WHITE PAPER

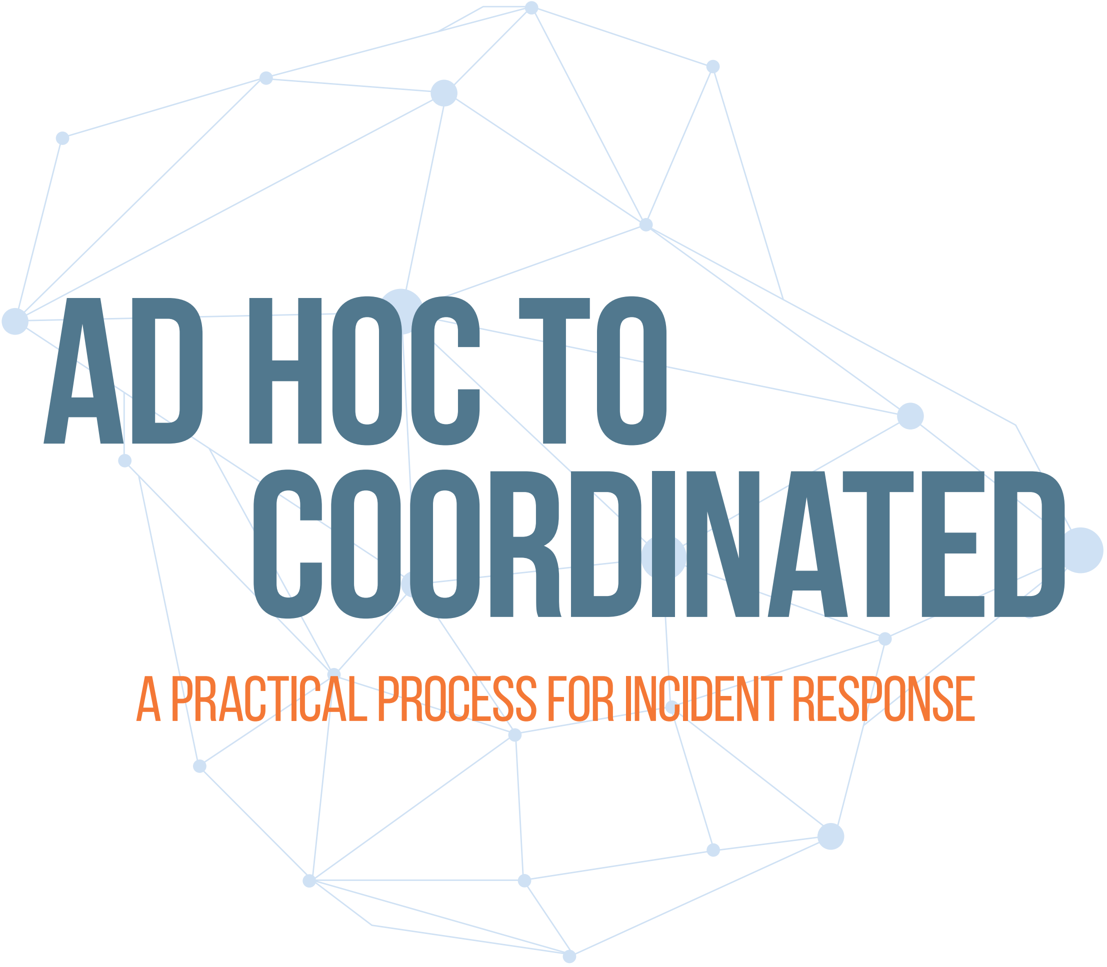

AD HOC TO COORDINATED

A PRACTICAL PROCESS FOR INCIDENT RESPONSE

If you're a security analyst working in incident response, you face a deluge of security alerts every day—so many that it's probably impossible to keep up. In fact, nearly one-third of organisations claim that they ignore more than 50% of all security alerts because they simply can't handle the volume.[1] Clearly outnumbered, incident response teams are overwhelmed with the sheer volume of security alerts received on a daily basis.

Investigating every single alert is impossible, and would be a waste of time considering the average rate of false positives exceed 40% in most organisations.[2] So, you need a practical way to act on security alerts that enables you to cover your bases in the investigation, prioritization, and verification to contain a threat.

You Need a Plan

There are many different security incident response plans. They range from the most formal, coordinated processes advocated by NIST and SANS to those created on the fly for a specific purpose—the ad hoc kind. Under pressure and trying to save time, IT teams often labor under ad hoc investigation techniques that don't contain the whole problem. To develop your incident response capabilities, you need a way to move beyond those techniques.

Using the Proofpoint four-step incident response process will help your team bridge the gap between makeshift efforts to a coordinated incident response effort.

Each step offers you increased situational awareness for better response. Learn about the benefits of each step, the barriers to address as well as the risks you may create when opting to save time by cutting corners.

These steps include:

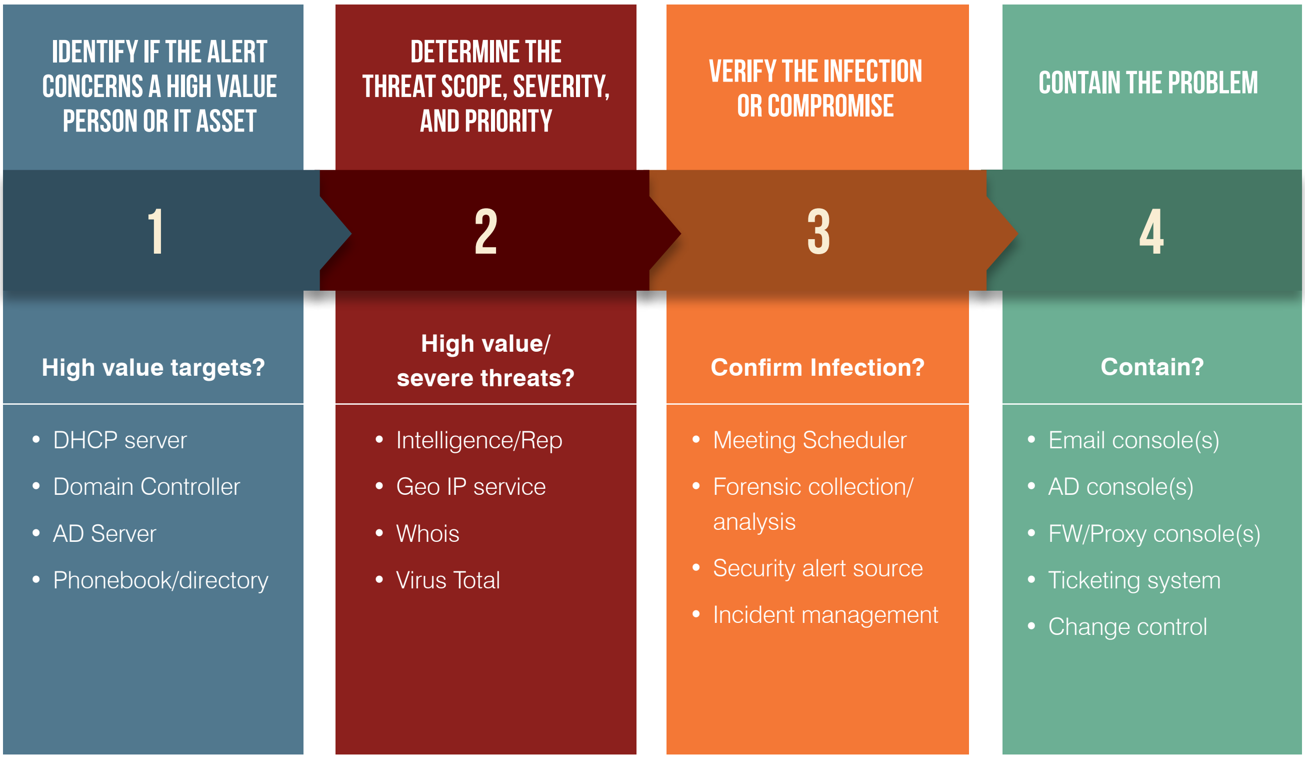

IDENTIFY IF THE ALERT CONCERNS A HIGH VALUE PERSON OR IT ASSET

1

High value targets?

- DHCP server
- Domain Controller
- AD Server
- Phonebook/directory

DETERMINE THE THREAT SCOPE, SEVERITY, AND PRIORITY

2

High value/ severe threats?

- Intelligence/Rep
- Geo IP service
- Whois
- Virus Total

VERIFY THE INFECTION OR COMPROMISE

3

Confirm Infection?

- Meeting Scheduler
- Forensic collection/ analysis
- Security alert source
- Incident management

CONTAIN THE PROBLEM

4

Contain?

- Email console(s)
- AD console(s)
- FW/Proxy console(s)
- Ticketing system
- Change control

[1] Enterprise Strategy Group (ESG) The Shift to Incident Response Automation and Orchestration, March 2016 https://drive.google.com/file/d/0B76__3ub6mzdM3V5UzBvSS1ZMVFmdTBoNFZfNEtKMVJQMmc0/view?pref=2&pli=1

[2]. Ponemon Institute, The State of Malware Detection & Prevention, March 2016 http://www.ponemon.org/blog/new-ponemon-study-on-malware-detection-prevention-released

INVESTIGATING THE TARGET OF ATTACKS

If a security alert concerns high-value systems or users, the need to respond is clear-cut. When an alert occurs, but there is no actionable data to reference, it can be difficult to decide whether the alert warrants further investigation. Time is a crucial factor, so finding ways to weed out unnecessary tasks and low priority alerts will accelerate your ability to contain the threat.

Figuring out the focus of the attack within your company is the first step. Is this a business-stopping event such as an attack on a transactional or public web server? Are key departments involved such as finance? Is intellectual property at risk, for example, a source code server? Or is a privileged user the target?

Common reference points for determining the details of an internal target include:

DHCP Server Logs

They allow you to identify the host name as well as the MAC address associated with an IP address. Query these logs so that you know if a set of security alerts relate to the same target or many. In coordination with querying domain controller and Active Directory (AD) server logs, determine what users are logged in, when they last changed their password, and the patch level of the endpoint. Use this information to branch out into other systems and figure out authentication with other systems.

Domain Controller Logs

By crosschecking domain controller logs against the timeframes in the security alert, you can determine what authentication took place with other systems.

AD Server Logs

Checking AD server logs is a standard part of incident response. AD server logs give you the details on the users and their group memberships. This helps you determine if a privileged user is the target. Secondarily, AD server logs may show recent escalation of permissions or group policy changes that could be an indicator of attack.

Phone Directory Services

Once you have established the priority of the security alert, it may warrant informing the target by phone about the situation.

Depending on various factors, this step usually takes between 5 to 30 minutes if all information sources are available. With preliminary information on the target defined, you can now take a closer look into the scope, priority, and severity of the threat.

INVESTIGATING THE TARGET OF ATTACKS

The Benefits of This Step

- You get context to make better decisions about response and priority. By knowing what is at risk, you can respond faster.
- You can identify users and adjacent systems that may be within the scope of the attack
- Identify whether the attack is moving laterally within your environment and using other systems as pivot points to launch an attack.

Barriers to Address

- Too much investigation time is spent capturing and organizing information about internal targets of attack.
- Waiting on other IT groups to gain access to logs for investigation
- No time to investigate all the possible places where high-value targets reside.

Risk of Skipping Steps

- You may overlook a user, endpoint, or entire department that is an integral part of an attack.
- Misjudging the scope of attack on adjacent systems and users.
- You cannot determine full extent of what or who is the target of the attack.
- The infection spreads while you determine the scope of targets within your environment.

The benefits of investigating the internal target of a security alert are clear. You now have the internal context to prioritise the threat and make decisions about investigating the matter further. If it is an AD or source code server under attack, or the alert is about several endpoints in the finance department, you know what to do. Obtaining full information on the extent of internal targets is vital so you can connect the dots between your internal targets and the threat severity and scope.

DETERMINING THREAT SEVERITY

By querying outside sources of information, security analysts get the situational awareness to understand the scope and severity of the attack so that they can better prioritise their response.

Pulling together information gathered on internal targets with external sources of threat intelligence enables you to answer these questions:

1. What is the severity and scope of the threat?
2. Who might be perpetrating the attack and where are they located? Is the attack coming from a country where the company normally does business? Is it in a part of the world known for bad actors? Are command-and-control (C&C) IP addresses involved in the alert? Are the IPs or URLs on a reputation list?
3. Is this part of a larger attack campaign or botnet? Are the attackers conducting other malicious operations?

Reference these sources of information to get details on the threat severity, where it came from, and where else the attack has been seen.

Geo IP Services

They are often seen as the first pass filter for prioritizing inbound threats. They pinpoint the location of the IP addresses listed in the security alert so you know if it relates to a specific actor group or suspect location. They also help you determine communication and data transactions have been going to a location known for malicious activity.

Intelligence & Reputation Feeds

Security analysts typically look to threat intelligence feeds to find current and historical metadata on IPs and hashes listed in security alerts. Threat intelligence feeds can provide background information on the attack and tell you if it involves a known C&C server, botnet, or if it is part of a larger campaign such as Dridex. But remember that using free sources of threat information is risky; the data is not verified and is usually not up to date.

Whois

Look here for insight into who is behind the domains and IPs used in an attack. It is commonly used to flag suspicious websites for investigation such as a domain with recent registration date, one that is close to expiring, or websites where the registrant and website are in different locations.

DETERMINING THREAT SECURITY

The Benefits of This Step

- You can prioritise security alerts effectively based on scope, severity, and background information of the threat.
- You can detect attack and campaigns earlier and contain them faster.

Barriers to Address

- Manual investigation cannot scale to meet the volume of security alerts, creating a backlog of security of security alerts left untouched.
- Free sources of threat intelligence produce more false-positive security alerts.
- Independent severity rating systems make it difficult to determine the real threat level.

Risk of Skipping Steps

- Skipping this step leaves you without the information you need to contain or prioritise the incident properly.
- External sources of intelligence do not take into account internal systems and potential impact.
- You cannot determine where attacker is in kill chain – and cannot preemptively respond.
- Free services can produce more false positive to investigate and slows down your incident response.

VirusTotal

Services such as VirusTotal are an incident response go-to resource that allows you to search on URL scan reports, file hashes, IP addresses, and domain names. It provides historical and current incident reports related to the IP addresses or domains listed in your security alerts.

With internal targets confirmed and outside research on the threat completed, you now have the information you need to understand, prioritise, and decide on next steps.

As one of the most time-intensive areas of incident response, verification is sometimes skipped in favor of going straight to containment activities. But skipping verification can lead to containment problems.

First, if you skip verifying the threat, you may be relying on re-imaging several if not many more systems than what is required to cover all possible scenarios. In this case, your team spends the majority of its time reimaging systems rather than investigating the threat. This leads to the second problem of missing the one endpoint or possibility a whole department that is also infected. This is why infections, backdoors, and other security risks persist within organisations—and it's how whack-a-mole containment processes start.

VERIFYING THE INFECTION OR COMPROMISE

With the reasonable suspicion of a compromised asset or user credentials, you must now get ahold of the suspect endpoint or server. This phase of incident response stands out for two reasons: it requires the most time and is the most crucial period for quick response. Without timely verification and containment, infections from malware spread or attackers get a head start in covering their tracks by deleting files and toolsets from compromised systems.

Before verification, contact the user whose endpoint appears to be infected. In the case of an IT asset like an AD or source code server, reach out to the IT operations team supporting the server.

Here is the play-by-play list of activities required to verify the incident:

Schedule a Meeting

Reach out to the employee and the IT team members to let them know about the suspected infection or compromise. Schedule a meeting to swap out a loaner PC or gain access to the suspected system so you can examine it and copy the hard drive.

Forensic Collection and Analysis

Isolate the endpoint and obtain a full disk copy to search for signs of malware.

Security Alert Source

SIEM systems import and normalise security alerts from all types of security devices. With a confirmed incident, look at the source

VERIFYING THE INFECTION OR COMPROMISE

The Benefits of This Step

- Identification of attacker artifacts as proof of compromise.
- You can determine the initial "beachhead" system and r oot cause that matches the security alert.
- Prevent future events by finding and mitigating all artifacts and evidence left behind, such as back doors left on systems.
- You can begin resolving confirmed infections or compromises.

Barriers to Address

- Coordinating an agreeable time to exchange a user's system with a loaner slows down the verification process.
- Reimaging system too soon may leave you unable to validate infection or compromise.
- Malware may have deleted its tracks and hid evidence of its presence.
- False positives lead to unnecessary work.

Risk of Skipping Steps

- Verification may not be performed on all infected systems
- You may prevent users outside the scope of the investigation from performing their daily work.
- The threat may be eradicated before you have fully investigated it—leaving you vulnerable to similar attacks in the future.
- Your systems may be re-infected and the attack may spread.

and what rules triggered the SIEM alert. Look for signs of other infected systems, unusual traffic patterns, or actions that only administrators would perform, such as changing group policies or escalating permissions.

Incident Management

To save future analysis time and increase knowledge about users and system information, document all employees and endpoints involved with a confirmed security incident. This is important system knowledge to retain. It reveals repeat offenders—the "clickers." You can also note tactics, techniques, and patterns of attacks to improve signatures and rules for detecting them in the future.

CONTAINING THE THREAT

Containment is why quick filtering and validating security alerts within and outside of your organisation is so important. You have to verify the threat before you can contain it. Any containment process has two goals: stop the spread and limit the damage as fast as possible.

For some companies, containment can be as simple as pulling the plug on the infected system and reimaging it. However, that is never an advisable approach, even as a stopgap measure, since you could disrupt employee productivity. A more likely scenario is that enough time has passed that the original system infected is now a pivot point used to infect other systems.

During containment, reference the following sources of information to trace the attack, identify all affected hosts, and contain the problem.

Email consoles

If the attack was email-borne, check the target and message ID to determine the recipient and payload. By searching for the message ID, you can remove it from mailboxes. As a containment measure, log in to the email gateway to disable and reset accounts for verified compromised credentials and any infected system.

Active Directory Services

Log in to AD services to lower permissions of affected employees and contain their identities to stop lateral movement and usage of compromised credentials.

WAF/FW/Proxy Dashboards

Block IP addresses, host names and URLs involved in the incident. If the attack is coming from a country where there is no legitimate business interest, it may be worth blocking regional locations. This is not always practical, but since attacks do tend to originate from obscure locations, it may be prudent to prevent future attacks coming from that area. Depending on the scope of attack, you may want to isolate certain network segments to better contain the threat.

CONTAINING THE THREAT

The Benefits of This Step

- Block access to infected sites, servers, compromised departments, servers
- Shut down attacker access to compromised systems
- Confidently return your organisation to normal business operations as soon as possible.
- Limit the effect of infections on neighboring IT resources.

Barriers to Address

- Lack of coordination and ownership between IT teams slows down response time.
- Change control procedures delays blocking of IP addresses and domains.
- Network and security teams may disagree on the right course of action.

Risk of Skipping Steps

- Delays in containment processes causes infections to spread.
- Blocking the wrong IP addresses disrupts employee work activity and hurts productivity.
- The attacker can use compromised identifies across other systems.
- Residual compromised files remain on systems, leading to reinfection.

Ticketing Systems

Ticketing systems are typically used to collaborate with other IT groups to contain an infection. This includes coordination of temporary systems until the infected endpoint is cleaned.

Change Control

In large organisations, incident details are noted in change control systems, including what you observed, when, and under what circumstances. Denote the IP addresses, domains, or geographic locations that were blocked and any other follow-up countermeasures.

Use this guide to evolve your incident response from an ad hoc reaction to a coordinated response. As you become more efficient, reducing time spent toggling between browser tabs and logging into numerous systems becomes more important. You'll want to move away from the practice of re-imaging systems as a containment measure and towards the strategic work of incident response while enabling better defense for your company.

Reaching the next level of incident response efficacy frees team members from the tedium of manually collecting and stitching together information to prioritise, validate, and contain incidents. Automation saves hours or even days per incident and minimises some of the less rewarding aspects of incident response.

The result: proactive defense and the ability to engage team member in more strategic work. Your team will spend less time and effort on investigating security alerts, especially the ones that prove to be false positives.

To transform your incident response process, you need Proofpoint Threat Response. It automates key areas of incident response, so you can prioritise, verify, and resolve threats 10 times faster.

Threat Response enables your organisation to:

Accelerate response

Get automatic collection of external threat information and internal target data. This gives security analysts full situational awareness to investigate and prioritise security alerts quickly.

Save hours per incident

Automated, built-in infection verification dramatically reduces time spent chasing false positives and confirming infections.

Instantly contain threats

Automated workflows trigger response actions to immediately quarantine and contain infected systems.

Don't waste your incident response team's time doing the mind-numbing manual work of investigating security alerts when automated solutions are available. Threat Response is the force multiplier you need so that your team can resolve incidents in less time with less effort.

Learn more about how Threat Response can automate your incident response process, save hours or days per incident, and help you to engage your team in more strategic work.

ABOUT PROOFPOINT

Proofpoint, Inc. (NASDAQ:PFPT), a next-generation cybersecurity company, enables organizations to protect the way their people work today from advanced threats and compliance risks. Proofpoint helps cybersecurity professionals protect their users from the advanced attacks that target them (via email, mobile apps, and social media), protect the critical information people create, and equip their teams with the right intelligence and tools to respond quickly when things go wrong. Leading organisations of all sizes, including over 50 percent of the Fortune 100, rely on Proofpoint solutions, which are built for today's mobile and social-enabled IT environments and leverage both the power of the cloud and a big-data-driven analytics platform to combat modern advanced threats.

proofpoint.™ www.proofpoint.com