THREATQUOTIENT™

Incident Pruning

Maintaining Control within Incident Response Investigations

Ryan W. Trost, Co-Founder and CTO, ThreatQuotient

Contents

Introduction

Incident response (IR) investigations are complex efforts, shifting between chaos and order, as the incident lead maintains investigation alignment with IR policies, while the team chases down every possible clue, leaving no stone unturned. Without incident pruning (or investigation pruning) investigations can spin out of control within a few minutes simply due to the number of possibilities — associated indicators, adversary aliases, MITRE ATT&CK tactics or techniques, victims, attributes, sightings, etc. This paper discusses some of the strategies to effectively prune an investigation to maintain security operations efficiency and focus.

What is Incident Pruning?

The term “incident pruning” — the process to remove “dead end” investigation paths during an incident that have been deemed benign, irrelevant, or out-of-scope — can elicit many questions. Do all investigations require incident pruning? Are there best practices wrapped around incident pruning? Does a team’s size, maturity and/or skill set dictate incident pruning? Can automation (e.g., SOAR) play a role within incident pruning or is the process strictly manual? And, do certain stages of an IR investigation use and benefit differently from incident pruning?

“Incident pruning” is the process to remove “dead end” investigation paths during an incident that have been deemed benign, irrelevant, or out-of-scope

The definition of incident pruning *itself* can raise questions. Is it truly best practice to remove investigation paths, or is it better to track and archive the information as part of the checklist of investigation actions to avoid duplicating efforts in the future? The answer depends on job roles. Dedicated IR teams find value in maintaining and de-prioritizing investigation data (incident thinning) so that it is available for future reference. However, security organizations without a dedicated IR team tend to find it more useful to delete data that isn’t immediately relevant to help security analysts maintain focus on higher risk efforts (incident deadheading). Recognizing the need for both approaches, this paper discusses incident thinning and incident deadheading.

How ThreatQ Investigations Applies to Incident Pruning

ThreatQ Investigations (TQI) can handle both incident thinning and incident deadheading methodologies. TQI has two features specifically designed to support both incident pruning approaches, including individual vs. shared investigation visibility and exploratory data points.

Incident Deadheading — Private vs. Shared Investigations

Investigations (or incidents) can be private or shared. Private investigations are not visible to other ThreatQ contributors, whereas shared investigations are accessible by any person with a ThreatQ account. ***Each analyst viewing an investigation has an independent view of the investigation which allows the capability to explore endless investigation paths without interfering with other analysts’ views of the investigation.*** This is key to incident pruning because an IR team lead (or incident authority) can delegate components of an investigation to individuals or smaller teams and as tasks are accomplished those results can be published to the global investigation. This is a cornerstone aspect of TQI because it allows the investigation lead to dictate a validation workflow or process before posting content to the larger team. This also helps ensure analysts are focused on their task at hand and not constantly distracted by other analysts' exploratory research paths.

This capability supports incident deadheading because the IR team lead can remove nodes from the incident, either when tasks return negative findings or when the investigation follows a different path and the node is no longer viable. This ensures complex investigations remain digestible by the team and focus is maintained on higher fidelity efforts and thought processes.

To illustrate [Figure 1], when an object is globally viewable in the investigation, the object has a 'solid fill color' (e.g., McAfee ESM - ThreatQ - MD5 Sighting). When the object is still isolated to a specific investigation, it has a 'translucent fill color' (e.g., McAfee ESM Sighting Investigation Report, Event ID, Match Field, and Average Severity) as demonstrated in the image below. To publish an investigation path to the larger group, an analyst must right click and select "Add to Investigation" which will propagate to all views of the investigation.

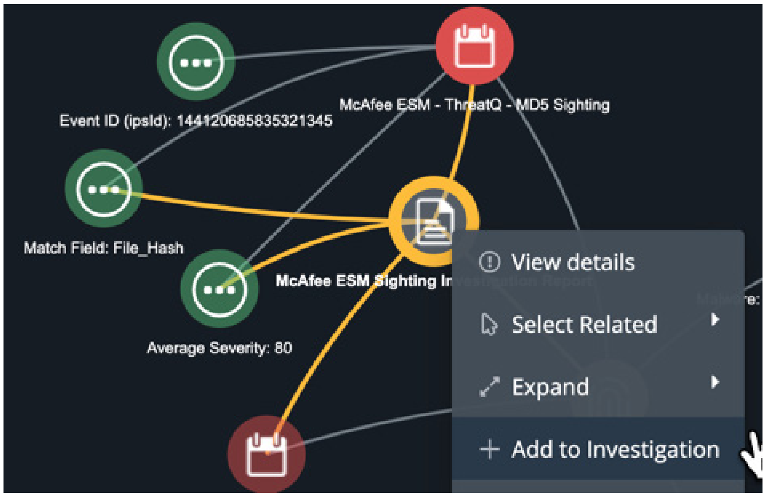

Figure 1: Global versus investigation specific object representation.

Incident Thinning — Exploratory Data Points

Incident thinning is the ability to declutter an investigation and return to a pre-approved state, similar to reverting back to a previous host snapshot. When "show exploratory data points" is checked it allows the team to document investigation paths in all directions as they think outside the box and selectively add only the relevant paths to the global [shared] investigation highlighted in Figure 2.

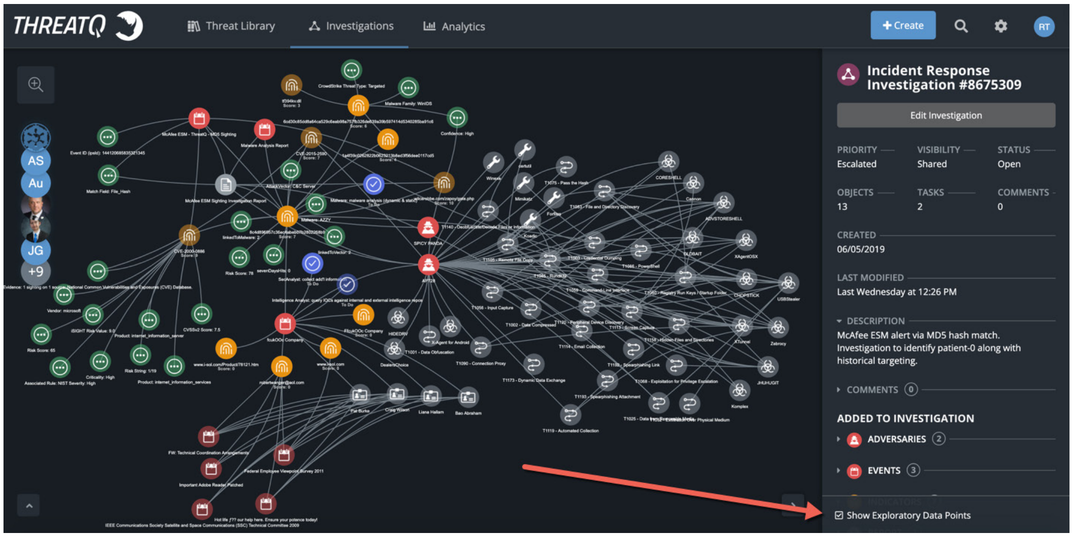

Incident Response Investigation #8675309

Edit Investigation

PRIORITY Escalated

VISIBILITY Shared

STATUS Open

OBJECTS 13

TASKS 2

COMMENTS 0

CREATED 06/05/2019

LAST MODIFIED Last Wednesday at 12:26 PM

DESCRIPTION
McAfee ESM alert via MD5 hash match. Investigation to identify patient-0 along with historical targeting.

COMMENTS 0

ADDED TO INVESTIGATION

ADVERSARIES 2

EVENTS 3

☑ Show Exploratory Data Points

Figure 2: Exploratory data points allows a team to explore various investigation paths.

After the team enumerates through the options and pinpoints the correct path, they can unselect "show exploratory data points" [Figure 3] and pivot back to only the core relevant pieces of the investigation. This is extremely beneficial when dealing with high volume datasets (i.e., MITRE ATT&CK™ having 155+ possible techniques, an abundance of indicators (direct and indirect), fast flux DNS, adversary aliases, potentially numerous malicious emails and/or victims, etc.).

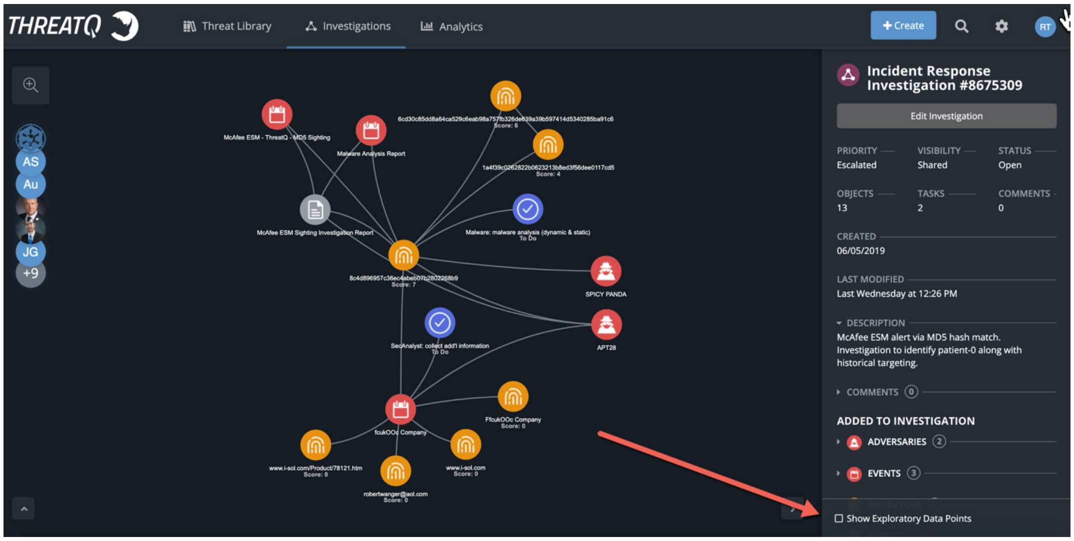
THREATQ
Threat Library
Investigations
Analytics
+ Create
RT
AS
Au
JG
+9
McAfee ESM - ThreatQ - MD5 Sighting
Malware Analysis Report
6cd30c85dd8a64ca529c6eab98a757fb326de639a39b597414d5340285ba91c6
Score: 6
1a4f39c0262822b0623213b8ed3f56dee0117cd5
Score: 4
McAfee ESM Sighting Investigation Report
Malware: malware analysis (dynamic & static)
To Do
8c4d896957c36ec4abeb07b2802268b9
Score: 7
SPICY PANDA
SecAnalyst: collect add'l information
To Do
APT28
FfcukOOc Company
Score: 0
fcukOOc Company
www.i-sol.com/Product/78121.htm
Score: 0
robertwanger@aol.com
Score: 0
www.i-sol.com
Score: 0

Incident Response Investigation #8675309

Edit Investigation

PRIORITY Escalated

VISIBILITY Shared

STATUS Open

OBJECTS 13

TASKS 2

COMMENTS 0

CREATED 06/05/2019

LAST MODIFIED Last Wednesday at 12:26 PM

DESCRIPTION
McAfee ESM alert via MD5 hash match. Investigation to identify patient-0 along with historical targeting.

COMMENTS 0

ADDED TO INVESTIGATION

ADVERSARIES 2

EVENTS 3

☐ Show Exploratory Data Points

Figure 3: Unchecking the exploratory data points reverts the investigation back to a pre-approved state.

Use Case — The Initial Steps of an Incident

In this use case, an incident response investigation is triggered from a SIEM alert. The premise of this use case isn't to provide step-by-step incident response instructions nor provide never before seen indicators, adversaries or tactics, but rather highlight several key features within

TQI that allow an IR team lead to properly prune an incident based on threat prioritization and resource responsibility. The investigation does contain several characteristics that will require incident pruning — both deleting and de-prioritization.

To set the stage, the investigation will be escalated to an incident which will extend visibility to fellow SOC teams encompassing security analyst, incident response, malware analysis, threat intelligence and vulnerability assessment teams.

Each team is responsible for a specific component of the investigation. As information is collected and validated it is "added" to the bigger investigation. As with any security team dichotomy/taxonomy, responsibility bleed over across tasks varies organization to organization depending on skill set, time, etc. The role/task delineation in this use case is just one example, but the IR investigation premise will remain the same across organizations. Also note there are numerous ways to approach an investigation so the sequence of events may also differ depending on the organization and situation.

Step #1: Security Analysts triage the SIEM alert

The security analyst will triage the alert to create an investigation [Figure 4] within TQI containing the SIEM alert and any immediate information gleaned from the alert. At this point, the investigation is just beginning and this visualization serves as the incident foundation.

It quickly surfaces that the alert was triggered by a spearphish attempt via a malicious attachment. The email in question is tracked down and added to the investigation which includes the MD5 hash, email subject and email sender.

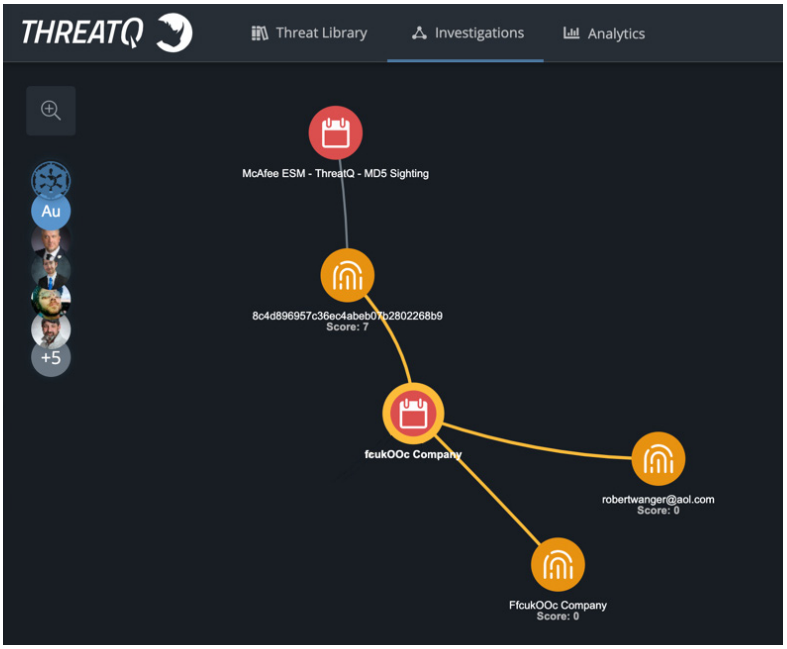

Figure 4: The SIEM sighting alert.

Step #2: Malware analyst performs analysis on MD5 hash

The malware analyst submits the corresponding MD5 hash file to an internal sandbox. Having successfully detonated the specimen, the malware analyst publishes the malware report and the command and control FQDN and URL to the larger global investigation [Figure 5]. The analyst also queries the hash in external malware repositories (i.e., VirusTotal) to glean any additional information or community commentary. Note: Most teams are very cautious NOT to use VirusTotal as the sandbox itself because submitted malware samples are published to the general public which likely includes the attackers who are monitoring VirusTotal for their malware.

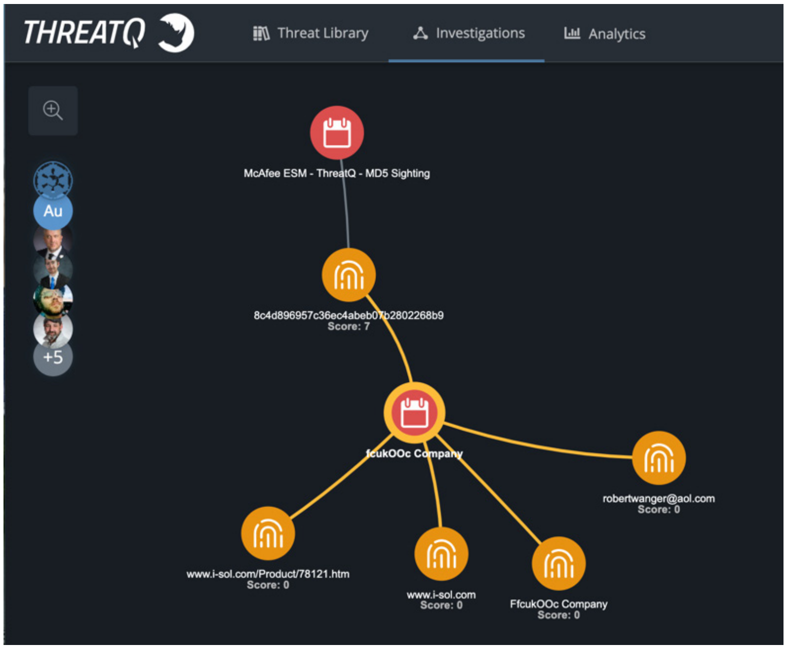

Figure 5: The results of the malware analysis report are added to the investigation.

Step #3: Query internal and external intelligence repositories

Since the email and malware analysis surfaced several indicators of compromise (IOCs) the security analyst and/or intelligence analyst queries the IOCs against internal and external intelligence repositories to gain any additional insights to assist with the investigation. Internal repositories include the ticketing system, historical SIEM alerts and threat intelligence platforms. External intelligence repositories include the commercial intelligence providers (i.e., FireEye Intelligence, CrowdStrike, Flashpoint, Digital Shadows, Intel-471, Recorded Future, etc.) or third-party query services (i.e., PassiveTotal, DomainTools, FarSight Security, etc.).

The queries identify three external intelligence repositories that provide some helpful information. Some information is a bit too detailed to publish to the global investigation, but provides good situational awareness for certain teams. For example, one of the commercial providers attributes the hash to an adversary, APT28 (aka Spicy Panda) so that creates a new path for validation and exploration for the intelligence team.

Up until this point, the investigation has been limited to one security analyst, one incident response lead and one malware analyst, each with a specific role and a small handful of IOCs within the investigation. As such, the investigation has not warranted incident pruning using either methodology. But as more information is gathered and the intelligence and vulnerability team get involved, the need for incident pruning is inevitable.

Step #4: Threat intelligence analyst digs into email sender

The intelligence analyst focuses on gathering additional intelligence on the attacker's email address (robertwanger [at] aol [dot] com). Typically, burner addresses don't yield highly actionable results, but sometimes it does help to identify previous attacks or registered domains which might help with existing IR efforts.

In this situation, chasing the attacker's email address identified four previous attacks. This initiates a deeper dive into attack vector patternization and targeting insights.

Step #5: Threat intelligence analyst performs victimology

By discovering the previous email attacks, the intelligence analyst can begin to perform some quick victimology analysis to try to determine "why" those individuals were targeted. Did the attacker focus on soft targets to increase their likelihood of success before turning their attention to the hard targets — the end goal? Or did they directly attack their final targets?

Including previous spearphish attacks and respective targets within the global investigation is valuable, but pulling in all associated information for those attacks can quickly consume the link analysis graphic. This is why it is so critical to selectively publish information to the global investigation. The intelligence analysts will want to keep it visible within their view but the IR team lead might not publish it to the global view [Figure 6].

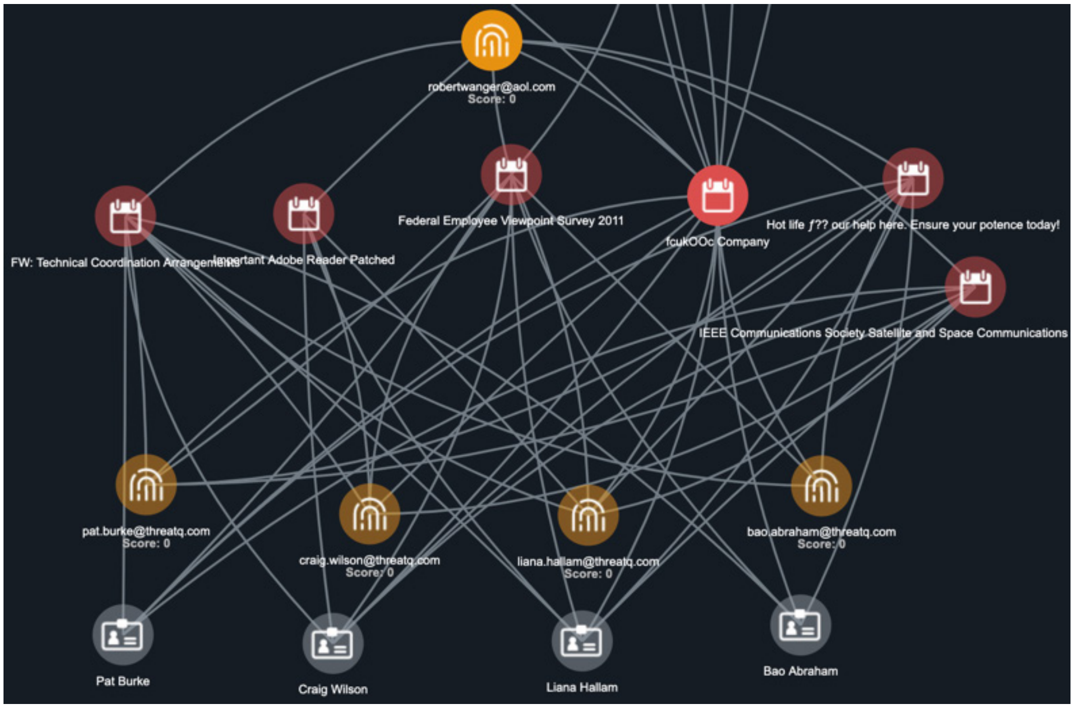
robertwanger@aol.com
Score: 0
FW: Technical Coordination Arrangements
Important Adobe Reader Patched
Federal Employee Viewpoint Survey 2011
fcukOOc Company
Hot life ƒ?? our help here. Ensure your potence today!
IEEE Communications Society Satellite and Space Communications
pat.burke@threatq.com
Score: 0
craig.wilson@threatq.com
Score: 0
liana.hallam@threatq.com
Score: 0
bao.abraham@threatq.com
Score: 0
Pat Burke
Craig Wilson
Liana Hallam
Bao Abraham

Figure 6: Alert in question is compared to previous targeted attacks to aid intelligence analysis.

This is the first incident pruning effort, specifically incident thinning, where the information exceeds what is desirable to include in the global investigation. The initial information is helpful for the intelligence team but not relevant enough (at this point) to publish to the larger team. If additional information about one of the previous attacks or targets is discovered that will help the current investigation, then the pertinent subset of information should be published to the global team.

Step #6: Threat intelligence analyst confirmed adversary attribution

In step #3, two premium intelligence providers have associated an adversary group, APT28 and SpicyPanda, to the hash value. Though this intelligence is slightly more trustworthy given the source is a paid intelligence feed, every piece of external intelligence should be validated to some degree. The intelligence team works to verify the information, comparing the externally provided adversary profile information to the internal investigation information to identify similarities.

This step includes incident pruning in either form depending on whether the attribution was already added to the investigation for global visibility. If the adversary attribution was included in the investigation and the team cannot verify its accuracy, the team can remove [deadhead] it from the investigation. However, if the intelligence team did not include it in the global investigation already, then incident thinning will deprioritize that piece of intelligence and it will not be broadcast to the larger investigation team.

Step #7: Incident Response analyst performs 'rear-view mirror search'

In an effort to cast an "APT28 indicator-centric" net to try and identify either a pattern of activity or potentially try to identify any internally initiated egress communication to APT28's malicious infrastructure, the intelligence analyst displays all the indicators associated with the adversary and runs the respective operation. Since throughout the lifecycle of the adversary ~1,995 indicators are associated to the adversary group, incident pruning efforts are required to maintain focus [Figure 7].

If the analyst displays all those indicators, the graphic will quickly become unusable. Instead, the analyst can filter through the list and include the indicators more relevant to the investigation, whether designated by doppelganger syntax, threat score, or chronologically, by clicking on the 'eye-ball' immediately before the IOC to review the IOC's summary.

Including the relevant indicators in the local investigation allows for the ability to trigger an Operation to query them against internal logs. If a fruitful result occurs it can be published to the global investigation for larger visibility. This is an example of incident thinning — selectively determining which objects/nodes to publish globally and which ones have false negative results, and therefore, are not relevant to the investigation and are not published globally.

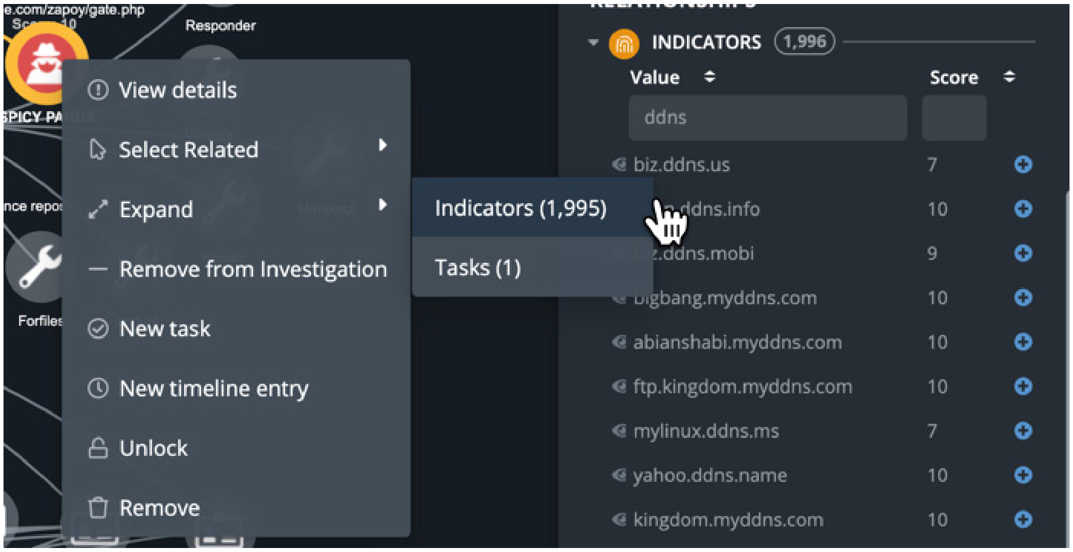

Value	Score
ddns	
biz.ddns.us	7
ddns.info	10
z.ddns.mobi	9
igbang.myddns.com	10
abianshabi.myddns.com	10
ftp.kingdom.myddns.com	10
mylinux.ddns.ms	7
yahoo.ddns.name	10
kingdom.myddns.com	10

Figure 7: A quick reference point for the number of indicators associated with the adversary group Spicy Panda.

Step #8: Threat intelligence analyst and IR analyst map adversary to MITRE ATT&CK (software, tool, TTP) to help uncover possible investigation paths

Although adversary attribution itself is not immediately actionable per se, it does help the team focus on the respective TTPs to ensure a unified defense. This step is a joint action between the intelligence analysts and IR team to help prioritize investigation paths. In this situation [Figure 8], APT28 has the following characteristics associated to them via MITRE ATT&CK:

- 50 TTPs [attack patterns]
- 14 malware software
- 6 malicious tools

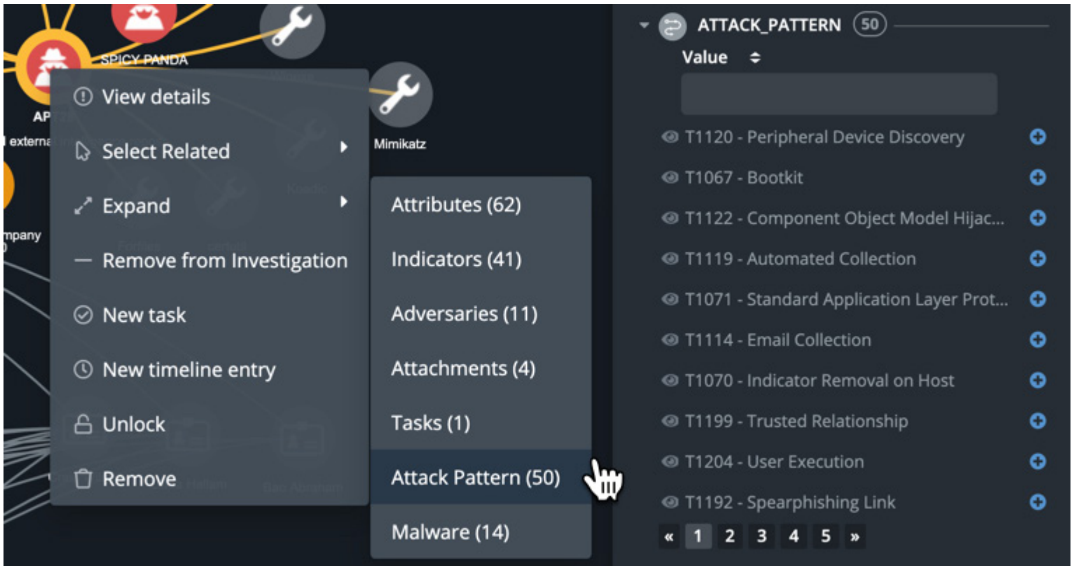

T1120 - Peripheral Device Discovery
T1067 - Bootkit
T1122 - Component Object Model Hijac...
T1119 - Automated Collection
T1071 - Standard Application Layer Prot...
T1114 - Email Collection
T1070 - Indicator Removal on Host
T1199 - Trusted Relationship
T1204 - User Execution
T1192 - Spearphishing Link

Figure 8: Selectively deciding which Attack Patterns to include in the investigation.

Fifty additional objects on the link analysis graph is busy but do-able; especially if only for an individual or team view versus publishing it to the global investigation for the whole team. But in which direction should the team allocate manpower? Once all three categories [TTPs, software, tools] are added to the investigation and a cursory review eliminates some of the "noise," it creates a smaller, prioritized list of investigation tasks. This is where the "gut feeling" of the intelligence team and IR team comes into play. If the adversary has gravitated to a certain style recently, that will be prioritized at the top. Otherwise, a "paper weight" litmus test can be used to try to identify which combination of TTPs, malware and tools have connecting edges. In a "paper weight" test, the analysts evaluate each node across the three categories to determine which nodes have the most correlations (e.g., connecting edges). This approach is not scientific but indirectly speaks to the nature of probabilities.

For instance, using the methodology to compare edges of Zebrocy and USBStealer finds that Zebrocy has six edge associations [Figure 9]:

- 1 association to APT28
- 5 associations to APT28's Attack Patterns

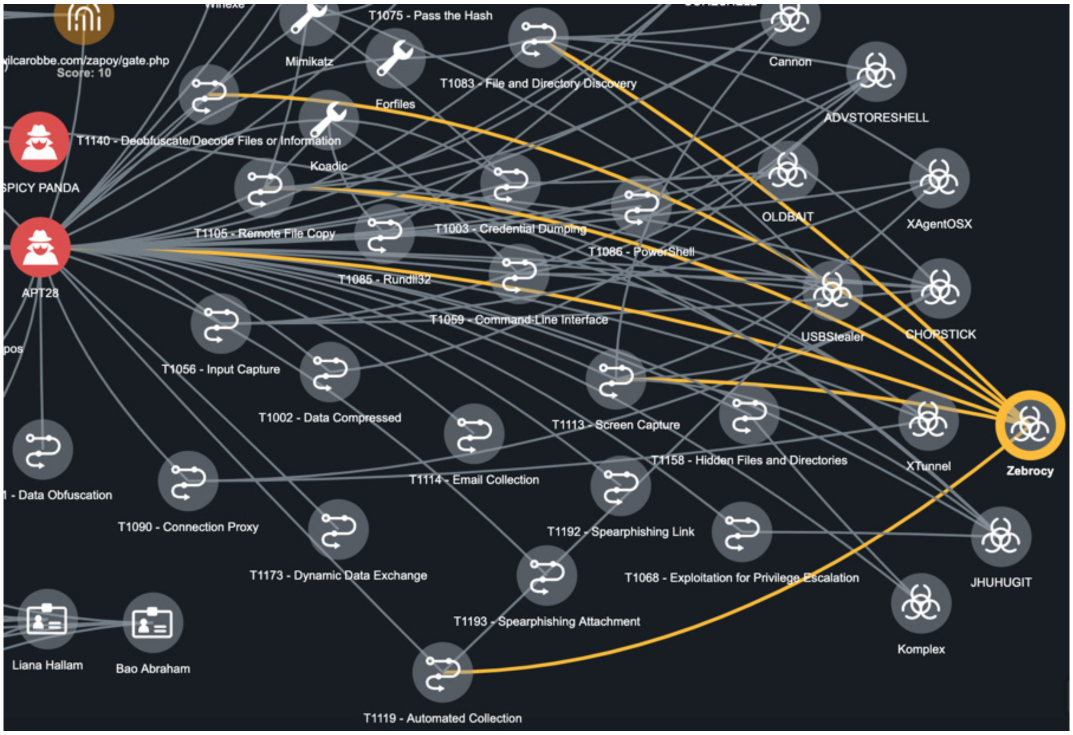

Figure 9: The attack pattern characteristics directly associated with Zebrocy malware.

However, selecting USBStealer reveals nine edge associations [Figure 10]:

- 1 association to APT28
- 1 association to MD5 hash (from SIEM alert)
- 1 association to SHA-256 hash
- 1 association to CVE-2015-2590
- 5 associations to APT28's Attack Pattern

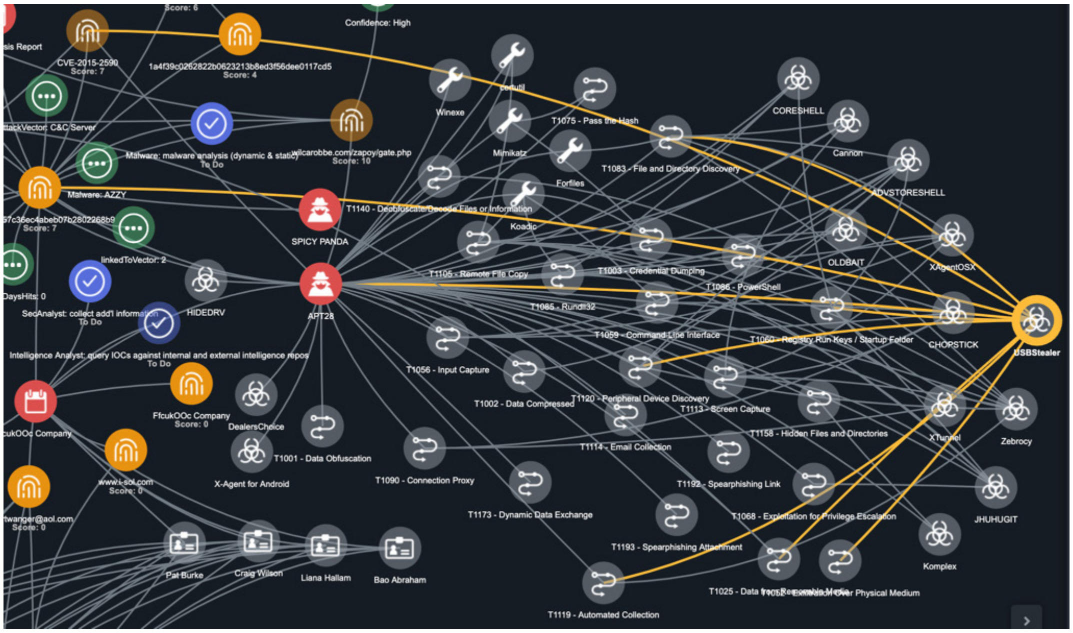
Score: 6
Confidence: High
sis Report
CVE-2015-2590
Score: 7
1a4f39c0262822b0623213b8ed3f56dee0117cd5
Score: 4
ttackVector: C&C Server
Malware: malware analysis (dynamic & static)
To Do
wilcarobbe.com/zapoy/gate.php
Score: 10
Malware: AZZY
57c36ec4abeb07b2802268b9
Score: 7
linkedToVector: 2
DaysHits: 0
SecAnalyst: collect add'l information
To Do
HIDEDRV
SPICY PANDA
APT28
Intelligence Analyst: query IOCs against internal and external intelligence repos
To Do
FfcukOOc Company
Score: 0
cukOOc Company
DealersChoice
X-Agent for Android
T1001 - Data Obfuscation
www.i-sol.com
Score: 0
twanger@aol.com
Score: 0
Pat Burke
Craig Wilson
Liana Hallam
Bao Abraham
Winexe
certutil
Mimikatz
Forfiles
Koadic
T1075 - Pass the Hash
T1083 - File and Directory Discovery
T1140 - Deobfuscate/Decode Files or Information
T1105 - Remote File Copy
T1003 - Credential Dumping
T1086 - PowerShell
T1085 - Rundll32
T1059 - Command-Line Interface
T1060 - Registry Run Keys / Startup Folder
T1056 - Input Capture
T1002 - Data Compressed
T1120 - Peripheral Device Discovery
T1113 - Screen Capture
T1158 - Hidden Files and Directories
T1114 - Email Collection
T1090 - Connection Proxy
T1173 - Dynamic Data Exchange
T1192 - Spearphishing Link
T1068 - Exploitation for Privilege Escalation
T1193 - Spearphishing Attachment
T1025 - Data from Removable Media
T1052 - Exfiltration Over Physical Medium
T1119 - Automated Collection
CORESHELL
Cannon
ADVSTORESHELL
OLDBAIT
XAgentOSX
CHOPSTICK
USBStealer
XTunnel
Zebrocy
JHUHUGIT
Komplex

Figure 10: The attack pattern characteristics directly associated with USBStealer malware.

The linkage to the hash values immediately prompts the team to focus on USBStealer as the malware software used in the attack. This creates a highly probable investigation path for the team to chase. At this point, the intelligence analyst and incident response team agree on the effort and globally publish the information to the investigation which immediately propagates to other analysts working on the investigation — primarily the malware team for verification.

Although it is still very early in the investigation of a relatively simple and straightforward spearphish investigation, analysts already have relied heavily on incident pruning, demonstrating that even simple investigations can gain complexity quickly. That complexity is not purely based on the attack itself, but also the supplemental metadata each investigation requires to ensure every stone has been uncovered. Incident pruning is a critical process for every IR team, whether through incident thinning to leverage the individual view of the investigation and only publish the most relevant information to the larger global investigation view, or through incident deadheading where the team removes false negatives or irrelevant information from the investigation in order to maintain a narrower investigation scope.

Conclusion

Security teams continuously look for ways to mature their process and improve their IR efforts. Incident pruning should be one of the first activities to consider, however, it is commonly overlooked. Often, when an event is escalated to an incident it immediately attracts a greater set of eyes among fellow teams and that invites an increase of possible tasks, ideas, and a seemingly endless number of foxholes to investigate. This requires an IR team lead, or someone with the 'bigger picture' authority, to deadhead the investigation (top down) to avoid inefficiency and confusion or perhaps an analyst to thin the investigation (bottom up) by only globally publishing relevant data points to the other teams. Without incident pruning (or investigation pruning) investigations can spin out of control within a few minutes simply due to the number of possibilities — associated indicators, adversary aliases, MITRE ATT&CK tactics or techniques, victims, attributes, sightings, and more. As this paper shows, ThreatQ Investigations can support both incident pruning approaches, deadheading and thinning.

Incident pruning should be one of the first activities Security teams consider to mature their process and improve their IR efforts.

Frequently Asked Questions

Do you think an investigation needs to reach a certain size or complexity in order to require incident pruning — whether incident thinning or incident deadheading? If so, how can it be quantified? # of analysts? # of investigation possibilities? # of victims/targets? Lateral movement?
Absolutely, however, the juxtaposition of "when" incident pruning is necessary is going to subtly vary from team to team and incident to incident. Analysts with several years of experience often have the ability to track several simple incident paths without having to rely on incident pruning, but that quickly spirals out of control once the investigation reaches a certain complexity.

There are several investigation characteristics that will initiate the need for incident pruning, including:

- Increasing the number of analysts beyond four or five
- Increasing the involvement of more than three roles
- Increasing the involvement of two or more analysts within the same roles (symptomatic of a larger or more complex investigation)
- Transitioning from simplistic, fact-finding efforts to more complex attack methods that extend beyond five-eight investigation paths (each path requiring its own micro investigation, elaboration and collaboration)
- When the attack extends beyond 10 infected victims requiring individual dissection of attack
- When the attack reaches a point of lateral movement across five hosts (either directly from patient-0 or sequentially)

The above are general rules of thumb to draw a line in the proverbial sand to ensure investigations run efficiently, but exceptions always exist at both ends of the spectrum. Some teams will initiate incident pruning much later in the investigation, whereas other teams will initiate incident pruning techniques almost immediately when an investigation is created.

Should every possibility within an investigation have some sort of “likelihood or probability score” to better rank exploratory possibilities?

Maybe not “every” possibility but absolutely it should be pretty close. Teams can perform this more explicitly by literally assigning a score to each node/object. Alternatively, teams can do this more subconsciously where possibilities are prioritized in more of a “waterfall approach” with the highest priority items that pose the largest threat and likelihood to happen appear at the top, and decrement down to the end of the action list. The waterfall approach caters to a link-analysis type investigation methodology because it is visually driven.

Does orchestration/playbook automation play a role within incident pruning efforts?

The human element will always remain vital in security operations but automation has a place. With respect to incident pruning, automation logic can be applied to perform incident pruning based on:

- A low threat or probability score provided by an analyst
- A stale node/object based on the lack of activity taken against an item over a prolonged period of time
- A lack of associated characteristics strongly suggests the item is a conceptual stretch within the investigation and is irrelevant to the case

Are there different stages of an investigation that require more pruning than others?

IR workflows vary from team to team so a good benchmark to use is the SANS Incident Response process which includes *Preparation, Identification, Containment, Eradication, Recovery, Lessons Learned*.

The SANS’ Incident Response checklist provides a checklist summary regarding each IR stage. Below is a description of the checklist phases along with how incident pruning effects each stage.

Preparation	As expected, this stage revolves around ensuring the team is ready for incident response actions — ownership and responsibilities for all systems, communication channels are agreed upon, the team understands the structure of an attack, standard operating procedures (SOP) are defined and agreed upon across, etc.
	The Preparation stage does not lend itself to incident pruning. Preparation is more of an anticipatory checklist for the team and outside of minor periodic adjustments remains static.

Identification

This stage encompasses the discovery of the intrusion, whether generated internally or externally, as well as, the research and categorization of the attack to determine the depth and breadth of the intrusion. The identification stage includes a significant amount of the work due to the aggregation of information gathered and analyzed, and numerous attack possibilities. This includes matching the information gathered to the MITRE ATT&CK framework to better understand the attack landscape and possibly learn which adversary is behind the attack. This stage also concludes the Mean-Time-to-Detect (MTTD).

The Identification stage does lend itself to incident pruning, whether through incident thinning or deadheading, because in this stage a significant amount of the work requires brainstorming and exploration of possibilities. To minimize investigation clutter and duplicating efforts, teams need to consciously sanitize and maintain a clean investigation to ensure effective and efficient defenses. The team's dichotomy/taxonomy will determine the degree of incident pruning necessary, but this stage will inevitably entail the most incident pruning across the SANS IR stages.

Containment

This stage is a concerted effort to stop further entrenchment. However, this stage is controversial because some teams believe the priority should be to isolate and remove the threat immediately, whereas others believe it is better to monitor the adversary prior to "pulling the ripcord" to study the adversary's movements and logic and to identify any "sleeper cells" hidden amongst the environment. Both are reasonable paths for an IR Team to pursue and are dictated largely by the skillset of the team and sophistication of the tools to help "safely" isolate the adversary to negate any additional harm. For example, the ability to re-route them to a controlled environment or deception technology (i.e,. Attivo Networks).

The Containment stage is an actionable phase where the security team is implementing steps to minimize additional adversary entrenchment. This stage includes updating signatures, deploying IOCs, correlating additional logs into a SIEM, pushing infected hosts into a controlled environment for monitoring, etc. The Containment stage is more of a checklist of counter measures to enable/deploy which includes incident pruning methodologies, specifically incident deadheading to eliminate counter measures that are not effective.

Eradication

This stage focuses on the removal of all malicious software, implants, remote access tunnels (RATs), internal command and control infrastructure, or backdoors from the environment. This is probably the hardest stage because adversaries can camouflage themselves extremely well.

The Eradication stage starts by having a list of all the infected hosts and all the malicious software that needs to be removed. This phase caters greatly to a checklist effort, and therefore, incident deadheading. However, rather than deadheading in the sense of deleting, leverage the Shared vs. Private feature to highlight what tasks need to be accomplished. In this case, which hosts are still infected and which ones have been returned to a safe end state.

Recovery

This stage focuses on restoring all business functions, from bringing servers back online to regaining employee productivity, and includes implementing any necessary defensive measure to ensure future attacks will not be successful. The Recovery stage also concludes the Mean-Time-to-Respond (MTTR) stopwatch.

Restoring business function, though critical, is not likely going to be enumerated within TQI. However, tracking the implementation of the necessary defensive measures will align with TQI and the efforts of the larger security team. In this situation, most teams will review the overall attack (i.e., align it to an attack framework like MITRE ATT&CK or the Cyber Kill Chain) to try to identify multiple defensive countermeasures for each attack phase. Typically, the countermeasures are not new technology deployments but rather correcting a misconfiguration, updating the proper signatures, funneling respective logs into the SIEM for correlation, or adjusting an internal workflow to ensure an alert receives attention significantly faster. Similar to the Eradication stage, this type of checklist effort, caters greatly to a deadheading approach.

Lessons Learned

This stage is unequivocally the most commonly overlooked process — primarily because either the existing daily tasks have piled up and the team is playing catchup, or the team is inexperienced and doesn't perform an after action debrief.

The Lessons Learned stage is crucial, but most of the heavy-lifting effort has already been performed so there is no need to prune any of the data. The focus of this stage is to summarize the incident and determine if the counter measures in place will be effective in the long run.

ThreatQuotient's mission is to improve the efficiency and effectiveness of security operations through a threat-centric platform. By integrating an organization's existing processes and technologies into a single security architecture, ThreatQuotient accelerates and simplifies investigations and collaboration within and across teams and tools. Through automation, prioritization and visualization, ThreatQuotient's solutions reduce noise and highlight top priority threats to provide greater focus and decision support for limited resources. ThreatQuotient is headquartered in Northern Virginia with international operations based out of Europe and APAC. For more information, visit **www.threatquotient.com**.