Fidelis™ Cybersecurity

The First 72-Hours

How to Approach the Initial Hours of a Security Incident

www.fidelissecurity.com

Contents

I've Got an Alert. Now What?

The initial signs that you have a security incident on your hands are rarely black and white. Perhaps you got a call from law enforcement that they've seen your confidential data in the wild. Or, maybe a trading partner reported unusual activity. Even when the alert comes from your own security operations center (SOC), the first questions you have to ask yourself are "Is this a real incident?" and "How should I respond?"

Over the course of responding to hundreds of critical security incidents we have seen organisations approach the initial hours of an incident in every conceivable way. In most cases, the knee-jerk reaction is to try and resolve the incident immediately. Whilst it's understandable that you would want to immediately take systems offline or block IP addresses, these actions can often be counterproductive and even increase the length and risk that the incident poses.

In our experience, a rapid response in the first 72-hours is critical. Whilst all organisations want to quickly resolve the incident, in order to know how you should respond it is important to first understand what type of incident you are dealing with.

Whilst all organisations want to quickly resolve the incident, in order to know how you should respond it is important to first understand what type of incident you are dealing with.

Understandably, most executives become uneasy when they hear that the first step may not be stopping the attacker. Whilst removing the attacker is obviously the end goal, security teams must first understand the nature and scope of the incident to identify the course of action that will be most effective at removing the attacker while balancing the risk to the organisation and the disruption that the response can cause.

Determining What Type of Incident You Face

How you respond during the first 72-hours depends first and foremost on how long the attacker has been in your environment. At Fidelis, when we respond to security incidents we typically see two types of situations. We call them incursion detections and persistence detections. In short, an **incursion detection** means the security team has detected the attacker within the first 48-hours of the initial compromise. A **persistence detection**, on the other hand, means the attacker has been in the network for months (or even years) and has established a foothold.

In most cases, the security team will turn to the SIEM and query it for indicators of compromise (IOCs) based on the original alerts or notifications that alerted the team to the incident in the first place. This could include IP addresses, domain names or MD5 hashes. If the organisation lacks a SIEM, the team need to query the various log sources manually.

The table below highlights the differences between incursion and persistence detections.

Comparison of Attack Detection

	INCURSION	PERSISTENCE
Nature of Intrusion	The attackers have gained initial access within the last 48 hours.	The attackers have established a foothold. In most cases they will have been in the environment for weeks, months or even years. One must assume they are using valid user credentials, have moved laterally and stolen data.
Identifying Root Cause	Easy to determine — can often be done in 1–3 days.	Difficult to determine — can take 2 to 4 months.
Investigation and Containment	Relatively easy to perform because the root cause is easy to determine. Typical length ranges from 1 to 3 weeks.	Drawn-out because information is slow to obtain, the point of infiltration is rarely identified and the root cause is more difficult to ascertain. Typical length ranges from 2 to 4 months.
Eradication Approach	Sequential Remediation: Incident responders mitigate compromised devices and eliminate malware as soon as it is identified.	Simultaneous Remediation: Incident responders execute all remediation activities at the same time. The eradication event can range from 1 day at small organisations to 3 days at large organisations.
Remediation Cost	The cost is typically built into existing security budgets.	Typically requires additional budgeting on top of pre-allocated security budgets.

Responding to an Incursion Detection

If you've determined you are dealing with an incursion detection then your immediate objective is to prevent the attacker from establishing a foothold. Because attackers typically establish a foothold by using custom malware, command and control and third-party application exploits, incident responders should be on the alert for web server exploitation, email phishing campaigns, social engineering attacks, the planting of web shells and SQL injections.

Response Roles and Responsibilities

Responding to an incursion detection is more straight forward than responding to a persistence detection. The table below, details the activities that each role typically performs in the course of responding to an incursion detection. As a general guideline, within the first 24-hours, response teams should only try to understand the attacker's motivations in their spare time. Whilst knowing the attackers' motives can help guide eradication and remediation activities down the line, it should be a lower priority than other activities.

Incursion Detection Response Activities

ROLE	0–24 HOURS	24–48 HOURS	48–72 HOURS
Network Administrators	• Identify IP addresses that are involved (both good and bad)	• Standby to assist the security team • Manually pull logs that do not feed into the SIEM and provide them to the security team	• Maintain standby status
System Administrators	• Quarantine affected systems as malware and IOCs are identified	• Stand by to assist the security team • Continue to quarantine affected systems	• Maintain standby status
Technology Administrators	• Obtain copies of malware	• Stand by to assist the security team • Analyse identified behaviour to determine if it is normal or abnormal	• Maintain standby status

Incursion Detection Response Activities *(continued)*

ROLE	0–24 HOURS	24–48 HOURS	48–72 HOURS
Security Team	• Identify security tools that detected the attack • Identify security tools that failed to detect the attack • Identify and remove malware • Set alerts for the IP addresses, user accounts and malware involved • Block malicious communication	• Perform SIEM lookups for good IP addresses, bad IP addresses, malware and user accounts to determine if systems were compromised prior to when the incident was detected • Update all security tools with the attack signatures • Fix security tools that failed to detect the attack • Determine why the attack succeeded and was not, detected or prevented	• Add a SIEM alert rule for the incident, including all related IOCs • Identify system, network and technology changes needed to prevent future occurrences • Ensure the security operation centre (SOC) has what they need to adjudicate reattempts
Incident Response Lead	• Initiate incident tracking and escalation • Start intra-security team dialogue to facilitate internal communication regarding events unfolding and incident resolution actions • Start internal reporting that documents the details of the incident as they are identified	• Provide status reports to leadership stakeholders • Ensure information incoming and outgoing related to the investigation is controlled and secure • Begin to put the big picture together	• Provide formal presentation to leadership stakeholders • Begin documenting lessons learned and remediation recommendations
Security Management	• Begin reporting to the executive team and other stakeholders to inform them on the progress of incident response activities • Begin searching for the reason for the attack	• Provide inter-departmental support to IR team • Notify other departments as needed to obtain cooperation and facilitate IR efforts • Begin dialogue with external stakeholders, legal and/or regulatory bodies as required	• Review initial lessons learned and short-term remediation recommendations • Assist security team in preparing incident-specific training and awareness campaigns designed to improve the security culture and promote good security workplace behaviour • Engage with third-party IR partner if necessary

Lessons Learned

In the past year, the Fidelis Incident Response team has identified, contained and eradicated incursion detections in less than three weeks. Response efforts rarely last more than a month. However, we did encounter several cases that took longer. In these instances, the delay occurred either because logs were not easily obtainable or we had to wait for managed service providers (MSP) or other external vendors who could not make changes to the systems or the network outside of pre-defined "change windows".

For example, at one organisation, the MSP insisted on waiting for the weekly change window to make changes to firewalls in order to block bad IP addresses. The delay gave the attacker three extra days to move around the network and extended the remediation effort by 3 weeks.

When administrative actions delay incident response, it is important to address the cause to avoid such delays in the future. In the case above, the victim organisation modified its agreement with the MSP to allow security emergencies to dictate changes to firewall settings at any time.

In another example, a client's SIEM was not configured properly. It was overloaded with data and on such outdated equipment that if a query did not return a result in 30 minutes, then the entire system would crash. In this situation, running queries on historical data was not feasible. Administratively, the client spent many days attempting to duct tape their system and we determined that it would cost less to copy the data than to try and make the current system work. Our team spent a week extracting the large databases and placing it into our own off-line database, which extended the investigation by two weeks.

Responding to a Persistence Detection

When you identify that you are dealing with a persistence detection, it means attackers have been in the network for months and often, for more than a year. Unfortunately, that means that not only have the attackers already infiltrated your network, they have also established a foothold with backdoors and RAT capabilities. You should also assume that they have obtained privileged account access to systems including VPN access. They will have mapped out portions of your network and planted mechanisms to exfiltrate data such as web shells, proxies and RATs.

Because of the length of time attackers have been in the network and the depth of their infiltration, traditional incident response plans and approaches are insufficient to remove the attacker, re-secure the network and recover from the incident. In most cases an organisation will want to engage the skills of an experienced external incident response team.

No matter who is leading the response, the first step is to determine the scope of the incident and understand where the attackers are and what systems they have access to. Only then, can you develop a coordinated eradication plan to eject the attacker from the environment. It is vital to learn as much as possible about the attacker before taking action. If you miscalculate and "miss" some of the persistence mechanisms the attacker has put in place they will often "go dark" and return weeks later using an undiscovered dormant backdoor.

Once you understand the scope of the incident, eradicating attackers involves planning and coordinating an event that will simultaneously turn out the lights on the attacker. It requires a tightly orchestrated set of activities that include resetting user accounts, removing malware, remediating affected systems, blocking known bad IP addresses and implementing custom signatures for endpoint active defence tools (e.g., antivirus). In larger organisations, this can involve upwards of tens-of-thousands of user accounts, over twenty systems and myriad families of malware.

Response Roles and Responsibilities

In the first 72 hours of a persistent detection, the incident response team should focus on gaining situational awareness and establishing a central hub for the subsequent investigation, analysis, containment, eradication planning and — ultimately — the eradication event. It is vital to act quickly because as time goes by valuable evidence such as logs and endpoint artifacts can be overwritten or deleted.

The table on the next page details the activities that IT and security teams generally perform in the course of responding to a persistence detection.

Persistence Detection Response Activities

ROLE	0–24 HOURS	24–48 HOURS	48–72 HOURS
Network Administrators	• Identify IP addresses that are involved (both good and bad)	• Standby to assist the security team • Manually pull logs that do not feed into the SIEM for turnover to the security team	• Maintain standby status • Monitor network operations for anomalous activity and report to the security team
System Administrators	• Document affected systems as persistence mechanisms • Identify malware and IOCs.	• Standby to assist the security team • Continue to document affected systems	• Maintain standby status • Monitor system operations for anomalous activity and report to the security team
Technology Administrators	• Obtain copies of malware and other persistence mechanisms	• Standby to assist the security team • Analyse identified behaviour to determine if it is normal or abnormal	• Maintain standby status
Security Team	• Identify security tools that detected the attack • Identify security tools that failed to detect the attack • Identify malware • Set alerts for the IP addresses, user accounts and malware involved	• Perform SIEM lookups for good IP addresses, bad IP addresses, malware and user accounts to determine if breach activity occurred prior to incident detection • Designate the attack as persistence detection • Determine why the attack succeeded and was not, detected or prevented	• Expand scope of investigation • Start coordination of activities that will lead to an expulsion event • Establish formal reporting procedures with the SOC about IOCs directly related to this incident
Incident Response Lead	• Initiate incident tracking and escalation • Start intra-security team dialogue to facilitate internal communication regarding events unfolding and incident resolution actions • Start internal reporting that documents the details of the incident as they are identified	• Provide status reports to leadership stakeholders • Ensure information incoming and outgoing related to the investigation is controlled and secure	• Provide formal presentation to leadership stakeholders • Provide formal information exchange procedures taking into account volume and security • Establish regular status update meetings

Persistence Detection Response Activities *(continued)*

ROLE	0–24 HOURS	24–48 HOURS	48–72 HOURS
Security Management	• Begin reporting to the executive team and other stakeholders to inform on the progress of incident response • Seek the reason for the attack	• Provide inter-departmental support to IR team • Notify other departments as needed to obtain cooperation and facilitate IR efforts • Begin dialogue with external stakeholders, legal and/or regulatory bodies as required	• Coordinate external stakeholder communication • Engage third-party IR partner if surge support is needed

Lessons Learned

The majority of the investigations the Fidelis Incident Response team performs are persistence detections. In our experience, the attackers often know more about the network than the system and network administrators responsible for managing the environment. In one case the attackers had been present in the organisation's environment for 500 days and had accessed every system and network without being detected.

To completely remove all traces of the attacker the team worked quickly to identify which systems and technologies were affected, what malware was involved and which user accounts the attacker had access to. This involved collecting Active Directory, firewall, proxy and security logs, going back as far as possible. In this case, the logs were being stored for 500 days and were only a couple of days away from being overwritten. If we had not gathered the information within the first three days, we would have lost the information that painted the picture of what happened in the past.

Ultimately, the investigation team was able to determine the malware that the attacker used to gain access and identify the compromised systems and user accounts. Because we acted quickly to collect and store critical logs, we were able to work the incident to a successful eradication and remediation.

Conclusion

Determining whether you have detected the initial incursion or an attacker that that has been persistent in your environment for weeks or months is easier said than done. There is often limited information and visibility in the first few hours is often spotty at best.

If you are not able to easily distinguish the type of incident you are dealing with we recommend treating it as an initial incursion and following your normal incident response processes. If at any point during the process you determine that the incident is a persistent detection, then you can change approaches and move from sequential eradication to a broader investigation aimed at identifying the full scope of the attack.

Fidelis Cybersecurity is creating a world where attackers have no place left to hide. We reduce the time it takes to detect attacks and resolve security incidents. Our Fidelis Network™ and Fidelis Endpoint™ products look deep inside your traffic and content where attackers hide their exploits. Then, we pursue them out to your endpoints where your critical data lives. With Fidelis you'll know when you're being attacked, you can retrace attackers' footprints and prevent data theft at every stage of the attack lifecycle. To learn more about Fidelis Cybersecurity products and incident response services, please visit **www.fidelissecurity.com** and follow us on Twitter **@FidelisCyber**.

 WP_Fidelis_First72_UK_1603