BLACK HILLS
INFORMATION SECURITY

THE INFOSEC SURVIVAL GUIDE

ORANGE BOOK

INCIDENT RESPONSE

Scan and sign up for FREE cybersecurity labs, webcasts, and swag!

CREDITS

Made by and for the community!

and our team at BHIS

Writers

Tom DeJong

Ronald A Mays Jr

Brian Baskin ------ @bbaskin

Patterson Cake ------ @SecureCake

Hayden Covington ------ @KilobyteTheDust

Erik Goldoff ------ @ErikG

Gerard Johansen ------ @IRProactive

Fletus Poston III ------ @fletusposton

Andrew Scott ------ @PurpleTeamLead

John Strand ------ @strandjs

Wade Wells ------ @WadingThruLogs

Technical Editors

Dave Blandford
Death Metal #1 Fan

Dale Hobbs
Does Not Look Good in a Tutu

Ashley Knowles
Principal Cheese Connoisseur & TTP Googler

Adam Rose
C[redacted]st
Is that a black highlighter?

Alyssa Snow
Shortest Security Consultant

Kaitlyn Wimberley
Second Shortest Security Consultant

Troy Wojewoda
Pencil Tester (Ya Know, Becuz of Mistakes N Stuff)

They're all security consultants... who apparently moonlight as comedians.

PROMPT# Crew

John Strand ------ That Nice Tall Man With the Cowboy Hat *per a very nice member of our community*

Jason Blanchard ------ Really Wants a Documentary About Himself for Some Reason *(currently interviewing himself for it)*

Deb Wigley ------ Weirdly Obsessed With 3D Printing *3D printed another printer*

Dani Diem ------ Draws Things? *for a living?*

Kassie Kimball ------ Just Really Nitpicky *those sneaky----spaces*

Look for all of us at cons and meetups!

Table of Contents

Brought to you by:

BLACK HILLS | Information Security

ANTISYPHON TRAINING — antisyphontraining.com

wildwesthackinfest.com

activecountermeasures.com

rekcahcomics.com

PROMPT# — promptzine.com

bhis.co

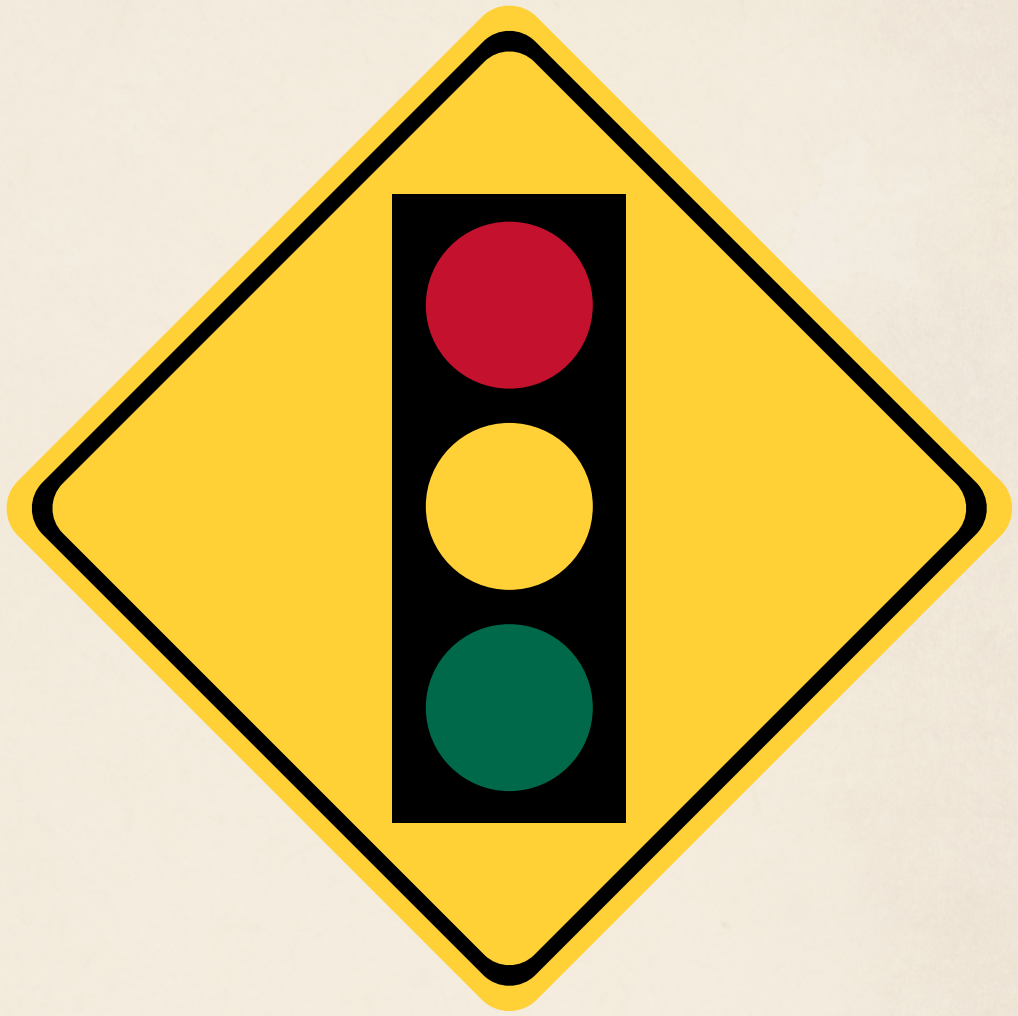

STOP

THINK

SOLVE

THE INCIDENT RESPONDER MANIFESTO

We are incident responders.

By nature, we are curious. We are problem solvers. We thrive on a good riddle. With our expertise and passion to uncover the truth, we hunt down the adversary—tracing where they've been and what they've done in order to stop them in their tracks and ensure that they never return to cause further harm.

Incident response is about answering the most critical questions a business can face:

Have we been compromised?

How and when did it happen?

How do we minimize the impact on business operations?

How can we prevent it from happening again?

Time is money. We need simple, effective, and repeatable processes that deliver quick and accurate answers. When we get it right, the business is positioned for expedient and successful recovery. If we get it wrong, the cycle of compromise repeats.

We serve and protect.

We are incident responders,
and we are dedicated to getting it right.

"GOLDILOCKS" ALERT REVIEW

Is It a True Positive?

written by Hayden Covington || linkedin.com/in/hayden-covington/

Security engineers, analysts, and incident responders all have one thing in common, and I'm not talking about 3 AM phone calls concerning incidents. I'm talking about triage: that challenging moment of urgency when assessments must be made and classifications communicated—for the right things to be decided on to prevent the bad stuff from happening.

We're all petrified about missing a critical event or misclassifying an alert, but when we're talking about incident response (IR), there are often hundreds if not thousands of alerts to parse through. It's easy to get caught up with one alert because it feels "too hot" or maybe not spend enough time looking into something that initially seems "too cold." I'll provide some tips, tricks, and techniques to help find that "Goldilocks Zone" of spending just the right amount of time on an alert, allowing you to quickly triage and move on to the next.

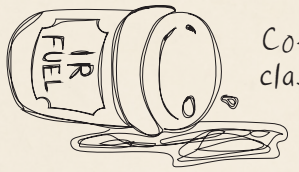

Coffee first, classify later.

Consider the Severities

The simplest way to triage events is to consider the severity of the alert. Initially, most Low-severity alerts should be entirely ignored. On any given case, time is limited, and the value return of going over hundreds of Low alerts is not even remotely comparable to that of reviewing higher priority findings.

Mediums will often be the bulk of your alert volume, with these being right in that uncomfortable middle of the "Probably nothing" of Lows and the "Definitely something" of Highs or Criticals. To move quickly on an IR engagement, I categorically relegate the Medium alerts to a later time on that case; and almost every time, the High and Critical alerts tell the real story, giving concrete direction on how to search the Mediums and Lows in a more targeted fashion.

Anomalies Against the Baseline

One of the most surefire ways to quickly classify an event as a true or false positive is to compare the activity against the normal baseline: **"Does this happen regularly on this host," "in this environment," or maybe even "in any of the environments I can observe?"**

While a certain execution or activity on one host may appear anomalous, once you discover that it happens on a number of hosts across multiple environments, either you've just found your answer, or you've discovered a much bigger issue... *(Hint: it's usually the former.)*

Actions on Objective

One of my favorite tactics is considering "actions on objective." If an attacker gains access to a host, they have an end goal in mind. Whether that goal is financially motivated, a desire to steal data, or even if they just want to observe activity in the environment—they broke in for a purpose. Someone isn't going to go through the effort of breaking into your house just to stand around.

For the attacker to achieve their goals, some form of activity must occur to get them there. That often looks like exfiltration or lateral movement. All that to say, if the activity you're looking at never actually goes anywhere, steals anything, or tries to override anything else, chances are it may not actually be evil, because it certainly isn't doing the attacker much good.

Detection Intent

Probably the most novel (and somewhat controversial) approach I've come across is considering detection intent first and foremost. This approach requires you to understand exactly what it is that the rule that fired is looking for, and to look for that one specific thing only, ignoring anything else you may see in the periphery.

That idea may sound counterintuitive but think about it: Threat detections are often straightforward and scoped to a specific TTP, and the biggest time-waster in investigations is when we follow rabbit holes of related activity. So, if instead we observe that the primary activity the alert fired on isn't malicious or relevant, that alert is now classified as non-malicious, and we move on.

While debatable in its effectiveness, you can't argue with its efficiency and straightforwardness. Ultimately, with this concept, you should be considering the alert itself and looking for the activity that most resembles what "success" would look like for the attacker in that context.

Ask Questions

To summarize the core concept of each of those suggestions into one idea, the best way to tell if the alert you're looking at is in the "Goldilocks Zone" is to consider the context around it procedurally:

- How high of a priority is the event?
- How often does this activity happen?
- Does the activity from the alert help the attacker further their goals?
- What would "success" look like for that attack?

If you consider each of those things in succession,
you should be able to breeze through a queue of alerts in no time.

CRITICAL YOUR FIRST-HOUR RESPONSE

It's a True Positive! ... Now What?

written by Fletus Poston III || linkedin.com/in/fletusposton || @fletusposton

You've just confirmed it: that alert isn't a false positive. Your heart rate kicks up a notch—this is the real deal. Whether it's suspicious PowerShell execution, unusual network traffic, or signs of credential compromise, you've got a genuine security incident on your hands. Let's go.

What happens in the next 60 minutes can make or break your entire incident response. Miss a critical step, fail to preserve evidence, or escalate to the wrong people, and you could turn a manageable incident into a career-defining disaster. Tread carefully. THIS IS NOT A DRILL!!!

The First 15 Minutes: Immediate Actions

When you confirm a true positive, your brain wants to jump straight into investigation mode. Resist that urge. Instead, take the following actions right away to lay the foundation for everything that follows:

Preserve Evidence First

Before you click a button or run another query —

Document the current state:

- Screenshot the original alert with full context visible.
- Capture system timestamps (note any time zone differences).
- Record which detection rule or signature triggered.
- Note your analysis tools and queries used for confirmation.

Preserve volatile data:

- If dealing with endpoint alerts, consider memory capture before the system state changes.
- Document running processes, network connections, and logged-on users.
- Save original log entries before anyone accidentally overwrites any evidence while "investigating."

Establish Incident Tracking

Without delay, create your incident ticket with these critical fields:

- **Incident ID:** Use your organization's numbering scheme.
- **Initial Classification:** Malware, data breach, unauthorized access, etc.
- **Affected Assets:** Systems, IP addresses, user accounts.
- **Discovery Time:** When the alert fired vs. when you confirmed it was legitimate.
- **Analyst:** The person taking point on this incident.

> **Pro Tip:**
> Many teams skip this step thinking they'll "create the ticket later." Don't. You need that incident number for all subsequent communications and actions.

Quick Impact Assessment

Spend 5 minutes answering these questions to determine everything else, including your escalation path, response urgency, and resource allocation:

- Is this affecting critical business systems?
- Could this be active lateral movement?
- Is sensitive data potentially at risk?
- Are customers or operations impacted right now?

The Next 45 Minutes: Assessment and Documentation

Now you can dig deeper, but you need to do it systematically. Your goal is to build a complete picture while documenting everything for the team members who'll join your response.

Comprehensive Alert Analysis

Expand your investigation scope:

- Look for related alerts in the same time window (±2 hours).
- Check for indicators across multiple data sources (SIEM, EDR, network monitors).
- Identify the attack timeline:
  initial access → persistence → lateral movement → impact

Document key findings:

- **Initial Vector:** How did this start (phishing, exploit, credential compromise)?
- **Affected Systems:** Complete inventory with hostnames, IPs, and roles.
- **User Accounts:** Any compromised or suspicious account activity.
- **Network Activity:** Unusual connections, data transfers, or command and control.
- **Potential Impact:** What could an attacker do with the current level of access?

Document EVERYTHING.
Ps: This is totally not a threat... maybe.

Communication and Escalation: Who Needs Immediate Notifications? (i.e., within the first 15 minutes)

Always notify the following:

- Direct supervisor/on-call manager.
- Other SOC analysts.
- Network operations (if needed).

Escalate to Security Leadership if:

- Multiple systems compromised.
- Critical business systems affected.
- Ongoing data exfiltration.
- Privileged account compromise.
- Indicators suggest APT.

Escalate to Executive Leadership if:

- Customer data confirmed compromised.
- Business operations significantly impacted.
- Potential regulatory violations.
- Media attention or disclosure is likely.

Involve Legal/ Compliance if:

- PII/PHI compromised.
- Breach notification requirements.
- Law enforcement involvement.
- Regulatory reporting deadlines.

Execution and Ongoing Management

Investigation Coordination

Daily standup structure:

- **Previous 24 hours:** What was discovered or completed?
- **Next 24 hours:** What are the priority investigation activities?
- **Blockers:** What's preventing progress or needs leadership support?
- **Key findings:** Any new IOCs, affected systems, or impact discoveries?

Evidence management:

- Maintain chain of custody for all forensic evidence.
- Use standardized file naming conventions.
- Store evidence in secure, access-controlled locations.
- Document who accessed what evidence and when.

Recovery Planning in Parallel with Investigation

- **System restoration:** Clean rebuilds vs. remediation approaches.
- **Access restoration:** New credentials, certificates, and authentication tokens.
- **Business process restoration:** How to resume normal operations safely.
- **Monitoring enhancement:** Additional detection capabilities to prevent recurrence.

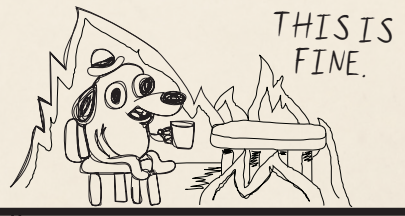

Leadership and Stakeholder Updates

- **Security leadership:** Daily — focusing on scope, impact, and containment effectiveness.
- **Executive leadership:** When status changes significantly or every 24–48 hours.
- **Business units:** When their operations or data might be affected.
- **Legal/compliance:** Regularly — on potential regulatory implications.

Leadership and Execution Resource Requirements

- **Dedicated incident response platform:** Don't rely on email chains and spreadsheets.
- **Forensic tool access:** Memory capture, disk imaging, and network analysis capabilities.
- **Communication tools:** Secure channels for sensitive incident coordination.
- **Documentation templates:** Standardized formats for consistent incident reporting.

Team Training Priorities

- **Incident classification:** How to quickly assess severity and impact.
- **Evidence preservation:** Proper forensic procedures that hold up under scrutiny.
- **Communication protocols:** Who to notify, when, and how.
- **Tool proficiency:** Regular hands-on training with investigation tools like DeepBlueCLI.

Final Thoughts

The moment you confirm a true positive, you're not just investigating an incident—you're protecting your organization's future. The actions you take in that critical first hour determine whether this becomes a contained incident or a business-threatening breach.

Key Principles to Remember

- **Preserve first, investigate second:** You can't go back and collect evidence you've overwritten.
- **Communicate early and often:** Stakeholders need information to make good decisions.
- **Document everything:** Your future self (and legal team) will thank you.
- **Contain quickly, but thoughtfully:** Rushed containment can destroy evidence or cause business disruption.
- **Measure success by these four things:** Time to containment, communication effectiveness, evidence quality, and learning integration.

Your organization is counting on you to get this right. With a reliable process, tools like DeepBlueCLI, and a systematic approach, you'll be ready when that next alert turns out to be the real deal.

Next Steps

1. **Review your current true positive workflow:** Do analysts know exactly what to do in the first 15 minutes?
2. **Test your escalation procedures:** When was the last time you exercised your communication plans?
3. **Audit your investigation tools:** Do you have the capability to investigate modern threats effectively?
4. **Train on DeepBlueCLI:** If you're not using this for Windows event analysis, start now.

Essential Reading for Report Writing

- **NIST SP 800-61:** Computer Security Incident Handling Guide (Section 3.4 on documentation).
- **CERT Resilience Engineering Framework:** Structured approaches to incident documentation.
- **Your organization's legal requirements:** Know what documentation standards apply to your industry.

Just some light reading...

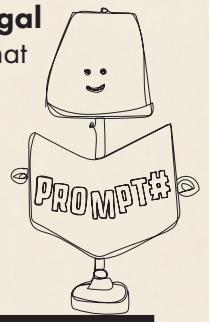

Tool Resources

PowerShell for Security Professionals
Advanced techniques for incident response.

Windows Event Log Analysis
Understanding what logs contain the evidence you need.

Learn more about **DeepBlueCLI** on **page 28**!

Looking for more practical cybersecurity content and incident response guidance? Check out additional resources and real-world scenarios on my YouTube channel at **youtube.com/fletusposton.**

KAPE 101

https://www.kroll.com/en/services/cyber/incident-response-recovery/kroll-artifact-parser-and-extractor-kape

written by Gerard Johansen || @IRProactive

Spend time performing forensic analysis on the Windows Operating System and you'll see a host of artifacts that can be used to identify adversary activity. From changes to the registry to the System Resource Utilization Monitor, Windows artifacts run deep. The challenge is locating, extracting, and parsing these artifacts in an efficient manner.

This is where the Kroll Artifact Parser and Extractor (KAPE) comes into play. KAPE gives analysts and incident responders the capability to collect specific artifacts and parse them into an easily analyzable format. Available free to individuals working in their own environment, KAPE is capable of handling a wide range of artifact extraction and parsing for faster, more efficient analysis.

The Artifact Extractor

The Artifact Extractor allows analysts to gather key artifacts related to adversary activity. In this example, our analyst will leverage KAPE to extract a Basic Triage package that contains a variety of artifacts. This is useful in a couple of different cases. For example, KAPE is leveraged by analysts who usually have local access to a system and want to extract key artifacts as part of an initial analysis. In other circumstances, KAPE can be run against a mounted disk image or a virtual disk, which allows the analyst to work with the file system. If it can be mounted and given a drive letter, KAPE can be used.

Pro Tip:
It is tempting to focus on individual artifacts, such as extracting the Windows Prefetch files to identify suspicious executions. Instead, focus on extracting a collection of artifacts that you can continually use as the investigation progresses to avoid repeatedly going back to extract more data.

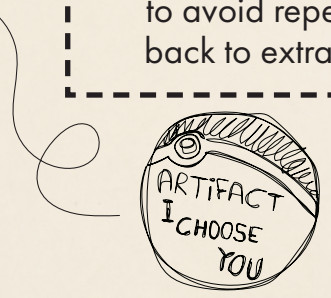

Why grab one artifact when you can catch 'em all?

Extracting Artifacts to a Virtual Disk

In the following case, the analyst is using KAPE to extract key artifacts and aggregate them into a virtual disk. This approach extracts key artifacts for follow-on analysis but also places them into a container that is easy to copy and share:

1. In the top left of the KAPE GUI, select **Use Target options.**
2. Complete the Target options by selecting the source. Select the **C:** directory. **Note:** KAPE will only extract the artifacts that are selected.
   a. For the Target destination, select an applicable directory. **Note:** The Flush option will delete all files in the directory before the output is processed. Keep that in mind.
3. In this example, use the **VHDX** option for the container.
   a. A good practice is to use the system name for the **Base name** field.
4. Once these are completed, click **Execute.**

You will be left with the system VHDX file along with an activity log.

The Parser

Next is the Parser. The Parser places the input into a format that can be manually analyzed or incorporated into an analysis platform.

Pro Tip:
You can parse multiple artifacts at the same time, which saves trips back to the GUI. If triaging a system, you will want to look at file data, such as the Master File Table and USN Journal. For potential binary execution, the Prefetch, Shimcache, and Amcache may be useful. A review of the Windows Event Logs can provide a comprehensive overview of system activity. Save time by selecting a group of files to parse at once for the best overview of a system.

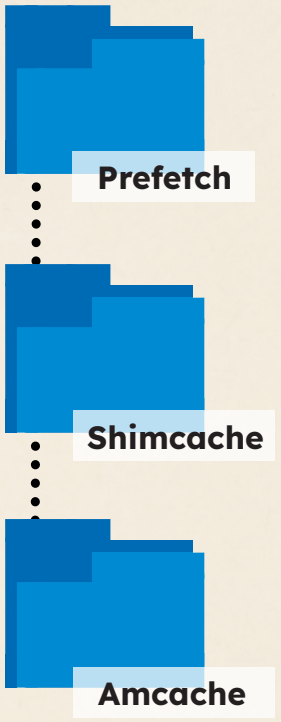

Parsing Key Artifacts for Analysis

In this third case, the analyst will use the triage package and parse out the Master File Table, Event Logs, and the USN Journal.

1. Double-click the VHDX file. This will mount the virtual disk.
2. Click on **Use Module** options.
3. For **Module source**, select the root of the drive letter associated with the VHDX file.
4. For **Module destination**, select an appropriate directory to output the files.
5. Select a few Targets such as the **Event Logs, Master File Table,** and the **USN Journal.**
6. Export format and select **CSV.**
7. Click **Execute.**

A quick check of the destination directory should have three sub-directories: Filesystem, Eventlogs, and Consolelog. The files within are ready for analysis.

Pro Tip:
When looking at the Export format for the parsed results, consider whether you will be using specific tools for the analysis. For example, there are modules that will output a JSON format that are useful for importing into SOF-ELK. To start off, using the default or the CSV output is easiest to work with.

A SIMPLE, USEFUL IR PLAN

That Actually Works

written by Patterson Cake || @SecureCake || @secure-cake

Incident response (IR) is no time for subtlety, ambiguity, or unnecessary complexity. Yet most IR plans are too long, minimally applicable, outdated, and ineffective for their intended purpose. An active threat demands a practical response plan that is succinct and broadly applicable, serving as a clear blueprint for emergency roles, responsibilities, communication, and authority.

Drafting Your Plan

As we draft our IR plan, remember that invoking the plan denotes an emergency! Aim to keep it under 10 pages and focus on the most critical phases of incident response: identification (finding all unauthorized access and activity), containment (stopping that access and activity), and eradication (preventing it from happening again). If we succeed in these phases, we will be poised for successful recovery.

IDENTIFICATION
↓
CONTAINMENT
↓
ERADICATION

"What about ____?" While most businesses share key similarities, each environment has unique needs. If you find yourself needing to add "____" due to your business model, just keep it simple, relevant, and current.

I'm a huge fan of modularity in an IR plan. Start with a core plan based on the simple outline below. Then expand it with appendices and playbooks tailored to your environment.

So, what should be in your IR plan?

Roles & Responsibilities

This is not the time to describe everyone's normal job duties. What we need is: "In case of emergency, this role has this responsibility." Consider the following emergency response functional role assignments:

- IR Commander (leads the response team)
- Technical Response Team (accesses and utilizes IT and cybersecurity resources)
- Scribe (documents EVERYTHING)
- Internal Communications Lead (liaises between engaged and informed groups)
- External Communications Lead (liaises with external resources)
- HR & Legal Counsel (serves internal and/or external needs)

Contact Info

Include primary and secondary contacts, along with full contact information for each role, including any relevant third parties. Review this at least quarterly to make sure it remains up to date. In addition to the previously mentioned roles and responsibilities, consider the following contact categories:

- Cybersecurity Insurance
- DFIR Retainer
- Law Enforcement
- Cybersecurity Vendor Support
- Functional Business Leadership

Escalations & Emergency Powers

Don't overcomplicate this. When confronting an active threat, expect disruptions to business IT systems and services, both from threat actor behaviors as well as IR containment and eradication procedures. Assign and empower your emergency response team members to act decisively in defense of the business!

Miscellaneous

There is no "miscellaneous" section in your IR plan! If you feel the need to include definitions (like "ransomware") or to outline the PICERL (prepare, identify, contain, eradicate, recover, lessons learned) process phases, consider IR-plan appendices or a separate "explanatory" resource (e.g. "Understanding Our IR Plan").

Playbooks

Resist the temptation to create a playbook for every "Threat of the Day/Week/Month." Remember our primary goal: to create an IR plan that is succinct, widely applicable, and serves as a clear blueprint for emergency roles, responsibilities, communication, and authority. And remember: **the IR plan should not stand alone!** It is just one critical element of your overall disaster recovery planning, integrating with your IT asset inventory (you can't protect or restore what you don't track), your asset criticality matrix (identifying which IT assets are most important to operations, driving recovery effort prioritization), and your business continuity plan (the process for evaluation and restoring affected IT assets).

Check out **"Containment and Eradication"** on **page 26** for further steps on preparing for an incident.

Find an example plan here:
https://git.new/simple-ir-plan

INCIDENT INVESTIGATIONS

Digging Deeper

written by Patterson Cake || @SecureCake || @secure-cake

If unauthorized access or actions occurred in your environment, your technical controls failed. Our priority at this point is to determine the "what"—what unauthorized access or actions occurred and what are we going to do about it?

Spoiler: The answer is not asking your "friend" Chad G. Peetee.

We need "actionable intelligence" to inform our next steps. Generally, this means identifying indicators of compromise (IoCs) and using those indicators to bolster our technical controls via emergent changes, such as manual blocklist entries in your AV/EDR, port/URL/IP blocking at your firewall or secure proxy, correcting misconfigurations, etc.

Finding Indicators of Compromise

Where do we find IoCs? That depends on your current visibility capabilities, the tools you have to interrogate your environment—for example, a security information and event management (SIEM) platform, endpoint detection and response (EDR) solution, or user entity behavior analytics (UEBA). You may need additional tools but start with what you already have in place.

With rare exception, everything we care about in our environments occurs on an endpoint, in the context of an identity. That's where we'll start.

Keep in "MIND"

Let's break down the components of a typical endpoint into four categories: (M)emory, (I)dentity, (N)etwork, and (D)isk (MIND). These are general areas we want to examine via our visibility capabilities. There are typically hundreds of thousands of forensic artifacts on an endpoint; most are irrelevant to our investigation. Let's focus on "attack surface" categories, where a threat actor and/or malware is most likely to impact our endpoints.

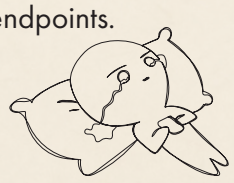

Not me crying myself to sleep and dreaming in artifacts.

Understanding Attack Surface

Understanding "attack surface" helps us pinpoint the forensic artifacts we care about most, prioritize acquisition and analysis, and expedite the discovery of actionable intelligence!

While malware and threat actor (TA) behaviors constantly evolve, the primary components of our endpoint operating systems (OS) remain largely constant. By understanding the attack surface for our primary OS's and mapping those to available forensic artifacts, we can consistently review these for indicators of unauthorized activity. Instead of changing our approach based on the latest malware/TA behaviors, we can monitor the impact on endpoints and simply evolve our attack surface investigative workflow.

Windows and Linux are the two primary OS's we'll encounter in our investigations. Although the forensic artifacts are different, the attack surface categories are the same:

- Running processes
- Identity (user/service accounts)
- Network (sockets, pipes)
- Tasks (scheduled actions)
- File system
- Settings/config
- Services
- Startup environment

Windows Attack Surface

To analyze the attack surface for Windows, prioritize acquisition and analysis of these artifacts:

- Tasklist (running processes w/command line)
- Windows Event Logs (EVTX)
- Master File Tables (MFT)
- Netstat (with user/process/parent process)
- ASEPs (auto-start extensibility points)
- DNS cache
- Scheduled tasks
- Artifacts of Execution (Prefetch, Shimcache, Amcache)
- Web actions/history

Linux Attack Surface

To analyze the attack surface for Linux, prioritize acquisition and analysis of these artifacts:

- Running processes w/command line
- Crontab
- Systemd/init.d
- SSH login logs
- Network sockets
- Shell history (bash, sh)
- Executable file listing w/hash

Just a busy little penguin... silently judging your logs.

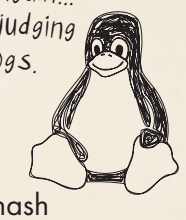

For a step-by-step workflow on acquiring, parsing, and analyzing these artifacts, check out "Rapid Endpoint Investigations":

https://github.com/secure-cake/rapid-endpoint-investigations

Your Investigation Workflow

Step 1: Don't freak out. Step 2: Freak out, but methodically.

So, we've had a "true positive" security event and need to understand what happened and what to do next (actionable intelligence). We want to prioritize endpoint analysis, narrowing our investigation to attack surface categories based on OS and starting with existing visibility capabilities to acquire and review targeted artifacts, augmenting with other tools if necessary as we search for IoCs.

Although we've been selective about the artifacts we're acquiring and analyzing, thousands remain to review. Fortunately, we have clues based on the event that instigated our investigation.

This "point of impact" guides our investigative workflow: date and time, endpoints, identities, etc. Focus on date/timestamp of confirmed impact then slowly expand your temporal analysis (forward/backward in time), looking for "attack extents" (the beginning and end of unauthorized access/actions).

As you find abnormalities, compare to "known good" endpoints and external resources (threat-intel data/libraries), and document potential IoCs. Validate these by extending your visibility-capability review to all related endpoints, confirming "normal vs abnormal."

Upon confirming a high-confidence IoC, implement changes to bolster your technical controls (AV/EDR blocklists, port/URL/IP blocklists, endpoint isolation, account session revocation, etc.). Continue this until you've identified attack extents on all impacted endpoints.

Once we've identified all unauthorized access and actions, and blocked or remediated these throughout our environment, we are poised to move into the "recovery" phase, working towards a return to business as usual.

HAYABUSA 101

https://github.com/Yamato-Security/hayabusa

written by Patterson Cake || @SecureCake || @secure-cake

Windows Event Logs (EVTX files) are a forensic goldmine, but they are voluminous and only a small percentage of events provide value during security investigations.

Enter Hayabusa, an open-source "Windows event log fast forensics timeline generator and threat hunting tool" by Yamato Security. We use this tool frequently for Windows endpoint investigations to filter out EVTX noise, consolidate events we care about, and prioritize our analysis.

Getting Started

To get started, download the latest stable Hayabusa release to your analysis system, acquire EVTX logs from your Windows target system, and execute a couple of simple commands.

Since Hayabusa utilizes over 4,000 Sigma rules and over 170 built-in detection rules, the first thing we want to do is update rules with the following command:

```
hayabusa.exe update-rules
```

Understanding Command Options

Next, let's review detailed help for the "csv-timeline" option:

```
hayabusa.exe csv-timeline --help
```

As you'll see, there are a lot of options for fine-tuning Hayabusa behavior and output. The general command syntax looks like this:

```
hayabusa.exe [COMMAND] <INPUT> [OPTIONS]
```

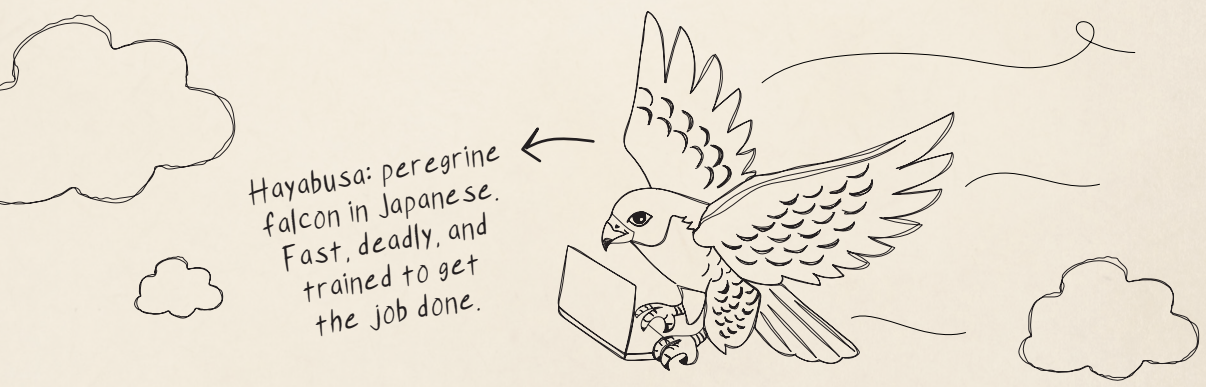

Creating a CSV Timeline

We'll create a CSV timeline for quick review in a CSV viewer/editor (Timeline Explorer, Excel, OpenOffice Calc, etc.) with the following command:

```
hayabusa.exe csv-timeline -d d:\case123\host-abc\
evtx-logs -w -o d:\cases\case123\host-abc\hayabusa-output-host-abc.csv
```

The command switches we used are as follows:

-d = specifies a directory, followed by a path to where we've stored our EVTX files

-w = "no wizard," just scan our EVTX files for all events and alerts

-o = specifies the directory and name for the output file

What we just instructed Hayabusa to do is scan all of the EVTX files in our specified folder, use all Sigma and built-in detection rules, look for rule "hits," and then write that output in CSV format, including timestamp, rule title, computer name, event channel, event ID, and event details to our designated folder/file.

Now you can review the CSV output file in whatever viewer/editor you prefer. A great place to start is sorting based on timestamp, then by rule "level" (critical, high, medium, etc.).

Further Learning

This is the tip of the proverbial Hayabusa iceberg. Check out these additional resources to learn more:

"Wrangling Windows Event Logs with Hayabusa and SOF-ELK – Part 1" blog:

https://www.blackhillsinfosec.com/wrangling-windows-event-logs-with-hayabusa-sof-elk-part-1/

"Wrangling Windows Event Logs with Hayabusa and SOF-ELK – Part 2" blog:

https://www.blackhillsinfosec.com/wrangling-windows-event-logs-with-hayabusa-sof-elk-part-2/

Welcome to the deep end. The water's cold, but the data's hot. (Yes, I do comedy.)

KNOW YOUR ENEMY

Threat Actor Standard Operating Procedure

written by Wade Wells || linkedin.com/in/wadingthrulogs || simplydefensive.transistor.fm

Threat actors. Individuals, groups, or organizations conducting malicious activity in cyberspace. Their motivations span profit, power, ideology, or amusement. Methods range from commodity phishing to zero-day exploits, from data theft to power grid disruption.

What unites them is intent: to compromise systems, networks, or people for advantage at someone else's expense. Tracking these threat actors isn't about knowing every alias or log. It's about sharpening defense, prioritizing resources, and lowering risk.

Hey now! So it's not paranoia if they're really after you.

Four Primary Groups

Thrill Seekers and Trolls
Curiosity, clout, and chaos motivate this group. Skills range widely, but these aren't the "script kiddies" of yesteryear anymore.

Nation-State
Their goals are espionage, influence, disruption, or destructive outcomes. They invest in long-term access, disciplined tradecraft, and region or sector-specific victimology. You will see living off the land, supply chain compromises, and careful timing around geopolitical events. They have unlimited resources and endless patience.

Ideologues (Including Hacktivists and Terrorists)
They aim for narrative impact as much as network access. Expect defacements, data leaks, doxing, distributed denial of service (DDoS), and action against symbolic targets or political moments. Targeting logic follows headlines and ideology more than profits.

Financially Motivated
Think ransomware operators, initial access brokers, fraud rings, cryptojackers, and data brokers. They do it for money. Their strategy is simple: minimize cost, maximize payout.

Keep in Mind

These groups do not stay in a lane; they mix and match all the time. Ransomware groups work with governments. Nation-states hack for ideological reasons. Even trolls look to make money with stolen data.

Not the singing, glittery, great-hair kind.

Profiling

Threat profiling is the structured practice of describing and assessing a threat actor so defenders can prepare accordingly. It answers who is behind the activity, what they want, and how they operate. Profiles provide technical indicators of tools, infrastructure, and malware families with softer attributes like motivation, targeting logic, and patterns of behavior.

If you know that a ransomware crew specializes in brute forcing exposed RDP servers, you can prioritize monitoring for those attempts. If hacktivists are active around a political event, you can prepare for DDoS attempts. Profiles also help analysts understand escalation paths—whether a financial group might pivot into data theft, or whether a nation-state campaign might lie dormant until triggered. Profiling transforms threat data into operational foresight.

How to Profile

Creating profiles and prioritizing threat actors often requires assigning measurable scores. Andy Piazza's "Threat Box" model rates actors based on Intent and Capability, along with modifiers of Willingness and Novelty.

- **Intent** reflects why an actor might target an organization, whether due to sector alignment, ideological motivations, or opportunistic targeting.
  - ◊ **Willingness** adjusts that **intent** based on external constraints—economic ties, diplomatic relationships, etc.
- **Capability** evaluates the actor's proven operational strength.
  - ◊ **Novelty** boosts scores for rare or sophisticated tools or TTPs.

Tidal Cyber's Profiling Toolkit and Curated-Intel's guide also provides frameworks to assess operational impact, regional activity, and intelligence confidence in models.

By blending structured scoring with threat profiles, organizations can make decisions based on data and invest in defenses where risk is highest.

Resources on Threat Actors & Profiling

Wiz Threat Actor Info
https://threats.wiz.io/all-actors
Malpedia
https://malpedia.caad.fkie.fraunhofer.de/actors
APT Groups and Operations
https://apt.threattracking.com
ATT&CK Groups
https://attack.mitre.org/groups/
Threat Box by Andy Piazza
https://klrgrz.medium.com/quantifying-threat-actors-with-threat-box-e6b641109b11
Curated Intel
https://github.com/curated-intel/Threat-Actor-Profile-Guide
Tidal Cyber
https://github.com/tidalcyber/cyber-threat-profiling

Are you dating a state-sponsored threat actor? Check the ANTISOC #PROMPT issue. You're welcome.

Why Tracking Matters

Tracking threat actors isn't abstract; it's aligning your defenses to reality. Every organization faces too many alerts with too few resources, making prioritization essential. By knowing who is most likely to target you, you can tune defenses around threats that actually matter.

Strategically

- **Prioritization of Risk:** Understanding which actors are relevant to your industry or region.
- **Resource Allocation:** Informing investment decisions.
- **Threat Forecasting:** Anticipating future campaigns and shaping strategic resilience.

Operational

- **Incident Response Playbooks:** Giving teams likely behaviors, infrastructure, and kill chain steps to anticipate during intrusions.
- **Threat Hunting:** Guiding efforts by TTPs unique to high-priority adversaries.

Tactical

- **Detection Engineering:** Converting specific techniques and indicators (e.g. MITRE ATT&CK® mappings, malware families, command-line patterns) into detections.
- **Control Validation:** Testing if current defenses actually catch known adversary tools.

MITRE ATT&CK® FRAMEWORK:

Know Your Enemy, Know Your Gaps

written by Fletus Poston III || linkedin.com/in/fletusposton || @fletusposton

In cybersecurity, knowing your environment's blind spots is critical. MITRE ATT&CK® gives you the framework to do exactly that.

What Is MITRE ATT&CK®?

MITRE ATT&CK® is a knowledge base of adversary tactics, techniques, and procedures (TTPs) based on real-world observations—a comprehensive playbook of how attackers operate, organized into 14 tactics from Initial Access to Impact, with hundreds of specific techniques underneath each.

ATT&CK documents what security teams have actually seen in the wild—from APT groups using WMI for lateral movement to ransomware operators abusing PowerShell.

Why It's Your Gap Analysis Goldmine

ATT&CK reveals where your defenses have holes. Map your current security tools against ATT&CK techniques and you'll quickly spot concerning patterns.

Common visibility gaps include:

- **T1055 (Process Injection):** Do you detect when malware hollows out legitimate processes?
- **T1021 (Remote Services):** Can you spot abnormal RDP or WMI usage patterns?
- **T1003 (OS Credential Dumping):** Would you catch someone running Mimikatz or similar tools?

For each technique, ask yourself: "Could an attacker do this without triggering alerts?" If the answer is "maybe" or "I don't know," you've found a gap.

Making ATT&CK Work for Incident Response

During active incidents, ATT&CK becomes your investigation roadmap. When you detect suspicious PowerShell execution (T1059.001), you know to look for what came before (like phishing [T1566]), what comes next (like credential dumping [T1003] or lateral movement [T1021]), and related techniques (like other PowerShell-based activities). This systematic approach prevents tunnel vision and helps you understand the full attack chain, not just individual alerts.

Pro Tip:
Create detection coverage maps showing which ATT&CK techniques your tools can identify. Update these quarterly as you add new capabilities or as attack methods evolve.

The Bottom Line

ATT&CK isn't just a reference guide—it's your strategic advantage. By mapping real adversary behaviors against your actual defensive capabilities, you transform theoretical security into practical protection. Know where you're blind, and you can start seeing clearly.

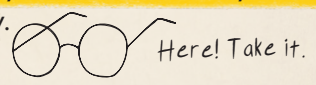

Check out additional resources and real-world scenarios on my YouTube channel at youtube.com/fletusposton.

COMMON IR FINDINGS:

As Told by Backdoors & Breaches

written by Patterson Cake || @SecureCake || @secure-cake

Some of the most valuable incident response outcomes are the "lessons learned."

Unfortunately, lessons learned from a cybersecurity incident can be very, very costly. Our motivation to create Backdoors & Breaches (B&B) was to facilitate proactive "lessons learned" conversations through fun and engaging tabletop exercises. We have run hundreds of B&B exercises with our community. Here are some of the most common "lessons learned" outcomes.

These are all B+B cards btw!

Crisis Management

When walking through cybersecurity scenarios, the transition from "alert" or "event" to "incident" often sneaks up on us. This transition is a critical opportunity to invoke incident-response protocols and marshal additional resources.

Action Item

Develop, document, and socialize crisis-management invocation criteria. Develop incident response protocols that are accessible, practical, and useful.

Tools:
Incident Response Plans (IRP)
Disaster Recovery Plans (DRP)

Isolation

Though most defenders can quarantine compromised systems, the criteria and authority for containment are often unclear!

Action Item

Develop and document isolation criteria that includes potential business impact based on systems, scope, and authorization authority.

Tools:
Host Firewall
Switch and Router Commands
Endpoint Detection and Response (EDR)

User and Entity Behavior Analytics (UEBA)

With rare exception, all cybersecurity incidents impact endpoints and identity. Whether using a discrete UEBA solution or investigating via an identity source like Microsoft Active Directory, we use this card frequently. Be ready!

Tools:
LogonTracer
DeepBlueCLI / OpenUBA
Hayabusa

Action Item

Develop, document, and test procedures for investigating user access and authentication throughout your environment.

Endpoint Analysis

If you're experiencing a cybersecurity incident, your technical controls failed! We're not blaming the controls—just acknowledging that something went wrong, something we'll need to address to contain and eradicate the threat. If your Endpoint Security Protection solution fails to give you needed visibility, what's your "Plan B?"

Tools:
DeepBlueCLI
Velociraptor / Incident Response Cheat Sheets
osquery

Action Item

Develop, document, and test technical procedures for endpoint analysis and forensic-artifact acquisition that complement but don't depend on your endpoint security protection solution.

BACKDOORS & BREACHES

How to Play and Where to Get Started

an Incident Response card game created by BHIS

Backdoors & Breaches is a cooperative, cybersecurity threat emulation game in which "Defenders" will work together to uncover the attack pathways used to hack into their environment. Taking the concept of traditional tabletop exercises, Backdoors & Breaches combines the structure of a card game with the flair of classic role-playing games to help organizations and individuals learn about the tactics, methods, and tools used in cyber attacks and defense.

Contents

Among the 52 unique playing cards in your Backdoors & Breaches: Core Deck, you will find:

You Will Also Need

- A crew of 2 or more (ideal number of players is 5-7)
- A d20 (20-sided die) OR a virtual dice-rolling app
- A healthy dose of imagination!

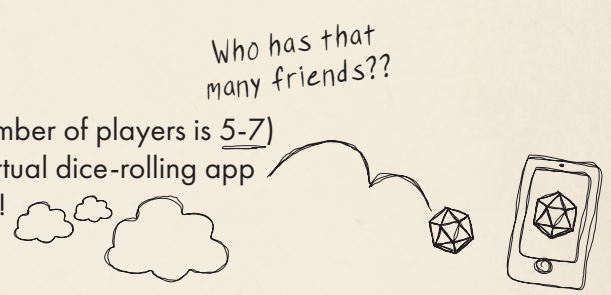

Getting Started

Overview

Using a secret array of 4 Attack cards, the "Incident Captain" will craft an imagined security breach and guide the "Defenders" through the scenario. Equipped with critical thinking, dice, and DETECTIONS, the Defenders will attempt to discover what the attackers are doing before it's too late! The gameplay of Backdoors & Breaches is cooperative. You either win as a team, or you lose as a team.

Objective

To win, the Defenders must reveal all 4 Attack cards before 10 turns have passed. Otherwise, they have failed to uncover the various avenues of the attack, and they lose.

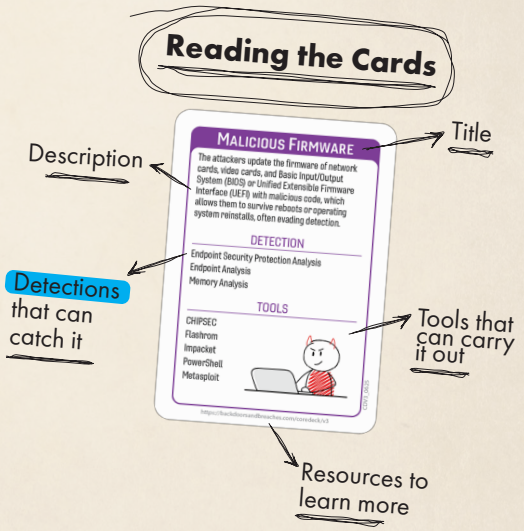

Determining Roles

Before you start, you must determine roles for each player: Incident Captain or Defender.

Choose 1 person to serve as the Incident Captain. This person will be responsible for crafting the starting scenario, answering questions, improvising situations, and is overall in charge of guiding the game process. Whoever you choose should have a wide breadth of cybersecurity knowledge and be a quick thinker. All other players will serve as Defenders. They form the team responding to the incident at hand.

INCIDENT CAPTAIN
1 CREATIVE LEADER

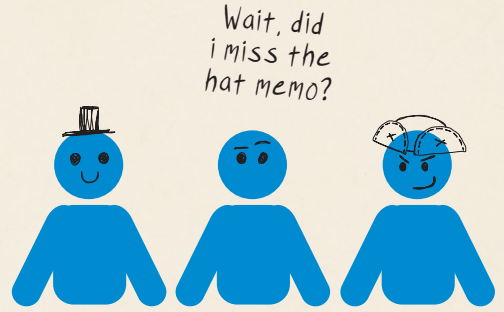

DEFENDERS
2+ ENTHUSIASTIC PLAYERS

Incident Captain Setup — Attacks

The Incident Captain chooses 1 card from each Attack card pile (INITIAL COMPROMISE, PIVOT and ESCALATE, C2 and EXFIL, PERSISTENCE) and keeps those cards hidden from the Defenders! Once the Incident Captain has all 4 Attack cards, you will not need the rest of the Attack card piles for the remainder of the game.

Defenders Setup - Detections

You will now deal the DETECTION cards into 2 rows: Established Detections and Other Detections. For Established Detections, deal 4 random cards face up. For Other Detections, place all remaining DETECTION cards face up in a row beneath.

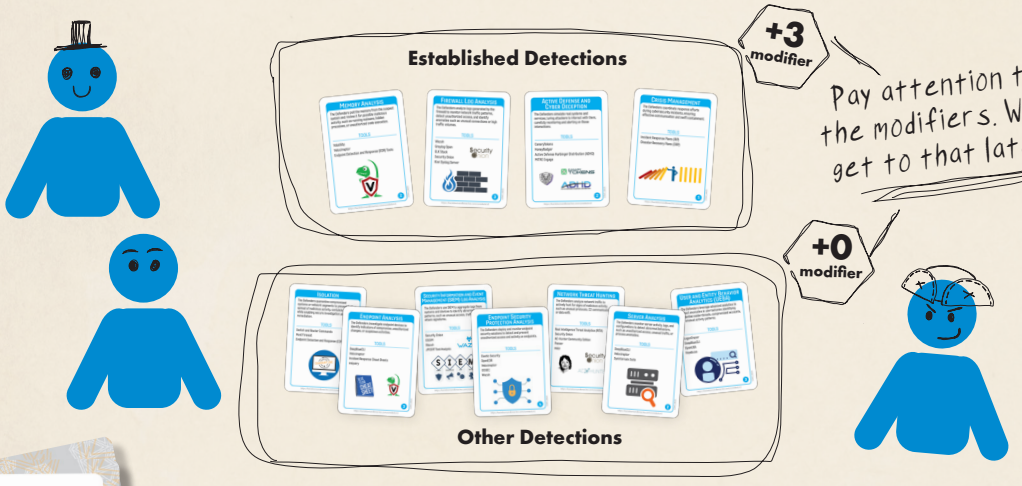

Injects Setup

Place the INJECTS pile to the side of your play area, face down.

Playing The Game

To begin, the Incident Captain must set the stage by crafting a breach scenario based on the 4 Attack cards. This should be detailed enough to give the Defenders a place to start, without giving away the specifics of any Attack cards.

[Incident Captain Tip: It is usually easiest to build the scenario from the INITIAL COMPROMISE card.]

Sequence of Play

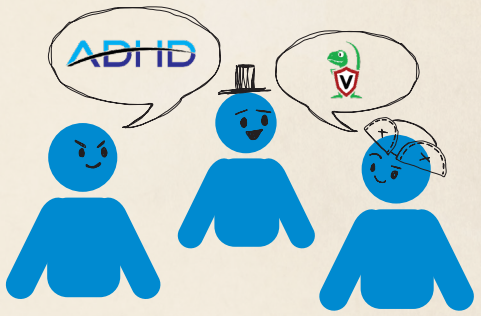

1. Discussion

The Defenders should discuss the current situation amongst themselves and decide which of the DETECTIONS they should attempt to use.

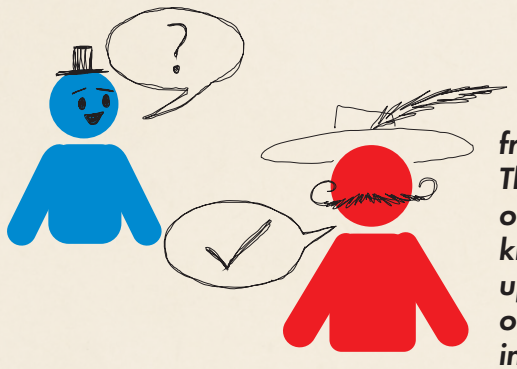

[Defenders Tip: The Defenders can seek clarity from the Incident Captain during this phase. They may ask the Incident Captain to expand on details that would make sense for them to know. This does not require any dice rolls. It is up to the Incident Captain to decide whether or not the Defenders would have access to the information they are seeking clarity on.]

2. Decision

Once the Defenders have reached a consensus, they declare which DETECTION they will be attempting, and roll the d20. You may only play 1 DETECTION per turn. Established Detections (top row) add a **+3 modifier** to the dice roll when they are played. These have an advantage as they indicate detections that your team is very experienced with. Other Detections (bottom row) **do not receive any modifiers.**

3. Rolling

When the Defenders wish to play a DETECTION card, they must roll the die to determine if the DETECTION is successful or if it fails to detect an attack.

Failure	**1-10**
Success	**11-20**

Remember to add any relevant modifiers to the roll! A roll of either a natural 1 or natural 20 (indicating the number on the die face before any modifiers are added) or 3 failures in a row will trigger an INJECT!

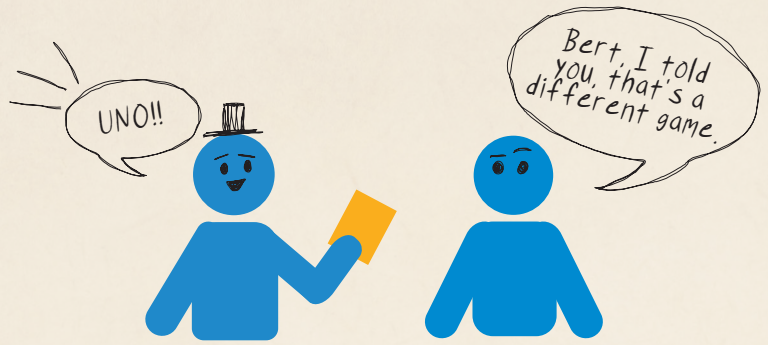

If an INJECT is triggered: Draw 1 card from the top of the INJECT pile and reveal it to all players. Follow any instructions that may be on the card, and have the Defenders discuss how (or if) this INJECT will affect their investigation.

INJECTS simulate the random events that can happen during a security incident. They add a bit of chaos to the scenario and spur important conversations. Some might not affect the game at all... or might end it. Either way, they're always unexpected.

4. Outcome

On a failure, nothing new is learned and the turn ends. On a success, the Incident Captain checks if the DETECTION played is listed under "Detection" on any of the Attack cards. If it is, they reveal that card to the Defenders. If the DETECTION could detect multiple Attack cards, it is up to the Incident Captain to choose only one card to reveal. (As in real life, when doing incident response, you find one thing at a time, not everything all at once.) After a DETECTION has been played, regardless of outcome, that card will have a 3-turn cooldown period during which it cannot be used again.

[Incident Captain Tip: If a DETECTION is unsuccessful, ask the Defenders for a reason—whether financial, political, personnel-wise, or technological—why the DETECTION would not be successful at that time.]

This is the most important part of the educational process of Backdoors & Breaches.

Ending The Game

The turn cycle repeats until whichever comes first:

The Defenders have revealed all 4 Attack cards
OR
10 turns have passed

Ready to Start Playing?

Play Online: *https://play.backdoorsandbreaches.com/*
Order Physical Decks: *backdoorsandbreaches.com*

CONTAINMENT & ERADICATION

Be Prepared with These Core Strategies

written by Patterson Cake || @SecureCake || @secure-cake

Once you've identified an active threat, stopping any ongoing unauthorized access and activities (containment) and preventing their recurrence (eradication) becomes top priority. Since complexity is the enemy of security, especially during a crisis, we need a simple, effective, repeatable plan. Let's simplify!

Understanding the Battlefield

Everything we care about happens on an endpoint (a network-connected device with an OS, user-interface, services, and/or data) through an identity (user/service accounts). Simply put, the threat actors are attacking your endpoints over the internet using valid credentials.

But who controls your network (ingress/egress) and identity (credentials)? Good news—you do! With that in mind, you can reassert "battlefield dominance" via two simple moves: disconnect the internet and rotate all credentials. This was my mom's go-to tactic.

No internet connectivity or valid credentials for the threat actors? No more command and control (C2) for authenticated actions! Some authenticated sessions may persist, but you can squash those next (revoke sessions, restart a service, reboot a system).

Simple Doesn't Mean Easy

You'll notice I've repeatedly said "simple," not "easy." If you haven't planned how to disconnect the internet and rotate all credentials, the process may end up being pretty blunt force. Unplug the firewall, if you must. Disable all but a couple high-privileged accounts and work your way toward changing passwords and re-enabling accounts. This is not ideal from a business-continuity perspective, but desperate times call for desperate measures. Trust me, allowing ongoing C2 and unauthorized authenticated sessions is ultimately far worse for business.

Disconnecting Internet & Rotating Credentials

If you prepared in advance by creating a C2 disruption runbook, it's still possible that you'll need "nuclear" options (like disabling all internet access.) However, with a predefined "Top 10" allowlist of business-critical URLs, you could quickly restore essential function. Can you disable the internet by physical location, for a specific network segment, or a handful of impacted systems? Map out your most granular options, document and test them, discuss invocation criteria and emergency authority, and you'll be ready to evict the threat actor from the battlefield!

Do the same for credentials: PLAN AHEAD. Map out your identity management systems (Microsoft Active Directory, Entra ID, AWS IAM) and prioritize the most privileged credentials, such as "Domain Admin," "Global Admin," "AWS root accounts," etc.

Consider rotating passwords for business-critical accounts while disabling non-essential accounts until further investigation. Don't forget accounts that exist outside of your primary identity management systems—printers, network-attached storage, or network devices may not seem high priority, but can be difficult to interact with at scale.

Outline a technical approach for each identity platform, document the process with invocation criteria and associated authority, and test until you're ready to disrupt any unauthorized authentication in your environment.

Prepare for Success!

While disrupting unauthorized ingress/egress and authentication is not the end of your incident response process, preparing for and implementing these steps positions you to take a deep breath, proceed methodically with your investigation, identify the threat actor's footprint and activity, and move along to successful remediation and recovery.

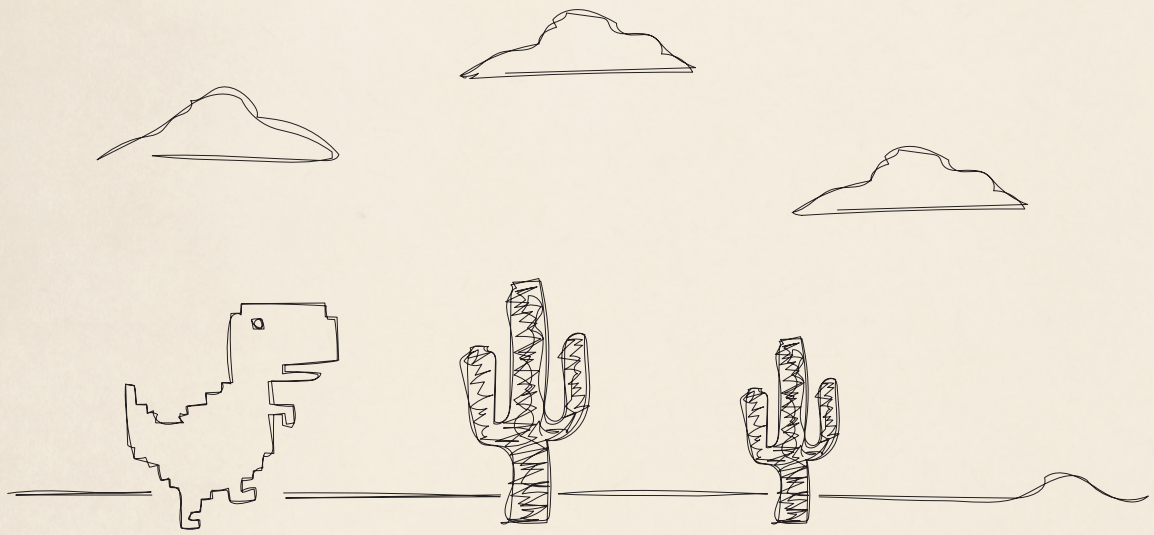

Try:

- Putting your phone away,
- Connecting with other humans,
- Enjoying the peace of nature.

DEEPBLUECLI 101

https://github.com/sans-blue-team/DeepBlueCLI

written by Tom DeJong || linkedin.com/in/dejongtom

DeepBlueCLI, a PowerShell module created by Eric Conrad, is a lightweight yet powerful threat hunting tool for blue teams. While its applications span incident response, security monitoring, detection engineering, threat hunting, and training, we will focus on its utilization during incident response engagements.

Key Benefits for Incident Responders

DeepBlueCLI can be used to detect a range of indicators, including initial compromise, persistence, privilege escalation, credential access, lateral movement, and defense evasion.

This tool excels in three primary scenarios:

- Rapid triage of Windows events log (EVTX) files
- Identifying threat actors without using a SIEM
- Quick scoping of a potentially compromised host

It's important to note that DeepBlueCLI is not a replacement for a full forensic parser, a timeline builder, or as a substitute for EDR telemetry.

Getting Started and Basic Usage

You can download or clone DeepBlueCLI from its official GitHub repository. The repository also includes sample EVTX files that you can use for testing and practicing commands. These sample EVTX files are located in the following directory within the repository.

```
DeepBlueCLI-master\evtx\
```

To start, run PowerShell as an administrator and navigate to the DeepBlueCLI-master directory.

The typical command structure is:

```
.\DeepBlue.ps1 <event log name>
```

or

```
.\DeepBlue.ps1 <evtx file path>
```

If you run into script execution errors when trying to run the tool, you may need to enter this command to allow the script to execute:

```
Set-ExecutionPolicy
Bypass -Scope Process
```

Processing Acquired EVTX Files

You can process the local Windows Security EVTX logs with these commands:

```
.\DeepBlue.ps1 -log security
```

or

```
.\DeepBlue.ps1 -log system
```

You can also use DeepBlueCLI to analyze EVTX logs acquired during an engagement with this command.

```
.\DeepBlue.ps1 .\evtx\
metasploit-sysmon.evtx
```

Note: metasploit-sysmon.evtx is a sample file included in the repository. Replace the path with the location of your own EVTX file.

Scanning Directories

You can use this command to scan multiple EVTX files in a directory for suspicious activity and output the results to a separate file:

```
.\DeepBlue.ps1 C:\Users\Security\Downloads\DeepBlueCLI-master\
DeepBlueCLI-master\evtx\*.evtx | Tee-Object DeepBlue-Output.txt
```

Converting to Different Formats

Another valuable feature of this tool is that it outputs results as PowerShell objects, enabling easy conversion into various formats using PowerShell's built-in utilities. Some of the formats that you can export the data to are CSV, JSON, XML, and more.

This command will convert results to a separate file in JSON format.

```
.\DeepBlue.ps1 C:\Users\Security\Downloads\DeepBlueCLI-
master\DeepBlueCLI-master\evtx\metasploit-sysmon.
evtx | ConvertTo-Json | Out-File Metasploit-Results.json
```

Conclusion

While DeepBlueCLI isn't a replacement for a SIEM or full forensic suite, it serves as an invaluable rapid triage tool for working with EVTX logs. Being able to quickly identify the "low-hanging fruit" early in the incident response will help direct focus towards the areas requiring your deep forensic analysis.

Saves time, sanity, and a small piece of your soul.

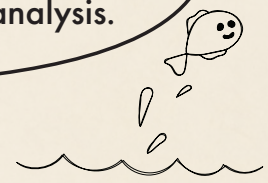

Further Learning

Check out John Strand's SOC Core Skills training course from Antisyphon Training (which features a DeepBlueCLI lab):

https://www.antisyphontraining.com/course/soc-core-skills-with-john-strand/

You can view the steps of that lab here:

https://github.com/strandjs/IntroLabs/blob/master/IntroClassFiles/Tools/IntroClass/deepbluecli/DeepBlueCLI.md

FORENSIC DATA

How to Acquire and Retain Vital Evidence

written by Ronald A Mays Jr || linkedin.com/in/ronald-mays-a8300914/

While "digital forensics" and "incident response" often appear together (as in the acronym "DFIR"), they represent two distinct disciplines.

What's in a Name?

The American Heritage Dictionary defines forensics as "The use of science and technology to investigate and establish facts in criminal or civil courts of law."

Contrast this with IBM's definition of incident response: "Incident response [...] refers to an organization's processes and technologies for detecting and responding to cyberthreats, security breaches or cyberattacks." Note the absence of any reference to legal proceedings—it simply describes the process by which an organization detects, halts, and recovers from an incident.

Legal Considerations

Although most incidents **never end up in court,** all should be **treated as if they might.**

Taking the proper steps before, during, and after response ensures that any data collected can be admissible as evidence.

It's important to know the rules of evidence for relevant jurisdictions. Global corporations should consider each jurisdiction they have locations in when developing standard operating procedures (SOPs). Always include legal counsel when developing your organization's SOPs.

Data Preservation

A common rule for digital evidence is the retention of data in its "original state." This is accomplished by using validated hardware and software to create an exact bit-for-bit copy of a piece of digital media. Yet during incident response, it is rare that full forensic images are created—typically, only a small subset of data is collected from victim devices for analysis. Before any action is taken, consult legal advisors to decide whether to collect full forensic images.

A good practice is to create a folder for the original data, generate and document a hash of its contents, and then create a working copy in a separate folder. Only perform analysis on the working copies of your data to preserve the integrity of the original data.

Documentation

Documentation is critical when handling potential evidence. A chain of custody (CoC) records details about what the data is, where it came from, who collected it, whom it was transferred to, a brief description of what was done to it, and so forth.

The National Institute of Standards and Technology (NIST) provides a sample CoC here:

https://www.nist.gov/document/sample-chain-custody-formdocx

Each person working with the data must keep detailed notes: record dates, times, and any actions taken. Binary files often need to be parsed by special tools to make their data human-readable—include any tool names and versions. The same should be done with human-readable data, such as text-based log files. While you likely don't need to document the version of Excel used to read a log file, if you added a column to show UTC date/time stamps in local time, that is worth noting.

The gold standard for documentation is this: notes should be detailed enough that someone with a similar background can recreate your steps and produce the same results.

This is the scientific method part of digital forensics. Processes need to be repeatable and generate the same results.

Tool Validation Procedures

Before an incident occurs, verify that your hardware and software are working. Test write blockers to ensure they do not alter the evidentiary media. Validate that your forensic software functions correctly (it's not unheard of for new versions to report the wrong date/time stamps when parsing binary data).

While there is no set standard for validation processes, some guidelines exist to help you develop validation SOPs.

Validation SOPs guidelines:

NIST:
https://www.dhs.gov/science-and-technology/nist-cftt-reports

SWGDE:
http://cet4861.pbworks.com/w/file/fetch/120024474/4861.2014-09-05_SWGDE_Recommended_Guidelines_for_Validation_Testing_V2-0.pdf

Additional Considerations for Incident Handling

When managing incidents, also consider:

- What are the data retention policies of your Endpoint Detection and Response (EDR) tools?
- How is your cloud data preserved?
- Who has access to sensitive or personal data?

These factors often depend on your license agreement with providers. Make sure data can be maintained in a secure, long-term manner, as it can take years until an incident makes it into the courtroom. Implement processes to export incident data from systems and services while preserving integrity and meeting long-term storage requirements that protect its availability and confidentiality. Access to confidential information must always be managed carefully.

Remember: Involve your legal team, validate your tools, work from copies (never original data), document everything meticulously, and maintain chain of custody.

"I'll remember it later..."
Narrator: They did not remember it later.

BUSINESS IMPACT PLANNING

Collaboration, Communication, & Emergency Powers

written by Erik Goldoff, CISSP || @ErikG || linkedin.com/in/egoldoff/

Incident response is not solely an IT/infosec duty; incident response is a business responsibility!

Incident Response Success 101

Let's summarize how important proper collaboration, communication, and execution (or delegation) of emergency powers are to your incident response success.

Various 'incidents' can affect your organization, so it's important to have a business continuity plan (BCP) in conjunction with your IR plan. A BCP is a documented outline of how an organization will continue functioning during and after an emergency or disruption. All incidents can impact or halt your business; your incident response goal is to mitigate the effects as quickly as possible. You'll need to collaborate with various parts of your organization to achieve this.

Ask these questions about how prepared your team is:

- You might need all hands on deck, including outside of business hours, for your staff who are focused on incident response tasks. Who has the authority to approve overtime and cancel time-off?
- If you have everyone working to resolve your incident, you won't be able to schedule off-site meal breaks. Who holds the decision-making power to have meals delivered, including documentation of any staff dietary restrictions?
- What about staff who have no transportation options outside of normal business hours? Can a van/bus and a driver be easily hired to make sure necessary staff are on-site when needed? What about those who need childcare outside of business hours?
- You might need to bring in a professional, dedicated digital forensics and incident response (DFIR) team to augment or even lead your efforts. Who has the authority to fund and sign the necessary contracts?
- Depending on the type of organization and scope of incident, there may be legal requirements for preservation of, and proof of, chain of custody for physical evidence. There could be requirements to notify law enforcement before taking any action. Who on staff has the proper background to know what these legal requirements are?
- There may be legal or organizational requirements surrounding who can make public statements and what level of detail may be disclosed. Who has the authority to make those decisions?
- Internal resources like file shares, PBX phone system, and/or email might become inaccessible. Are there alternate communication details accessible, with rules around how and when they can be used?

If all else fails, you can always send a strongly worded carrier pigeon.

- You might lose utilities, including electrical power, data and voice circuits, HVAC, or others critical to operations. Who has access to all the proper contacts, including escalation lists?
- It may be necessary to quickly set up a temporary office to keep critical functions operating. Who has authority to investigate temporary location leases, quick internet access, and equipment purchases or rentals?

Build and Maintain a Reliable Business Continuity Plan

Based on your organization's size, you could be lacking a robust, practiced, reliable BCP. Having contact information readily available for decision makers in all relevant business units is crucial. The time to establish a plan is not after you need it, but now!

Start by assigning an incident commander, or a team of commanders, who can hand off responsibilities during extended response times. They shouldn't be subject matter experts (SME) in any particular area, as you'll want your SMEs focusing their efforts within their field. The incident commander should be trusted with coordinating communications between all the different required entities so that your SMEs can be unburdened with those details and free to work on the incident.

These incident commanders should have hard-copy access to the BCP and Escalation List. Moreover, they must be empowered to notify the financial decision makers about unplanned budget requirements, contact all relevant decision makers for each responsibility, bring together the proper team leaders at the proper time, and coordinate status updates among those leaders.

Conclusion

If you work in infosec and think of incident response as solely your responsibility, it will be a difficult experience when it's "go time." You'll be most successful if you have the discussion now: assemble contact lists, appoint the necessary incident commanders and other roles, and practice (run tabletop exercises) for various scenarios. Proper proactive planning, well, you know..

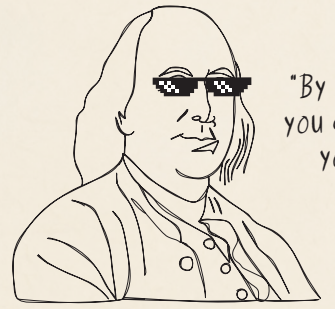

"By failing to prepare, you are preparing to..." you get the idea.

Check out **"A Simple, Useful IR Plan"** on **page 12** for more information on creating your IR plan.

VELOCIRAPTOR 101

https://github.com/Velocidex/velociraptor

written by Andrew Scott || @PurpleTeamLead || linkedin.com/in/andrew-scott-81b902250/

Docs: https://docs.velociraptor.app/
Downloads: https://docs.velociraptor.app/downloads/

We raptor a present for you!! (Yeah, I'll see myself out...)

Velociraptor is an open-source digital forensics and incident response (DFIR) platform that gives security teams deep visibility into Windows, macOS, and Linux endpoints. Built by Velocidex and later acquired by Rapid7, Velociraptor is lightweight, scalable, and designed for flexible investigations.

Velociraptor and VQL Basics

At its core, Velociraptor uses a server to collect information and communicate with endpoints with the client installed. Velociraptor uses the Velociraptor Query Language (VQL) to collect and analyze evidence. Clients perform most of the heavy lifting locally, streaming results back to the server, allowing for efficient scaling across thousands of endpoints with built-in throttling and concurrency controls.

While experienced users often craft large hunts or custom artifacts, newcomers can start with a few straightforward queries to answer common investigation questions.

Wait a minute, is VQL the sequel to SQL? Because if Jeff Goldblum shows up warning us about uncontrolled queries, I'm listening.

Typical Workflows You'll Actually Use

- **Single Host Triage:** Collect built-in artifacts (e.g., Evidence of Execution, Listening Ports), pivot in a Notebook with ad-hoc VQL, and download artifacts for deeper forensics.
- **Hunt at Scale:** Target by label/OS/query and push artifacts to thousands of clients; review and post-process with server-side VQL (e.g., hunt_results()).
- **Near-Real-Time Monitoring:** Deploy Event artifacts using watch_evtx() (and, where appropriate, watch_etw()) to stream detections such as process creation or registry changes.

Deploying Velociraptor

Velociraptor can be deployed in client/server mode or run as a stand-alone binary for ad-hoc investigations. In enterprise setups, you install the server on Windows or Linux, generate a client config, and use the same binary to configure endpoints. Velociraptor provides a robust graphical user interface (GUI) as well as a command-line interface (CLI).

Endpoints run the Velociraptor client, which connects securely back to the server over TLS and allows analysts to query or collect artifacts in real time. For smaller investigations, the portable binary can be run directly without installation, making it flexible for incident response.

Practical VQL Queries

VQL uses an SQL-like structure: SELECT <expressions> FROM <plugin>(named_args) as the core, where the plugin call is the row-producing data source.

The user can optionally add WHERE, GROUP BY, HAVING, ORDER BY, and LIMIT/OFFSET clauses, and can use variables, functions, and regex in expressions to shape and filter the results.

Checking Running Processes

Using VQL, you can quickly list running processes and their command lines.

```
SELECT Name, Exe, Cmdline,
Username FROM pslist()
```

What to Expect: Common processes like explorer.exe, svchost.exe, and chrome.exe. Suspicious results may include unusual command lines, unsigned executables, or administrative tools run by unexpected users.

Tracking New Process Creation

Windows security logs capture process creation events (ID 4688). Parse these logs to reveal suspicious commands.

```
SELECT * FROM parse_
evtx(filename="C:\\Windows\\
System32\\winevt\\Logs\\Security.
evtx") WHERE EventID = 4688
```

What to Expect: Timestamp, process name, command line, and parent process. Look for tools like powershell.exe, wscript.exe, or cmd.exe spawned by Office apps or unusual parents.

Watching for New Services

By monitoring for Event ID 7045, you can detect malicious new services in near real-time.

```
SELECT * FROM watch_
evtx(filename="C:\\Windows\\
System32\\winevt\\Logs\\System.
evtx") WHERE EventID = 7045
```

What to Expect: Details of newly installed services such as service name, binary path, and user context. Malicious services may have random names or point to executables in temp directories.

Reviewing User Accounts

With one query, you can enumerate local users and SIDs.

```
SELECT Name, UUID FROM
Artifact.Windows.Sys.Users()
```

What to Expect: Standard accounts like Administrator, Guest (often disabled), and domain users. Investigate unfamiliar accounts, especially those with admin privileges.

Detecting Suspicious Scheduled Tasks

Velociraptor makes it easy to find scheduled tasks.

```
SELECT * FROM Artifact.Windows.
Tasks.ScheduledTasks()
```

What to Expect: Task name, path, and the command executed. Look for tasks pointing to executables in unusual folders or with strange names like 'Updater123'.

Hunting With Built-In Artifacts

Velociraptor's built-in artifacts can be used to find information on clients. Start by selecting the wrench icon on the left of the GUI (Artifacts), choose which artifact you would like to hunt, then select the crosshairs icon.

Generic.System.Pstree

- **Purpose:** Displays processes in a parent/child hierarchy.
- **Use Case:** Identify suspicious execution chains (e.g., Office spawning PowerShell).

Windows.System.Services

- **Purpose:** Enumerates installed Windows services and startup configs.
- **Use Case:** Detect persistence mechanisms or hijacked services.

Windows.EventLogs.EvtxHunter

- **Purpose:** Parses Event Logs for security-relevant events.
- **Use Case:** Spot brute-force attempts, logon anomalies, PowerShell use, lateral movement.

Windows.Sysinternals.Autoruns

- **Purpose:** Collects autorun/persistence registry entries.
- **Use Case:** Detect malware startup entries or persistence.

Windows.Network.Netstat

- **Purpose:** Collects active TCP/UDP network connections and their associated processes.
- **Use Case:** Spot suspicious outbound connections (e.g., to rare external IPs or unusual ports) and tie them back to processes like PowerShell or rundll32.

Windows.Network.ListeningPorts

- **Purpose:** Lists processes bound to network ports.
- **Use Case:** Identify backdoors, C2 connections, or unexpected services listening on a machine.

CALLING REINFORCEMENTS

Your Incident Response Plan, Before The Fire Starts

written by John Strand || @strandjs || linkedin.com/in/john-strand-a1b4b62/

Many people ask: At what point do I bring in law enforcement during a cyber incident? When should I get outside help such as another consulting firm to support incident response? The honest answer is simple: before you ever need them. The best time to build reinforcements is not during a crisis. It is now.

It's the same idea as that old proverb: "The best time to plant a tree was 20 years ago. The second best time is now." Cybersecurity works the same way. Incident response is messy, emotional, and chaotic. You don't want to be flipping through your contacts for a random phone number when your domain controller is encrypted and your CEO is standing behind you asking for updates. You want relationships, people who already know you and are ready to help.

Build Your Network Before You Need It

Start with your own community. Get involved with organizations such as InfraGard and HTCIA. Attend local cybersecurity meetups, conferences, and regional response groups. Shake hands. Collect business cards. Swap stories. These small moments build the trust you will rely on when your network goes sideways at 2 a.m. on a Saturday.

Identify who your law enforcement contacts are now, such as local FBI field office personnel, state cybercrime investigators, or specialized digital forensic units. Introduce yourself. Learn what they do and how they operate. Because during an incident is a terrible time to be making introductions.

When to Call Law Enforcement

Whenever you are handling an incident, reach out to law enforcement early. They might not jump into every ransomware attack or stolen laptop case, but it is still important to keep them informed. They have something you do not have: legal authority. If you need warrants, subpoenas, or data preserved from third-party providers, law enforcement has the tools to make that happen. If nothing else, contacting them shows you are acting responsibly, and that matters later.

Line Up the Right Legal Team

Your next reinforcement is legal. You should already have a law firm on an incident response retainer before something goes wrong. But don't just pick any law firm. Choose one with real experience handling cyber incidents.

Traditional attorneys often respond to breaches by locking everything down and preventing collaboration. That approach creates chaos. Experienced breach counsel understands that incident response requires teamwork, not isolation.

They know how to:

- Work with law enforcement
- Coordinate with security firms
- Maintain legal privilege where appropriate
- Guide sensitive communications
- Draft clear breach notifications
- Stay cool when everything burns

If they have been through this before, they will not fall apart when it is your turn.

Cyber Insurance Requires Work

If you are in an incident and asking whether your cyber insurance covers it, it is already too late. You either have coverage or you do not. And getting that coverage is becoming harder.

Insurance companies are nervous. Ransomware costs are rising and payouts have become massive. Even if you do get coverage, read the fine print. Most policies require due diligence—you must demonstrate that you were following reasonable security practices. If you have ignored vulnerability management, asset inventory, access control, logging, and backups, your insurer may deny your claim.

Cyber insurance is not a magic safety net. It is a contract that assumes you are taking security seriously.

Build Relationships, Not Just a Plan

If there is one theme here, it is this: In a crisis, you do not rise to the occasion. You fall back on the strength of your relationships. When an incident hits, you should not be scrambling to track down experts. You should already know who to call and they should already know you.

Start building your network today. Law enforcement contacts, legal counsel, insurance providers, and trusted incident response partners should already be in your corner. When the day comes, you will not be calling strangers. You will be calling reinforcements. Friends. Colleagues. People who already have your back.

Start now.

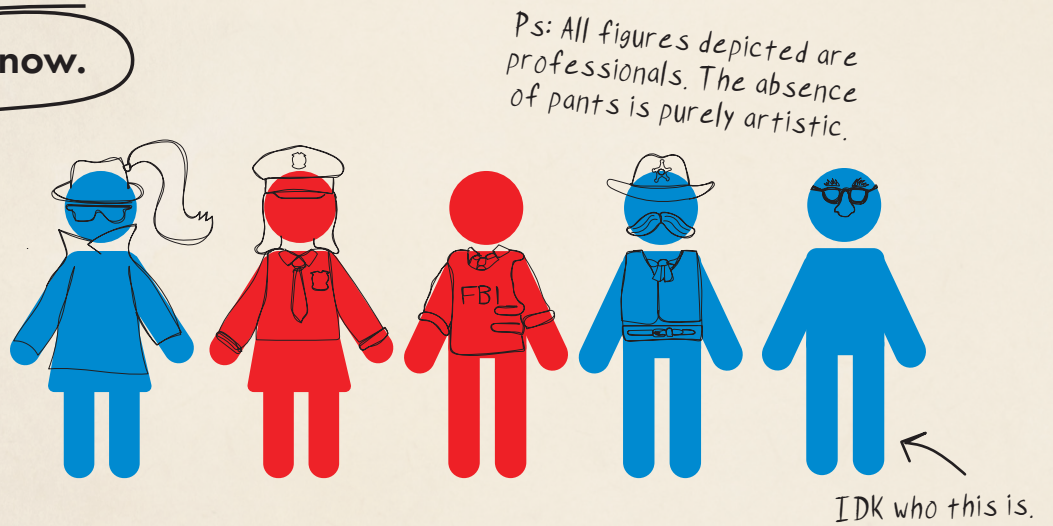

AFTER THE DUST SETTLES...

Lessons Learned & Continuous Improvement

written by Brian Baskin || @bbaskin

When the alerts stop, the logs quiet down, and the last incident bridge call ends, there's often a deep silence. This is a surreal moment for the incident responder who wants to believe their work is done while preparing for the long hours that may yet continue. This is a crucial period that should be dedicated to improving defenses for future attacks.

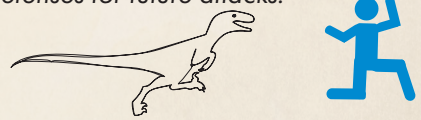

Responder Recovery Comes First

Incidents are exhausting. They create stress, late nights, missed family events, and frayed tempers.

The first step is to take a breath. It's natural to want to jump into fixes, but vital to recognize the mental health impact of stress and exhaustion. "Post-mortem" mistakes can be as costly as the incident itself, especially as adversaries may return weeks later with identical tactics.

Recovery is not just about systems; it's about people. Leadership should ensure appropriate downtime, rotate responsibilities, and normalize conversations about burnout. Mental health support helps prevent long-term damage to your team's resilience.

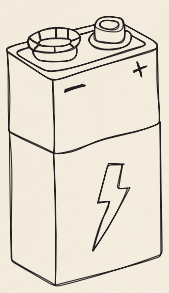

Developing the Attack Narrative

The next major task is developing an attack narrative. How did the adversary enter the network? What was stolen? What was their dwell time? Answering these questions requires extensive logs from servers, identity and cloud management tools, and security products.

Were your detection tools too slow? Were there delays in confirming suspicious activity? Was a crucial alert buried by false positives? Start by mapping the attacks to the MITRE ATT&CK® framework to identify security gaps. If you don't have enough visibility to answer these questions, that gap becomes one of your top action items.

Honest Assessments

Be brutally honest about why the attack worked. Classifying an attack as just "phishing" or an "unpatched vulnerability" doesn't give necessary context. Why did a phish work so well? Was a particular person targeted? Understanding the human factor helps tailor future awareness training and to adjust controls.

Take a page from Google's SRE (Site Reliability Engineering) Handbook which, while calling out how to learn from an incident, puts emphasis on blameless recovery, separating the events from the individual. Some incidents originate from a forgotten test service overlooked by security teams and exposed to the internet years before. What the service was and who created it does not matter as much as how it became internet-accessible and went undetected.

Recovery doesn't need to become personal.

Yeah well, the ransomware made it really personal.

Implementing Improvements

Once facts are laid bare, improvements can begin. This doesn't mean patching one application or rewriting one detection rule; it's implementing enterprise-wide controls such as multi-factor authentication to cloud services, improving log collection for certain applications, and ensuring detection coverage for high-value assets. It also means streamlining escalation paths so that containment decisions happen within minutes, not hours. Clear ownership mapping, asset tagging, and updated emergency contact lists are small investments that pay off big.

Don't just focus on mitigation tasks. Turn the incident into a training opportunity by implementing tabletop exercises within your teams (even at executive level). Regular drills build muscle memory, improve cross-team collaboration, and reveal weaknesses before an incident hits. Refer to CISA's Tabletop Exercise Packages and check out Backdoors & Breaches. Also consider hiring a third-party penetration company to test your recovery and mitigation solutions. Hire a third party...or, you know...us.

"Never let a good crisis go to waste." Emergencies create short windows where previously rejected proposals get attention. If a $50,000/year service previously seen as a luxury could have prevented a $400,000 loss, position it as a necessary investment. Pitch upgrades as risk-mitigation necessities backed by numbers, especially if they would have reduced the hours of labor involved in the response.

Most importantly, response plans should be living documents that are updated as changes are implemented and referenced when planning new security controls. This keeps lessons fresh and ensures improvements don't stagnate or get forgotten amid shifting priorities.

You can't stop all attacks, but you should be able to respond to them all. Focus on ensuring that your team can find alerts, make assessments, and mitigate risks quickly and with little impact. These skills are best learned from doing, and those who fail to learn from the past are destined to repeat it.

Additional Resources

NCSC Cybersecurity Response and Recovery:
ncsc.gov.uk/files/NCSC_A5%20Response%20and%20Recovery%20Guide_v3_OCT20.pdf

Cybersecurity Incident & Vulnerability Response Playbooks:
www.cisa.gov/sites/default/files/2024-08/Federal_Government_Cybersecurity_Incident_and_Vulnerability_Response_Playbooks_508C.pdf

CISA's Tabletop Exercise Packages:
www.cisa.gov/resources-tools/services/cisa-tabletop-exercise-packages

Backdoors & Breaches:
backdoorsandbreaches.com

Google SRE Handbook on post-mortem analysis:
sre.google/sre-book/postmortem-culture/

Atlassian's incident management for high-velocity teams:
atlassian.com/incident-management/postmortem

LFI (Learning from Incidents):
learningfromincidents.io

WHO IS BLACK HILLS INFORMATION SECURITY (BHIS)?

Established in 2008, Black Hills Information Security has created a network of companies in the infosec industry dedicated to providing affordable, outstanding products and services that cover all of your information security needs from pentesting to training.

Each company helps to support the infosec community in their own way—offering free educational content, open-source tools, or even donating to various projects.

We Offer

- Penetration Testing
- Red Teaming
- Active SOC
- Blue Team Services
- Purple Teaming
- Threat Hunting
- Incident Response
- Consulting
- Training
- IR Tabletop Exercises
- ANTISOC (Continuous Penetration Testing)

At Black Hills Information Security (BHIS), we help you find and fix vulnerabilities before attackers do. But we don't stop there. We want to teach you how attacks happen and how to defend against them. With our hands-on expertise, open-source tools, and a real-world approach, we make security practical, actionable, and continuous.

We've worked with

- Credit Unions
- Banks
- Investment Firms
- Higher Education
- Health Care
- Medical Devices
- Insurance
- Law Firms
- Local Government
- Municipalities
- Real Estate
- Retail
- Technology
- IT
- Software
- Utilities
- ICS/SCADA
- Defense/ Aerospace

From the smallest mom & pop shops to the biggest Fortune 5 companies, our top priority is helping you understand and achieve your security needs.

OFFENSIVE

Our team of 40+ pentesters conduct more than 1000 security assessments every year.

Knowledge transfer from our team to yours empowers you to mature and grow, so we take special care in our reporting. Our reports provide you with not only what was successful in an engagement, but also highlight your current strengths by showing what efforts failed.

Our experienced testers help you understand and fortify your own system.

- PENETRATION TESTS:
  - ◊ Red Teams
  - ◊ Internal
  - ◊ External
  - ◊ Assumed Compromise
  - ◊ Web App
  - ◊ Mobile
  - ◊ Wireless
- ANTISOC (Continuous Penetration Testing)
- AI SECURITY ASSESSMENTS
- HIGH-PROFILE RISK ASSESSMENTS

DEFENSIVE

Ever wondered what it's like to have a security partner truly in your corner? That's exactly what we offer with our ACTIVE SOC service.

If you need a primary SOC or just a second set of eyes, we're here to help you detect and respond to threats 24/7.

- ACTIVE SOC:
  - ◊ Continuous Monitoring & Alerting
  - ◊ Log Analysis
  - ◊ Attack Surface Monitoring
  - ◊ Adversary Emulation
  - ◊ Cyber Deception
  - ◊ Threat Hunting
  - ◊ Purple Teaming
- PURPLE TEAMING
- BREACH ASSESSMENT
- ATOMIC CONTROLS ASSESSMENT
- ACTIVE DIRECTORY REVIEW
- BHIS EXPERT SUPPORT TEAM

INCIDENT RESPONSE

With experience as both red and blue teams, our IR team knows the ways to hunt down threats and analyze the evidence because we've been on both sides.

Whether you've already been breached, or you're looking to prevent it, we've got you covered.

- INCIDENT RESPONSE:
  - ◊ Training
  - ◊ Log Collection & Analysis
  - ◊ IR Retainer
  - ◊ Monitoring
  - ◊ Consulting
  - ◊ Checklists & Playbooks
  - ◊ Tabletop Exercises

BLACK HILLS | Information Security

ANTISYPHON TRAINING — antisyphontraining.com

wildwesthackinfest.com

ACTIVE COUNTERMEASURES — activecountermeasures.com

REKCAH — rekcahcomics.com

PROMPT# — promptzine.com

bhis.co

ANTISYPHON TRAINING

Learn What's Bad; Do What's Good

you heard me

*We're here to **disrupt** the traditional training industry by providing **affordable education that doesn't suck.** Whether you're a total newbie or a seasoned pro, dive into interactive, hands-on sessions with certified instructors, and build real world skills **while earning cool badges.** From pay-what-you-can to corporate* options, on-site live training to on-demand courses, and everything in between—**we're all about making your learning journey effective, engaging, and ridiculously fun.***

L Live OD On-Demand
EX Excluded from Corporate Catalog

Pay-What-You-Can Courses

- [ ] Active Defense and Cyber Deception L OD
  John Strand
- [ ] Cybersecurity Foundations Training L OD
  John Strand
- [ ] Professionally Evil CISSP Mentorship Program L
  Kevin Tackett & the Secure Ideas Team
- [ ] Professionally Evil ICS & OT Fundamentals L
  Giovanni Cofre
- [ ] SOC Core Skills L OD
  John Strand

Pay-What-You-Can Workshops

As part of our PWYC initiative, we offer a variety of 4-hour workshops that introduce you to foundation-level cybersecurity topics. These workshops get you started with the basics so you can pivot to the next level courses!

Learn more here:

antisyphontraining.com/pay-what-you-can/

---

*Our **Corporate Catalog** brings the same energy, just scaled for business.
At **$1,500 per seat with a 5-seat minimum**, you'll get access to **over 40 courses**, the **Antisyphon Cyber Range** and a dedicated corporate dashboard where managers can track progress, measure impact, and see their team's skills grow in real time.

For more information, contact us at **info@antisyphontraining.com**

learn by doing for goodness sake!

Full Course Catalog

- [ ] Active Directory Security & Hardening L OD
  Jordan Drysdale & Kent Ickler
- [ ] Advanced Endpoint Investigations L OD
  Alissa Torres
- [ ] AI for Cybersecurity Professionals L
  Joff Thyer & Derek Banks
- [ ] Blue Team Foundations with the Atomic Controls OD
  Bryan Strand
- [ ] Breaching the Cloud L OD
  Beau Bullock
- [ ] CIS Controls Mastery Course for MSPs OD EX
  John Strand
- [ ] Cybersecurity Incident Command L
  Gerard Johansen
- [ ] Cyber Threat Intelligence 101 L OD
  Wade Wells
- [ ] Defending M365 and Azure L OD
  Kevin Klingbile
- [ ] Enterprise Forensics and Response L OD
  Gerard Johansen
- [ ] Enterprise Security for All L OD
  Bob Hewitt & Rich Fifarek
- [ ] Foundational Application Security L OD
  Kevin Tackett
- [ ] Security Foundations L OD
  Hayden Covington
- [ ] HackerOps L OD
  Ralph May
- [ ] Hacking Active Directory L OD
  Dale Hobbs
- [ ] Hacking and Defending Satellite Infrastructure L
  John Strand & Ralph May
- [ ] Hands-On AI Security Risk Assessment L
  Jake Williams

Rumor has it, these instructors are so good, even malware asks them for advice!

- [ ] Hiring Handbook: How to Build a Team that Gets Stuff Done L OD
  Kip Boyle
- [ ] Incident Response Foundations L OD
  Derek Banks
- [ ] Intro to Android Pentesting L
  John Strand
- [ ] Intro to IoT Hacking L
  Rick Wisser & Dave Fletcher
- [ ] Intro to Network Threat Hunting L OD
  John Strand
- [ ] Intro to Penetration Testing of Non-Western IT Infrastructures L
  Steve Borosh
- [ ] Intro to Pentesting L OD
  John Strand
- [ ] Introduction to Industrial Control Systems L OD
  Ashley Shirley
- [ ] Linux Command Line for Analysts and Operators L OD
  Hal Pomeranz
- [ ] Linux Forensics L OD
  Hal Pomeranz
- [ ] Management 101 L
  John Strand & CJ Cox
- [ ] Modern WebApp Pentesting L OD
  BB King
- [ ] Modern Webapp Pentesting 2: Webapp Internals L OD
  BB King
- [ ] Network Forensics & Incident Response L OD
  Troy Wojewoda
- [ ] Next Level OSINT L
  Mishaal Khan
- [ ] Offensive Tooling for Operators L OD
  Chris Traynor
- [ ] Offensive Tooling Foundations L OD
  Chris Traynor
- [ ] OWASP Top 10 OD
  Jim Manico

More! More! More!

- ☐ Penetration Testing: Beyond the Basics L
  Tim Medin
- ☐ PowerShell for InfoSec L OD
  Carrie Roberts
- ☐ Practical iOS Application Security L
  Cameron Cartier & Dave Blandford
- ☐ Practical OWASP TOP 10: 2021 L OD
  Kevin Tackett
- ☐ Practical Window Forensics L OD
  Markus Schober
- ☐ Python for the Security Practitioner L OD
  Joff Thyer
- ☐ Red Team: Initial Access L OD
  Michael Allen
- ☐ Regular Expressions, Your New Lifestyle L OD
  Joff Thyer
- ☐ Reporting for Pentesters L OD
  BB King
- ☐ Secure Coding and API Hardening: Secure Design, Development, and Threat Modeling L
  Tanya Janca
- ☐ Securing the Cloud L OD
  Andrew Krug
- ☐ SELinux L OD
  Hal Pomeranz
- ☐ SOC Course L
  Christopher Crowley
- ☐ Threat Hunting with Velociraptor L
  Eric Capuano & Whitney Champion

Go ahead, get that camera out.

Scan the QR code to access the full course catalog.

On-demand classes are being added regularly.
Please check the Antisyphon Training website for the most current information.

Level Up Your Team

- Customized team training tailored to your budget
- Access to 40+ on-demand courses across cybersecurity, SOC, cloud, pentesting, and leadership tracks
- On-site and virtual private instructor-led training options
- Hands-on labs and Cyber Range access for real-world, scenario-based training
- Real-time reporting on progress, completion, and competency development
- Mapped training paths aligned to DoD 8140/8570, NICE Framework, and NIST standards

Sign up today!

Those firewalls won't build themselves!

antisyphontraining.com

We asked members of the BHIS Discord what they enjoy about the server. Here's what they had to say:

Putting yourself out there is scary, but so worth it.

> The BHIS community is supportive and engaged, and that has helped me grow as a practitioner and student of cybersecurity. I've gotten the opportunity to share things I have learned with others even newer than me to cybersecurity, or IT generally. That ability to give back so easily is really encouraging... I believe anyone joining here and interacting with honesty and interest can reap these same benefits, speaking as a member of multiple minorities.
>
> *– siriciryel*

> The info, talks, and community. For lack of a better word, the culture. There is acceptance, understanding, and an overall genuine sense of concern for the well-being of each other. And the shenanigans.
>
> *– PacificRed*

Let's be honest, mostly shenanigans.

> The story I usually tell people is how being active in this server had me fully prepared for the CrowdStrike incident in July 2024. A user here reported it and posted the fix 5 hours before stuff really started hitting the fan around the world... At first, I was like "why do they seem so panicked?" and then realized the gravity of what they were reporting and was very grateful to get to be the hero for my clients and avoid any downtime.
>
> *– dypsis*

100% real quotes. (I think...)

> I've had overwhelmingly positive interactions with the people here, which is quite the milestone for the internet.
>
> *– phailB0t*

Join our community of amazing humans

BLACK HILLS Information Security | 890 Lazelle Street Sturgis, SD 57785 | Contact Us: 701-484-BHIS or info@promptzine.com
PRINTED IN CANADA.

Read past issues of PROMPT# and Infosec Survival Guides!

MADE BY AND FOR THE COMMUNITY

Visit our online store for cool swag, comics, and more!

Everything You Need to Survive Incident Response (...well, almost)

Alright, maybe not everything... but definitely a solid start! The Infosec Survival Guide: Incident Response is packed with tips, tools, and insights for incident responders who are just getting started or want to sharpen their fundamentals.

This guide is a collaboration between BHIS and our amazing infosec community. Many of the articles inside were written by community members, for community members, sharing real-world experiences, lessons, and knowledge earned deep in the IR trenches.

Articles Covering:

- Alert Review
- Your Critical First-Hour Response
- A Simple, Useful IR Plan
- Incident Investigations
- Threat Actor SOP
- How to Play Backdoors & Breaches
- Forensic Data

Tool "101" Cheatsheets:

- KAPE
- Hayabusa
- DeepBlueCLI
- Velociraptor

No mentions of my contributions I see... ok cool cool cool. I'm cool, you cool? We're all cool.

...and more!

Hungry for more? Explore other editions of the **Infosec Survival Guide** and the **PROMPT#** zines to keep leveling up your knowledge in all things information security!

Brought to you by:

BLACK HILLS | Information Security

ANTISYPHON TRAINING — antisyphontraining.com

wildwesthackinfest.com

activecountermeasures.com

rekcahcomics.com

PROMPT# — promptzine.com

bhis.co