Beyond 'Check The Box': Powering Intrusion Investigations

Jim Aldridge

11 March 2014

Introduction

Many organizations have implemented a range of security products intended to facilitate security monitoring and incident response activities. However, few organizations effectively deploy, configure, and maintain a suite of tools that enables them to solve real problems when investigating an incident. This paper begins by describing a typical investigation's goals and then outlines five important incident response capabilities that organizations should build and maintain to prepare for the day when they are faced with a significant intrusion.

Investigations' Goals

The majority of the investigations in which I have been involved began with a third party, typically a government agency, notifying my client of the breach. Less often, security personnel detected unusual activity and performed triage procedures to confirm the incident. When an organization declares an incident, business, IT and security leadership become concerned with the following questions:

1. What information was exposed?
2. Do I need to notify regulators or customers?
3. What is the extent of the compromise?
4. How much money did we lose?
5. How did the attacker gain entry?
6. How do we effectively stop the attack and remove the attacker?

An investigation seeks to answer the first five questions based on available evidence. The investigation's findings enable the organization to design an effective remediation plan, which addresses the sixth. The remediation plan should be developed in parallel to the investigation and should contain activities designed to:

- Enhance the security team's visibility and capabilities to detect indicators of compromise.
- Enhance responders' abilities to respond to intrusions.
- Inhibit attackers' activities in future intrusions.

Refer to my BlackHat 2012 whitepaper titled "Remediating Targeted-threat Intrusions" [i] for additional discussion of remediation strategy. Note that in some cases it is advisable to begin containment immediately, even as a remediation plan is being developed.

Effective, near-term eradication of the attacker requires the organization to develop a detailed understanding of the intrusion. For example, if the attacker infected ten systems with backdoors, one needs to identify all ten systems and rebuild them. Identifying only nine of the ten is not sufficient. Organizations can also apply this understanding to improve their longer-term success rate against similar attackers. For example, if the attacker dumped password hashes from a domain controller, properly deploying an application whitelisting tool could make that action more difficult in the future.

The following are important factual questions that an investigation should attempt to answer; these answers will form the basis for answering the big-picture questions posed previously.

- When and what was the earliest evidence of compromise?
- How did the attacker gain entry?
- What is the latest evidence of attacker activity?
- What systems are (or were previously) under the attacker's control?
- What systems did the attacker access?
- What actions did the attacker execute on the systems with which he interacted?
- How does the attacker maintain access to the environment?
- How does the attacker operate inside of the environment?
- What tools has the attacker deployed?
- What accounts did the attacker compromise?

Answering these questions requires specific data points: an organization's ability to successfully attempt to investigate the intrusion is significantly dependent on the availability and quality of these data points.

Because an effective remediation effort is largely dependent on an effective investigation, the degree to which organizations master these capabilities significantly influences their ability to manage security threats.

Five key capabilities

In my experience, the following capabilities are important when conducting an enterprise-wide investigation:

1. Mapping an IP address to a hostname.
2. Identifying the systems to which a specified account authenticated.
3. Determining the systems that communicated with a specified Internet IP address.
4. Tracking domain name resolution attempts.
5. Identifying forensic artifacts across the environment.

These capabilities are necessary, but not sufficient, to effectively investigate intrusions. The most important ingredient is having dedicated, skilled, and experienced responders. I have worked with organizations that initially possessed none of these capabilities at the outset of the investigation but quickly developed them due to savvy response by team members.

This list is not intended to be comprehensive. Rather, these are the capabilities that are most needed and least frequently implemented effectively, in my experience.

The remainder of this document further explains each of these capabilities' application to an incident investigation.

Capability #1: Mapping an IP address to a hostname

Dynamic Host Configuration Protocol (DHCP) is typically used to assign IP addresses to user workstations and VPN clients. DHCP logs can be critical to making use of other information that the investigation uncovers. Organizations should develop the capability to rapidly determine a hostname, given an IP address and timestamp. For VPN pool IP addresses, organizations should be able to identify the username, remote IP address and remote hostname associated with a VPN connection given the IP address and a timestamp.

Imagine that your organization has just become aware of a network intrusion and has begun an investigation. Picture yourself as the person responsible for answering the key investigative questions outlined in the Investigations' Goals section. The investigation is progressing, and you receive an update on the analysis of one of the three systems that your team has identified so far as infected with backdoor malware.

- On this particular user workstation, files associated with a backdoor were created on October 4, 2013 at 13:24:53 UTC.
- Registry keys that provided the backdoor's persistence mechanism were last modified at 13:28:00 UTC on that same date.
- A Windows System Event Log entry for the start of the "PsExeSvc" service at 13:22:00 UTC looks suspicious.
- A corresponding Windows Security Event Log entry records a successful network logon event at that same time. The event identifies the IP address 10.224.124.90 as the source of the connection.

One important next step in the investigation is to investigate the system that had the IP address 10.224.124.90 on October 4, 2013 at around 13:22:00 UTC. It is possible that a legitimate tool or service connected using PsExec around that time. It is also possible the attacker was using that system as a pivot system at that time. In order to investigate the system, you'll need to confirm the hostname.

MANDIANT | ADDRESS 2318 Mill Road, Suite 500, Alexandria, VA 22314 | PHONE 703.996.3012 | WEBSITE www.mandiant.com

If you cannot identify the hostname, and don't identify other evidence, you may not be able to follow this thread of events further into the past. You may miss the opportunity to identify:

- Attacker tools present on that system.
- Indications that data was stolen and information that can help you determine what was stolen.
- IP addresses and domain names that may help you identify other infected systems.

Capability #2: Identifying the systems to which a specified account authenticated

All but the least sophisticated attackers leverage legitimate accounts to move within an environment. Strong authentication event query capabilities can enable investigators to track attacker activity across multiple systems and unravel an intrusion's sequence of events.

Capturing logs across the environment – from all servers and workstations - is challenging. Organizations should develop the following capabilities to investigate authentication events to help investigators identify compromised systems:

- Given the username and a time range, identify the systems to which the account authenticated. For accounts the attacker used during that timeframe, the system names that result from this search may have been compromised.
- Given a system name or IP address and a time range, identify the systems to which the given source system connected. This search is useful when there is evidence that the attacker was active on a particular system during a particular time window, and can identify systems compromised through lateral movement.
- Given a system name or IP address and a time range, identify the sources of successful or failed connection attempts to the system. This search can be useful to determine the source of lateral movement when investigating a compromised system, and can help identify systems the attacker was using as pivot points.

At a minimum, authentication events pertaining to domain accounts, logged on domain controllers, should be captured in a searchable manner. Such events should be logged on all systems, regardless of whether they are fed into a SIEM, so that they will be available to investigators for analysis on a per-system basis.

Turning your attention to the investigation, you would like the team to become more comfortable that their list of compromised systems is comprehensive. You're aware that as-yet-unidentified systems the attacker accessed could contain artifacts of data theft or application exploitation activities. These could be significant to the remediation of this incident. For example, perhaps the attacker compromised application credentials stored in a database. He connected to the database server by tunneling RDP through one of the implanted backdoors. Failure to identify that file server as an "accessed" system and analyze it may lead you to miss the opportunity to identify the exposure of those application credentials.

The following relevant data points that have been developed to-date:

- The attacker has used the "DOMAIN\privileged_service" account for lateral movement.
- The investigation has identified five systems that the attacker accessed, likely while performing internal reconnaissance.

One way to identify additional attacker access, given this information, is to look for authentication events that look anomalous given the source, destination, or account used.

If all Windows system logs were fed into the security information and event management (SIEM) tool, it would be possible to search this information to identify successful login events. Unfortunately, though your organization has a SIEM tool, its scope is limited: only domain controller logs are fed into the system. Fortunately, Kerberos Service Tickets logged on domain controllers can provide the information you need.

When an Active Directory user receives access to a resource, for example a user mapping a remote file share, the domain controller issues a Kerberos Service Ticket[ii]. When the domain controller issues

MANDIANT | ADDRESS 2318 Mill Road, Suite 500, Alexandria, VA 22314 | PHONE 703.996.3012 | WEBSITE www.mandiant.com

the Service Ticket, if appropriate auditing options are enabled[iii], the Domain Controller will log an event that includes the following important fields[iv]:

- User name and domain name
- Client Address (IP address where the user resides[v])
- Service Name (contains system name to which the user authenticated[vi])

By querying Kerberos Service Ticket audit events generated by domain controllers, we can identify:

- The IP addresses of systems from which the "DOMAIN\privileged_service" account was used for authentication. These systems may be compromised and the attacker may have used them as pivot points.
- The names of systems to which this account authenticated within a given timeframe, or from a particular source system. These systems may also be compromised; perhaps the attacker dumped password hashes from them, connected to them to steal files, or installed a yet-undiscovered backdoor on them.

The availability of sufficient audit logs to facilitate this type of query can greatly help investigators scope an intrusion.

Capability #3: Determining the systems that communicated with a specified Internet IP address

One way that targeted attackers maintain access to an environment is through the use of backdoor malware. This type of malware communicates from infected systems to command and control (C2) infrastructure hosted on the Internet, for example at a hosting provider or another compromised organization. The attacker remotely controls systems through the backdoors by issuing commands through his C2 infrastructure.

Organizations should have the capability to search metadata associated with Internet-bound traffic and identify communications with a specified Internet IP address, or to search for any outbound traffic that originated from a specified internal system within a given timeframe. The following information should be available and searchable: timestamp, source IP address and port, destination IP address and port, size of data transferred (packet size or flow bytes).

The ongoing investigation has identified that a backdoor named "service.exe" is hardcoded to communicate with a C2 server located at 1.2.3.4. Analysts have confirmed that four systems are infected with this backdoor.

With this information, you would like to identify all systems that have communicated with 1.2.3.4, as this may indicate that they are or were infected. Though full packet capture would be helpful, that capability is not feasible for most organizations to implement at scale. This need can be met effectively by capturing and making the following types of data searchable:

- Netflow data from border firewalls
- Firewall logs showing packets that successfully traversed the border to the Internet
- Web proxy logs

This network connection metadata can also be used to detect some changes in attacker activity. For example, if the four systems infected with the "service.exe" backdoor are beaconing regularly, and the typical traffic exchange is 800 bytes, a sudden increase in traffic size into the multiple-megabyte range could indicate that the attacker is actively using the backdoor.

Depending on your network topology, it may also be important to log network address translations (NAT) on border firewalls so that internal source IP addresses can be determined for identified traffic. Reflecting on Capability #1 – "Mapping an IP address to a hostname" – consider that the information provided by the firewall log searches may produce data with which you cannot take action if you do not have DHCP logs. Ensure that you can quickly tie the connections back to a hostname.

MANDIANT | ADDRESS 2318 Mill Road, Suite 500, Alexandria, VA 22314 | PHONE 703.996.3012 | WEBSITE www.mandiant.com

Capability #4: Tracking domain name resolution attempts

Targeted attackers frequently configure backdoor malware using fully-qualified domain names (FQDNs) rather than IP addresses. This provides them flexibility. If one of their C2 servers is discovered or shuttered, they can update the domain record and quickly point all of their implants to the new C2 server. They can also increase the difficulty of your investigation by pointing the domain record to a C2 server only when they are actively using the backdoor, and leaving it pointing to a "domain parking" IP address or RFC 1918 address at all other times. Organizations should be able to query these logs to determine what source systems attempted to resolve a given FQDN.

The investigation continues to progress, and one of your network analysts has identified a piece of traffic that concerns her. Your team has identified evidence that the implanted backdoors are configured with a combination of eight IP addresses and four FQDNs. Based on the IP addresses, and current resolution of the FQDNs, it appears that the attacker is hosting C2 servers within three distinct organizations.

As part of the investigation, you have implemented some new tools to gain network visibility into certain types of traffic. Among other things, this capability allows you to monitor traffic to and from interesting IP blocks. As you identified organization #1, #2 and #3 as hosting C2, the analyst began using this system to flag all the IP addresses associated with each organization.

Today, the network analyst observed a single successful DNS query response that originated from organization #2's IP address space. Examining the payloads of the DNS responses indicates that a workstation in your environment successfully resolved "bad.organization-two.com", which resolves to an IP address located in organization #2's IP address space.

Quickly checking your DHCP logs, you determine the hostname associated with the source of the DNS query for "bad.organization-two.com" and conduct a quick triage of that system. Analysts determine that the system is infected with a previously-unseen backdoor that was configured to communicate with a C2 at "bad.organization-two.com". Other systems may be presently infected with that backdoor: you would like to query logs for any system that resolved "bad.organization-two.com", as those systems are likely infected.

Even though you are monitoring firewall logs for connections to the known IP addresses, it is possible the attacker has infected some systems with the backdoor variant that communicates with the C2 using the FQDNs, and that these FQDNs are currently "parked", so there is currently no TCP/IP traffic to log at the firewall.

Having historical DNS query logs would enable you to quickly determine if there are other systems of interest that have resolved one of those two C2 domains. It is possible that there were systems previously infected, where the attacker has since removed the malware or the systems have been rebuilt in the normal course of IT operations.

Most large organizations have multi-tiered DNS architectures that make this type of logging a challenge. Most architectures point DNS clients at a local or regional domain controller for DNS resolution. This means that Windows Server's DNS debug logging feature[vii] must be used to generate a flat text file. To monitor DNS logs at enterprise scale, the organization must have a SIEM to collect this information. Purpose-built DNS appliances can help simplify the DNS architecture to be more conducive to logging.

When contemplating DNS logging, the volume of data is a significant concern. A large percentage of DNS resolution requests tend to be for internal resources, or well-known domains associated with popular websites, your own organization, service providers with whom you work, and your customers, suppliers, and business partners. Rather than trying to log all DNS queries, first monitor the queries to determine which domains can be filtered out and not sent to the SIEM to lessen storage requirements. For example, filtering out the top 25 queries, which constitute 50% of the volume of DNS queries, would halve storage requirements.

MANDIANT | ADDRESS 2318 Mill Road, Suite 500, Alexandria, VA 22314 | PHONE 703.996.3012 | WEBSITE www.mandiant.com

Capability #5: Identifying forensic artifacts across the environment
An effective investigation will identify indicators of compromise (IOCs) as it progresses. The term "indicator of compromise" may refer to artifacts associated with a specific intrusion or attacker, for example artifacts that result from infection with a particular piece of malware. The term can also refer to more general conditions that indicate malicious activity, such as a Windows service that is configured with a DLL that resides in an unusual path. To be effective at scale, investigators need a way to quickly search for the presence of these indicators across the *entire* environment.

On one infected system identified during your investigation, analysts identified evidence of the following sequence of events:

- The attacker used the "DOMAIN\privileged_service" account to connect to a file server and browse files.
- The network session to the file server was initiated from SERVER_A, which is infected with the "backdoor_A" remote access Trojan.
- "Backdoor_A" persists by hijacking the legitimate "Routing and Remote Access (RRAS)" service to point to a malicious DLL named "sneaky.dll".
- The backdoor communicates with a C2 server located at "bad.domain.com", which currently resolves to 23.34.56.78.

Based on that evidence, the team can define the following indicators of compromise[viii] to describe that activity:

- Security event logs contain a successful authentication event by the "DOMAIN\privileged_service" account.
- The registry key for any service DLL file name contains "sneaky.dll".
- A file has the MD5 hash 3a185c77d533d12544bfc6a24d7d2a75 (matches the malicious service DLL).
- A file was compiled on December 4, 2013 at 05:22:13 UTC and has a size of 20,241 bytes (may catch variants of the malicious service DLL that are very similar, but not exactly a hash match).
- An executable or DLL imports from "ws2_32.dll" and also imports all of the following functions: "RegisterServiceCtrlHandlerA", "RegQueryValueExA", "OpenServiceA", "InitSecurityInterfaceA" (may catch variants of the malicious service DLL, but will likely have false positive hits).
- A running process has a mutex named "733f0_fd3t" (may catch other pieces of malware by the same author, who may prefer to use this name for a mutex).
- The system's DNS cache contains "bad.domain.com".
- The system has an established TCP connection to 23.34.56.78.

Ideally, investigators would have the ability to search the whole environment quickly for matches to these indicators to identify this particular malware, variants of the malware, or suspicious uses of a known compromised service account that could indicate the attacker accessed a system.

Going beyond searching for known malware, investigators should have the capability to perform analytics for attacker TTPs across the environment. While not strictly a search for IOCs, this capability can help identify activity that was associated with unknown malware, or that was unrelated to malware. In the previous example, the attacker installed malware as a service. The capability to review *all* services in the environment for anomalies may help identify a second (currently unknown) backdoor, also configured as a service. Capabilities to review other persistence mechanisms, such as registry keys and scheduled tasks en masse for anomalies are also helpful.

IOCs are important tools for capturing and applying investigative context. Consider an example that has nothing to do with malware or tools. An attacker accesses a system through a web proxy and their remote system name, which may be unique, is captured in event logs. This sometimes occurs when less disciplined attackers proxy remote desktop protocol (RDP) connections from their "home" workstation to pivot systems inside the target environment, using implanted backdoors' SOCKS proxy functionalities. Capturing and searching for this condition, i.e. the unique system name, would help identify systems to which the attacker has connected.

MANDIANT | ADDRESS 2318 Mill Road, Suite 500, Alexandria, VA 22314 | PHONE 703.996.3012 | WEBSITE www.mandiant.com

Investigators will encounter challenges when trying to implement the capability to search for IOCs at scale. Without resorting to a commercial tool designed for this purpose, several tools will need to be utilized. Ingenuity, perseverance, and scripting skills will also likely be necessary. Computer Incident Response Teams (CIRT) should define use cases for searching the environment for IOCs based on prior investigations and third-party case studies. Then develop playbooks and scripts in advance so that you will be positioned to act quickly when investigating an incident.

In addition to the logs discussed previously, organizations should consider the following techniques to search for IOCs if purpose-built tools are not available:

- Most organizations already have antivirus tools widely deployed. These tools are typically only useful to detect exact hash matches to known malware. Before submitting malware to antivirus vendors, it is important to consider the potential impact on your investigation and containment effort. During an investigation, particularly of an active, targeted threat, submitting malware hashes to antivirus vendors is typically not advisable. Look for capabilities to leverage custom signature files in "alert only" mode.

- System and configuration management software tools can often be used to search for file system and registry artifacts. During an active response scenario, keep in mind that the data you are querying may be days or weeks old. Also consider that the data may only reflect the latest snapshot from the management tool; e.g., if a registry key associated with a backdoor was present last month (meaning that the system was compromised), but is no longer present, searching the configuration management database may not yield a positive result.

- Email system message tracking logs can be leveraged to search for known spear-phishing subject lines searches, attacker source IP addresses, and attacker-specific mailer software strings.

- Network intrusion detection systems (NIDS) and SIEM tools can be used to create customized alerts for network activity. When coupled with effective malware analysis, custom NIDS rules can help identify variants of attacker tools based on elements of C2 protocols that are unlikely to change between variants.

- Some vulnerability scanners have the capability to search for IOCs such as artifacts associated with running processes, current network connections, registry keys, file (hash /filename matches).

- PowerShell and Windows Management Instrumentation interface (WMI) can be leveraged[ix] to remotely query file and registry information.

Conclusion

To enhance your CIRT's capabilities in these areas, begin by developing a set of incident response use cases. Review prior incident reports and post-mortem analyses and document your own prioritized wish list of capabilities. Use data from your past incidents to illustrate the impact of not having the desired capabilities. Tailor the list to your environment and operations.

If you don't have sufficient data based on your organization's prior experience, apply others' descriptions of targeted attack scenarios to your environment and current capabilities. If your capabilities have not been tested with a real, widespread, targeted intrusion, conduct a simulation. The simulation should go beyond a table top exercise and include walking through the technical investigation steps. This will either prove that your current processes are sufficient or highlight areas for improvement.

Use the simulation's results to define requirements for new roles, processes, and tools. Once the CIRT team has defined these requirements, develop an accurate understanding of your current capabilities against those requirements.

MANDIANT | ADDRESS 2318 Mill Road, Suite 500, Alexandria, VA 22314 | PHONE 703.996.3012 | WEBSITE www.mandiant.com

Finally, implement performance-based CIRT metrics. Two of my favorites are:

- Mean-time-to-remediate: Once you identified the incident, how long did it take until it was contained?
- Mean-time-to-detect: How long were the attackers inside your environment before you detected them?

This will help to quantify incidents' impact to your organization and help drive changes to improve the incident response process. If I were a decision-maker, those types of metrics would help me make more informed decisions – for example a decision to allocate resources toward improving these five capabilities.

---

[i] Remediating Targeted Threat Intrusions whitepaper: https://www.mandiant.com/library/BH2012_Aldridge_RemediationPaper.pdf

[ii] For a technical explanation of Kerberos and the Active Directory authentication process, refer to the following references: https://blogs.technet.com/b/askds/archive/2008/03/06/kerberos-for-the-busy-admin.aspx, http://technet.microsoft.com/en-us/library/bb742516.aspx

[iii] Reference on Kerberos audit logging options: http://technet.microsoft.com/en-us/library/cc787176(v=ws.10).aspx

[iv] References on Kerberos ticket event log messages:
http://www.ultimatewindowssecurity.com/securitylog/encyclopedia/event.aspx?eventid=673,
http://www.ultimatewindowssecurity.com/securitylog/encyclopedia/event.aspx?eventid=4769, https://www.sans.org/reading-room/whitepapers/forensics/windows-logon-forensics-34132

[v] When the Remote Desktop Client is used to authenticate to a Windows 2003 server where Network Level Authentication (NLA, http://technet.microsoft.com/en-us/library/cc732713.aspx) is not used, the Service Ticket event will reflect the target system's IP address rather than the source workstation's IP address. In this case, the investigator will need to hope that the target of the RDP logon has sufficient logs to identify the source of the logon. If instead the user uses a network logon method, such as "net use" to map a share, or PsExec, the Service Ticket event will reflect the source workstation's IP address (i.e., the system from which the attacker is originating).

[vi] Depending on the logon method used, the domain controller may log an event that indicates the username and client source IP, but does not indicate the service/system to which the user authenticated. These EID 540/4624 events will typically be noticed proximate to interactive logons where the user account authenticates to the domain controller to retrieve group policies and logon scripts.

[vii] Microsoft Windows DNS log reference: http://technet.microsoft.com/en-us/library/cc776361(v=ws.10).aspx

[viii] The OpenIOC framework is one way to codify indicators of compromise. The following presentation and blog post describe OpenIOC:

- "Identifying and Sharing Information with OpenIOC" by Doug Wilson: http://scap.nist.gov/events/2011/itsac/presentations/day2/Wilson%20-%20OpenIOC.pdf
- "OpenIOC: Back to the Basics" by Will Gibb and Devon Kerr: https://www.mandiant.com/blog/openioc-basics/

This blog series provides case studies on how to define indicators of compromise based on investigative activity:

- "Investigating with Indicators of Compromise" by Will Gibb and Devon Kerr: https://www.mandiant.com/blog/openioc-series-investigating-indicators-compromise-iocs-part/, https://www.mandiant.com/blog/investigating-indicators-compromise-iocs-part-ii/

[ix] PowerShell and WMI references:

- "Live Response Using PowerShell" by Sajeev Nair: https://www.sans.org/reading-room/whitepapers/forensics/live-response-powershell-34302
- "Incident Management in PowerShell: Recovery, Lessons Learned" by Wolfgang Goerlich: http://www.poshsec.com/category/incident-response/
- "WMIC for Incident Response" by Mike Pilkington: http://digital-forensics.sans.org/blog/2010/06/04/wmic-draft
- "PowerShell for ForSec & Incident Response: A Brief Musing" by Claus Valca: http://grandstreamdreams.blogspot.com/2013/09/powershell-for-forsec-incident-response.html
- "The Power of PowerShell Remoting" by Mike Pilkington: http://digital-forensics.sans.org/blog/2013/09/03/the-power-of-powershell-remoting
- MSDN's WMI reference: http://msdn.microsoft.com/en-us/library/aa394572(v=vs.85).aspx

MANDIANT | ADDRESS 2318 Mill Road, Suite 500, Alexandria, VA 22314 | PHONE 703.996.3012 | WEBSITE www.mandiant.com