacker's research may also involve the following activities:
  - Analyzing the target organization's current or projected business activities, organization and products.
  - Researching conferences attended by employees.
  - Browsing social media sites to more effectively identify and socially-engineer employees.

- **Initial Compromise:** The attacker successfully executes malicious code on one or more systems. Initial compromise often occurs through social engineering (typically spear phishing), by compromising popular web sites used by a target industry (also known as strategic web compromises), or by exploiting a vulnerability on an Internet-facing system.

- **Establish Foothold:** The attacker establishes a means of remote access to a recently compromised system. This occurs immediately following the initial compromise. Typically, the attacker establishes a foothold by installing a persistent backdoor.

- **Escalate Privileges:** The attacker obtains greater access to systems and data than he started with. Privilege escalation is often obtained through password hash dumping (followed by password cracking or pass-the-hash attack), keystroke/credential logging, or by leveraging privileges held by an application.

- **Internal Reconnaissance:** The attacker explores the victim's environment to gain a better understanding of the environment, roles and responsibilities of key individuals and the location of key information.

- **Move Laterally:** The attacker uses his access to move from system to system within the compromised environment. Common lateral movement methods include accessing network shares, using the Windows Task Scheduler to execute programs, using remote access tools such as PsExec, or using remote desktop clients such as RDP, Dameware, or VNC to graphically interact with target systems.

- **Maintain Presence:** The attacker ensures continued access to the victim environment. Common methods of maintaining persistence are to install multiple unrelated backdoors, gaining access to the VPN and legitimate credentials, installing web shells, and implementing backdoor code in legitimate applications.

- **Complete Mission**: The attacker accomplishes his goal, which oftentimes includes stealing intellectual property, financial data or business information (e.g. mergers and acquisition specifics). Once the attacker has completed the mission, he typically maintains access to the environment in case he is directed to complete a new mission in the future.

Caveats

The approach this paper presents makes the following assumptions:

- The victim organization is inexperienced in dealing with targeted intrusions.
- The security team has poor visibility into host and network activities across the environment.
- Prior to the organization becoming aware of the intrusion, the organization does not know whether they have targeted threat activity in their environment.

Organizations with mature incident response capabilities and experience with targeted threats may successfully follow an abbreviated process that differs from the approach outlined here. These organizations typically have the visibility necessary to identify intrusions rapidly, which they effectively and quickly contain. For example, such organizations typically have a 24x7 security operations center (SOC) team dedicated solely to this problem. Organizations that have not experienced this type of attack before do not generally identify these type of intrusions promptly. Consequently, the attacker is

MANDIANT | ADDRESS 2318 Mill Road, Suite 500, Alexandria, VA 22314 | PHONE 800.647.7020 | WEBSITE www.mandiant.com

likely to have been deeply embedded in the environment for years, rather than months or days. This scenario warrants approach outlined in this paper.

Additionally, the approach outlined here is generally not necessary post-remediation. If the investigation has effectively determined the scope of the compromise and the organization has executed the appropriate remediation activities, the day after the remediation event will represent a "clean slate." At this point, organizations should respond immediately to contain any signs of targeted threat activity before the attack can progress along the lifecycle.

The difference between organizations that are successful in the long-term and those that are not is how quickly they can detect and contain new intrusions. Consider the following example. Company A has no idea a targeted attacker has compromised their network. In a matter of days to weeks, the attacker has executed the complete lifecycle, stolen data, and has placed multiple types of backdoors in the environment to maintain a presence until he next needs information. In contrast, Company B identifies the traffic associated with the backdoor dropped in the initial compromise phase within hours. They have the visibility across the environment to identify other systems that the attacker infected and are able to stop the attack prior to the attacker entering the "move laterally" phase. That is a win.

Approach for targeted threat remediation

Organizations that successfully remediate targeted intrusions execute a three-phased remediation plan:

1. Posture for remediation (while investigating the incident to determine its scope).
2. Execute one or more remediation events.
3. Implement strategic changes.

The centerpiece of the plan is a remediation event during which a series of containment, eradication and recovery activities are executed over a short, defined period. The organization takes steps to deny the attacker visibility during this event to prevent the attacker from regrouping. Prior to the event, the organization executes posturing activities necessary to prepare for the event and to enhance the organization's security posture. During the posturing phase, the organization takes care not to disrupt the attacker or alert him to the upcoming event. Activities that cannot be executed before or during the remediation event form the basis for strategic planning.

Although not covered in this paper, it is imperative to determine the scope of the incident prior to executing the remediation event. Investigative activities should be conducted in parallel with remediation posturing. The goals of the investigation are to:

- Determine if an attack is ongoing.
- Confirm the initial method of intrusion and its timing.
- Determine the scope of the compromise, i.e. the affected systems and accounts.
- Understand the attacker's tools, tactics and procedures, e.g. characteristics of their backdoors and command and control infrastructures.
- Determine information exposure.

Mandiant has helped many organizations use this approach to remediate incidents successfully. However, it is important to note that this approach is not necessarily applicable to all incidents. An experienced incident response lead determines the remediation approach required based on a variety of inputs including, but not limited to, the attacker's capabilities, the company's operational readiness and the extent of the compromise and external factors.

It is also worth noting that certain circumstances may make it necessary to contain an attacker from accessing certain systems prior to the remediation event. As in the financial example presented earlier, it may make sense to act when there is a near-certain probability that the attacker is about to steal information with a severe impact to the organization.

MANDIANT | ADDRESS 2318 Mill Road, Suite 500, Alexandria, VA 22314 | PHONE 800.647.7020 | WEBSITE www.mandiant.com

Developing a remediation plan

During an incident, it is in the organization's interest to implement posturing activities quickly because the remediation event should not be executed until these activities are complete. Time and resources are limited. Consequently, organizations should plan to implement the countermeasures that most directly address the attack lifecycle. Avoid implementing strategic initiatives during this phase. When choosing initiatives, it is also important to tailor the plan to fit the organization: one size does not fit all.

Figure 1: Remediation planning matrix provides a graphical reference to the remediation planning process. The targeted attack lifecycle illustrates the different phases an attacker will execute. The matrix provides a guide to help focus remediation planning on the activities that will provide the best value.

De-prioritize any initiatives that do not directly address "detect", "inhibit" or "respond" for one of these phases. For example, do not have an hour-long meeting to discuss whether Windows should be configured to allow only NTLM v2 authentication if all workstations in your Active Directory have the same local administrator password. When time is limited, tangents have an opportunity cost.

The characteristics across the bottom of the figure remind one to consider countermeasures in context. For example, an organization should probably not initiate a project to aggregate NetFlow data from WAN routers to better track attackers' lateral movement if they do not plan to have the right personnel to make use of that information.

Organizations can also use this matrix to help them identify gaps in their security programs.

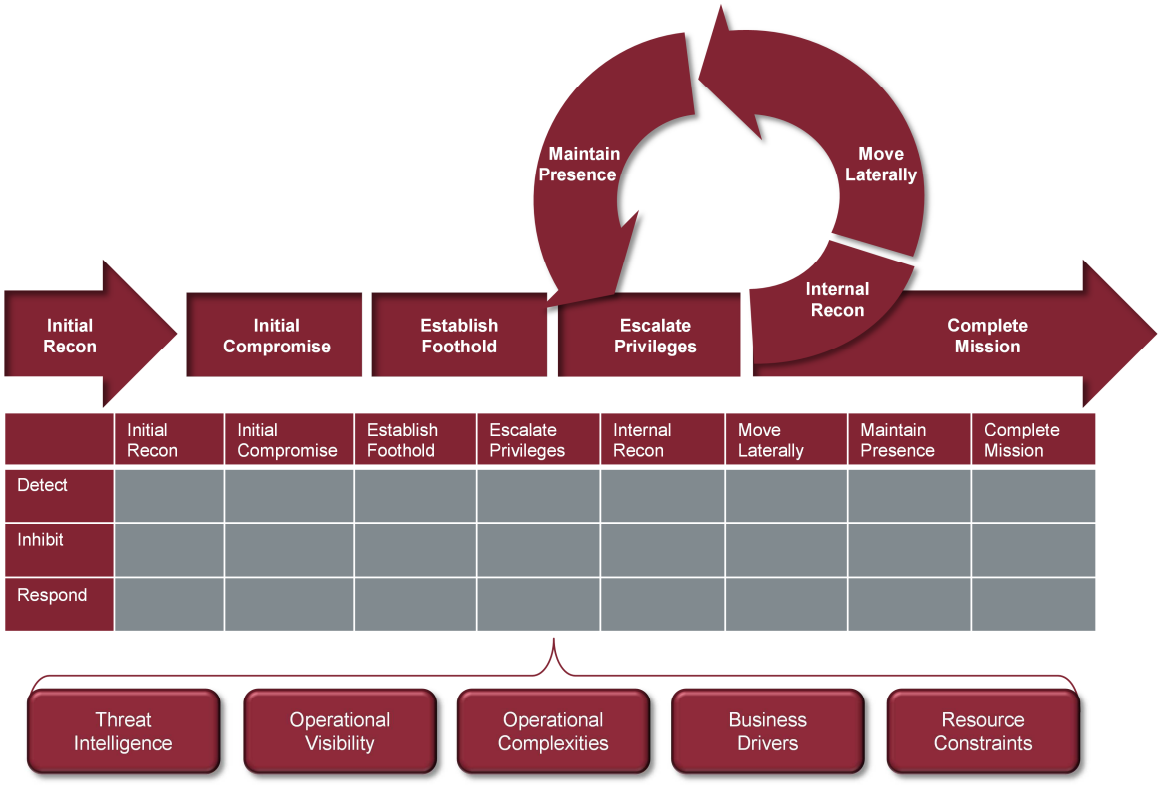

Figure 1: Remediation planning matrix

MANDIANT
ADDRESS 2318 Mill Road, Suite 500, Alexandria, VA 22314
PHONE 800.647.7020
WEBSITE www.mandiant.com

Posturing phase

As soon as a victim organization identifies an incident, they must begin planning the remediation event. During the posturing phase, they must avoid actions that are likely to alert the attacker to the impending remediation event. For example, administrators should not change compromised accounts' passwords, block C2 infrastructure or rebuild compromised systems. During the posturing phase, organizations execute the following key activities, at a minimum.

Plan the remediation workstream
Effective planning is critical to the success of the remediation event. Plan the remediation event in tandem with the incident investigation. At a minimum, these preparatory activities should include the following actions:

- Assign and empower a dedicated remediation lead.
- Brief executive management and obtain support for the remediation effort.
- Create a project plan that anticipates the details associated with planning and execution for each remediation event action.
- Assign accountability and deadlines for the planning tasks.
- Re-prioritize existing security and IT initiatives.
- Track the time spent on remediation planning to capture the level of effort expended on the incident response more effectively.

Preparing for a remediation event is challenging and time-consuming. In a period that typically lasts 4 - 8 weeks, the organization must prepare to execute the remediation event over a short timeframe, typically a weekend.

Enhance logging and monitoring
Increasing the quality and quantity of security event feeds and monitoring that information for indicators of compromise serves two primary purposes at this stage: supporting the investigation and rapidly alerting the organization to attacker activity post-remediation. Consider implementing a tool to centralize these logs and facilitate searches. If such a tool is not already in place, execute the following activities at a minimum:

- Log and retain information regarding DHCP leases. Given the IP address and timestamp, develop a process to determine the hostname. This process should take less than one hour.
- Log and retain logs of DNS queries generated by internal systems for Internet domain names. Given a domain name and timestamp, develop a process to determine the source hostnames that resolved that domain in less than one hour.
- Increase security audit logs on Windows systems to include authentication, system and process tracking events. Domain controller security logs are particularly important. Given a domain username and a timeframe that is within the limits of available data, develop a process to determine all systems to which that account successfully authenticated. This process should take less than 24 hours.
- Log and retain events relating to outbound network traffic at the border firewalls and web proxies. Given a destination IP address or URL, and a time period, develop a process to determine all sources systems that communicated with that site. This process should take less than 24 hours.
- Log and retain VPN authentication logs. Given an account name, develop a process to determine all available attributes related to that account's VPN authentication history in less than 24 hours.

Develop a process to monitor the above logs for indicators of compromise identified through the investigation. Design and implement alerts to detect attacker activity. Prior to the remediation event, this will help the investigation team identify compromised systems and scope the incident. Even if the investigation team has other methods of identifying attacker activity, it is important to have multiple sources of data that provide this information in case one fails. After the remediation event, monitoring these alerts will help to provide prompt notification of a re-compromise.

MANDIANT | ADDRESS 2318 Mill Road, Suite 500, Alexandria, VA 22314 | PHONE 800.647.7020 | WEBSITE www.mandiant.com

Prepare enterprise-wide password change
If the attacker obtained privileged access to Active Directory, it is essential to change all passwords in the environment and not just passwords to known compromised accounts. In many cases, the attacker has obtained all Active Directory user account password hashes including service accounts, and the local Administrator account password(s).

This means an organization must change all passwords, including those for shared, application, and resource accounts. This helps ensure the attacker cannot continue to use an unknown compromised account, should the investigation fail to identify a backdoor. Having changed passwords is also important when the attacker re-compromises the environment, as he may immediately try to move laterally using the passwords harvested from a prior compromise. Though it is difficult to execute, the password change is an important part of making the attacker's mission as difficult as possible.

The password change is typically the most difficult of all the remediation activities to plan and implement successfully. As such, it generally deserves a dedicated project leader and a workstream of its own, which includes the following tasks:

- Inventory service accounts and their dependencies per application.
- Develop a plan for resetting privileged and standard users' Active Directory credentials. This plan should include accounting for users that are traveling.
- Develop a plan to ensure that all Windows local administrator accounts' passwords are unique per-system, or plan to disable the local administrator accounts. Note that password-vaulting software may be required to feasibly implement this recommendation in a large network.
- Develop a plan to reset compromised database, application and non-Windows accounts' passwords.

Focus on the most impactful defensive measures
During the posturing phase, security and IT teams will be overworked in response to the investigation and preparatory remediation activities. Focus only on security initiatives that are critical-path activities for the remediation. Revisit the attack lifecycle and plan countermeasures that will most effectively inhibit attackers' activities. Many organizations execute the following actions during this phase:

- Prepare to implement application whitelisting on servers from which an attacker could harvest password hashes en masse or that provide critical capabilities. These servers typically include Active Directory servers, file servers, software deployment servers and Exchange servers.
- Implement host-based firewall rules on workstations to block workstation-to-workstation network traffic.
- Patch third-party client applications such as Java, Flash, Acrobat and web browsers.
- Reduce the number of elevated-privilege accounts.
- Implement two-factor authentication for remote access.

Remediation Event

During the remediation event, which is typically 24 to 48 hours in duration, the organization removes the attacker from the environment. Depending on the size of the environment and the extent of the compromise, it may be appropriate to execute a series of smaller remediation events rather than one large event. In Mandiant's experience, the most successful organizations executed the following sequence of events:

- Isolate the environment from the Internet.
- Block egress traffic to known malicious C2 IP addresses and domains.
- Block dynamic DNS providers.
- Perform an enterprise-wide password reset.
- Rebuild or replace compromised systems.

MANDIANT
ADDRESS 2318 Mill Road, Suite 500, Alexandria, VA 22314
PHONE 800.647.7020
WEBSITE www.mandiant.com

- Implement technical countermeasures such as securing local administrator accounts, implementing application whitelisting, blocking workstation-to-workstation communication and denying "uncategorized" web traffic.
- Validate that the preceding tasks were implemented correctly. This is extremely important because of the complexity of many IT environments and the number of individuals that may be involved in implementing these countermeasures during the event.
- Reconnect environment to the Internet and validate business functionality once all of the preceding tasks have been validated.

Once the investigation and remediation event are complete, the organization has returned the environment to a known state. The attacker's ability to interact with the environment has been mitigated. Appropriate security alerts have been configured to notify the security team of attacker activity. The organization has implemented countermeasures to both inhibit the attacker and to enhance their detection capabilities.

At this point, the organization should enact immediate containment measures at the sign of any resurgence of attacker activity. This approach is now advisable because the organization is operating from a known starting point with zero compromised systems. From this point forward, the organization's success in combatting targeted threats long-term will be proportional to the organization's ability to rapidly detect and respond to indicators of compromise. Responders must identify and eliminate the threat before the attacker can escalate his privileges and move laterally to compromise unknown systems.

Strategic planning

After the remediation event, most organizations still have a significant amount of work ahead of them to develop an effective security posture. It is a challenge to maintain the security-focused momentum generated by the incident after this point. Successful organizations realize that long-term success will require significant investments in resources and in many cases fundamental changes to the way the organization views information security. Strategic initiatives typically address the following four areas.

Investing in people
Organizations with successful security programs recognize the importance of having personnel with the right skillsets focused in the right areas. To handle targeted threats effectively, this means having personnel dedicated to security monitoring and "hunting" targeted adversaries within the environment. These personnel should have strong technical foundations in the areas of networking, operating systems and forensic analysis. Penetration testing experience is also helpful.

It may be appropriate for smaller organizations to outsource monitoring to a third party. These organizations should ensure that the third party focuses on monitoring for the right threats. Medium to large organizations typically have multiple full time employees dedicated to this area of security, and may augment internal resources with expertise and data feeds from organizations that specialize in dealing with targeted threats.

Organizations should structure security personnel's roles and responsibilities to provide them time (and the mandate) to focus on activities such as tuning security tools and investigating suspicious events on a daily basis. Incident responders' must continuously update their skills and understanding of attacker techniques and investigative methodologies. To facilitate this process, organizations should define a continuing education plan and fund participation in technical training to help maintain skills.

Creating an 'investigation-ready' environment
To maintain visibility, organizations instrument the environment as if they were preparing for an incident investigation. Reflecting on the attack lifecycle, consider what monitoring activities would best enable security personnel to detect and track attacker activity.

Part of this effort consists of an ongoing dialogue between the incident response (IR) team, the security operations team, IT and the business. The organization should define a formal incident

MANDIANT | ADDRESS 2318 Mill Road, Suite 500, Alexandria, VA 22314 | PHONE 800.647.7020 | WEBSITE www.mandiant.com

response plan. The IR team should not write this document in a vacuum. Rather, this document should represent the result of a collaborative process between all the stakeholders in the incident response process. The following activities will also require collaboration across the broader organization outside of technology:

- Create an inventory of the systems that store sensitive data. Designate a business and IT point of contact for each.
- Define the incident response team's structure and responsibilities.
- Define touch points between the incident response team and the business. Determine internal notification and communication processes.
- Define meaningful outcome-based metrics (e.g., time to contain an infection) and obtain stakeholders' buy-in. Capture data and report to stakeholders on these metrics periodically.

Improving visibility is another key component of this effort. Provide incident responders with tools they can use to "hunt" targeted adversaries across the network. Integrate log sources into a tool that provides real-time query capabilities. Aggregate all available logs into a central repository that the IR team can query, including: network logs (DHCP, DNS, firewall, web proxy, NIDS); Active Directory, server and workstation logs; and HIPS and antivirus logs. Based on observed or predicted attacker behavior, the IR team should tune the log management system to reduce the number of events that they must review.

Threat intelligence can help reduce the size of the proverbial haystack, and successful organizations integrate such intelligence from multiple sources. This information helps the IR team better decide how to prioritize incidents. For example, a perceived attack by a relatively unskilled, though targeted, attacker should warrant a different response than a known successful attack by an advanced attacker.

To identify gaps in visibility, conduct tabletop exercises during which the IR team responds to a mock incident. During each step of the process, determine whether the information responders need to be effective was available. Define monitoring use cases for the IR team based on the lessons learned from these exercises. This will help to keep the IR team focused on the monitoring efforts that provide them the best visibility and enable them to be most effective. These use cases will also help to justify investments in people and technology to improve the organization's IR capability.

Define playbooks for incident response, starting with the most common activities the incident response team will need to execute. These playbooks should include touch points outside of the IR team, e.g. with business unit IT teams or outsourced service providers. The playbooks should also contain time requirements for completion, which will help to improve IR metrics. The process of working with all participants in the IR process is especially important in highly outsourced or decentralized environments.

Enhancing authentication and authorization

Targeted attackers leverage weaknesses in authentication throughout the lifecycle of their attack. Many organizations place undue trust in vulnerability management as their primary means to enhancing security. Implementing measures to reduce and control account privileges throughout the environment is a more effective means of inhibiting an attacker. Enhanced capabilities to monitor account usage help to detect anomalies that may indicate a successful intrusion and help responders during an incident.

This area's fundamental importance to security leads organizations to implement initiatives such as the following:

- Upgrading Windows workstations to Windows 7, which implements User Account Control (UAC).
- Implement application whitelisting for privileged and high-risk users' workstations.
- Removing local administrator rights from the vast majority of users.

MANDIANT | ADDRESS 2318 Mill Road, Suite 500, Alexandria, VA 22314 | PHONE 800.647.7020 | WEBSITE www.mandiant.com

- Reducing the number of privileged domain-wide service accounts. Fine-tuning the permissions allowed to the remaining accounts using Active Directory Group Policy Objects (GPOs) to restrict network and terminal services logons where they are not required.
- Implementing a set of accounts designed for use during an incident response. These accounts are normally disabled. This helps to reduce the risk caused by authenticating to compromised systems with privileged domain-wide credentials.
- Implementing multi-factor authentication coupled with an architecture that limits the allowed sources for authentication of the most privileged accounts.
- Implementing a means for randomizing and automatically changing local administrator accounts (or disabling the account altogether).

Improving the network architecture
Many organizations have a flat global network. This is especially true for organizations that grew through acquisitions. Targeted attackers understand that a flat network will enable them to move laterally from one system to almost any other system in which they are interested. Additionally and with few exceptions, most organizations do not need to allow workstations to communicate directly with other workstations.

Organizations implement initiatives such as the following to strengthen network security:

- Prevent workstations from accepting inbound network connections, except from a well-defined list of allowed management subnets that are necessary for software deployment, helpdesk and other operations.
- Documenting and understanding critical applications' network data flows.
- Periodically validating network device rule sets.
- Implementing web application firewalls to reduce the risk of web application vulnerabilities.
- Implementing web proxies for all users, restricting access to "uncategorized" web sites.
- Building restricted, high security zones for critical data and applications.
- Implementing highly secured, isolated "jump servers", from which all administrative activities must originate. This includes limiting the sources from which the most privileges can authenticate, e.g. a domain administrator can only authenticate to a domain controller when originating from a jump server.

Conclusion

To be successful remediating targeted intrusions by persistent adversaries, organizations must change the way they view incident remediation. Understanding the targeted attack lifecycle helps focus responders on the most effective countermeasures. Rather than addressing perceived risks, inhibit the activities the attacker is actually executing and improve your ability to detect and respond to them.

	ADDRESS	PHONE	WEBSITE
MANDIANT	2318 Mill Road, Suite 500, Alexandria, VA 22314	800.647.7020	www.mandiant.com