VERSION 2.0

PIVTR-D

A field manual for incident responders.

Habib Tora | September 2026

P	I	V	T	R	D
Preparation *readiness*	**Identification** *alerting, detection, hunting*	**Verification** *event, or incident?*	**Triage** *rank against objectives*	**Response** *scope, contain, eradicate, recover*	**Debrief** *lessons learned, post-mortem, RCA*

What this is, how to use it

Taken straight from Joshua Wright's excellent updated IR book, molded into a tldr guide for the common incident responder. The lifecycle, the decisions that go with it, and the paperwork nobody enjoys writing at two in the morning but everyone is glad of at the after-action review.

Find the phase you are in and go to its playbook in Section 3. Each one says what starts the phase, what finishes it, what record it leaves behind, and what to check along the way. Section 2 holds the decisions that should already have been made, so look there before improvising one at speed. Section 1 orients anyone who walks onto the bridge halfway through.

The checklists that go with these playbooks travel in a companion workbook, **pivtr-d-checklists.xlsx**, one tab per waypoint with columns for owner, date and notes. The record templates travel in **pivtr-d-record-templates.docx**. This document carries the model, the decisions and the reference tables.

1. The lifecycle

Read it left to right. Scope, contain, eradicate and recover repeat until scoping stops turning up new evidence. Decision makers sit behind the whole sequence rather than inside any one phase, because every gate needs them.

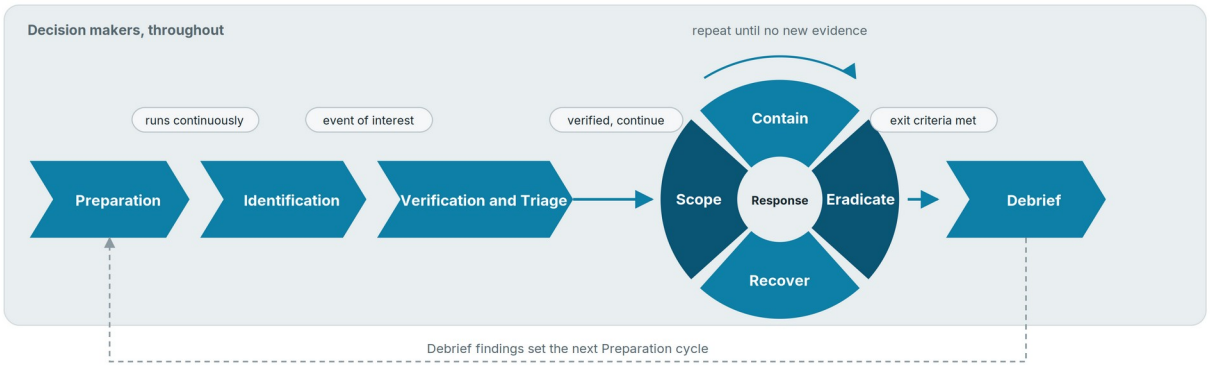

Each checkpoint sits above the junction it governs. Identification is drawn as continuous because it is. The dashed return path is the one most organizations never build.

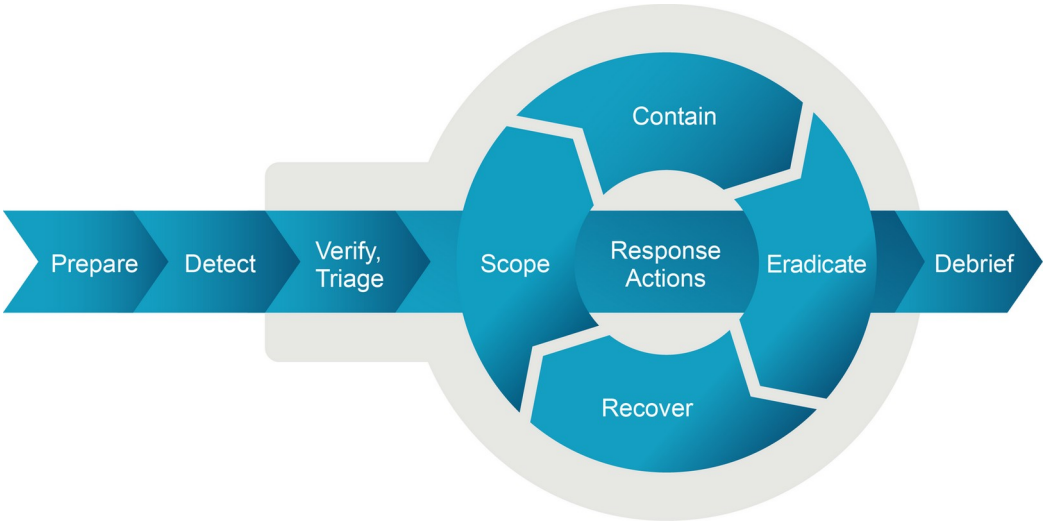

Figure 22. The source model. Detect becomes Identification so the letters match PICERL, and the four activities in the ring are grouped under Response. Nothing else changes.

Why the loop is there

Figure 23 draws four activities as steps: isolate and gather evidence under contain, analyze and remove persistence under eradicate, each running once and in order.

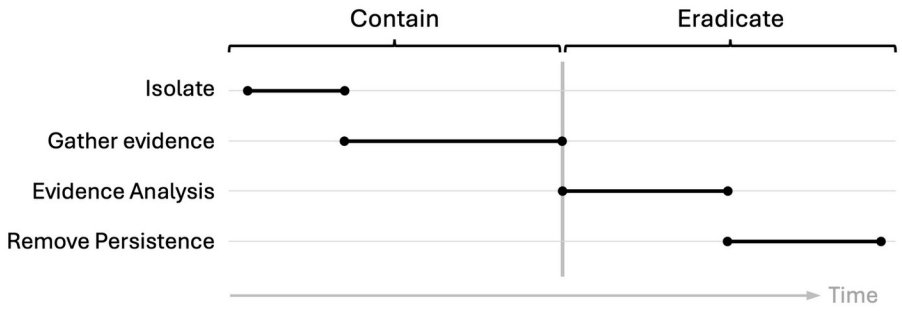

Figure 23. The stepped view.

Figure 24 draws the same four as they actually run. Isolation happens three times, evidence gathering starts and stops four times, and analysis runs underneath it all from early on. Plan for the second picture. Status reporting built on the first one contradicts itself by the third update.

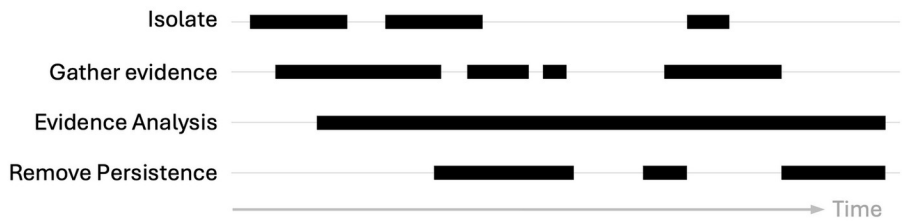

Figure 24. The same work, as it actually runs.

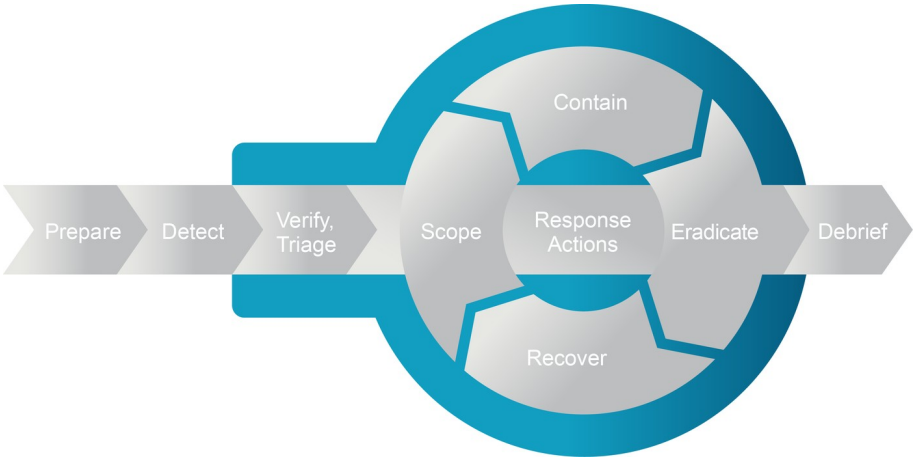

Figure 25. The highlight covers the whole model, not one stage of it.

Where the model came from

PICERL is the cycle most responders already know: preparation, identification, containment, eradication, recovery, lessons learned. All of it is here.

What PICERL leaves implicit is verification, triage and scope. They happen in every incident whether or not a model names them, and the unnamed activity is the one that gets skipped under pressure.

Scope, contain, eradicate and recover share the R because they repeat. Giving each its own letter would claim they run once each in order, which is the mistake this model exists to correct.

Letter	Phase	What it covers	In PICERL
P	Preparation	Readiness, authority, tooling, exercises, hunting	Preparation
I	Identification	Alerting, detection and hunting, all running continuously	Identification
V	Verification	Establish that the event is an incident requiring a response	Named here. Left implicit.
T	Triage	Prioritize the incident against organizational objectives	Named here. Left implicit.
R	Response	Scope, contain, eradicate and recover, repeating until scoping produces nothing new	Containment, Eradication, Recovery, plus scope named here
D	Debrief	Close the incident and convert it into changes	Lessons Learned

Verification and triage run together because both need the same conversation with decision makers. The hyphen marks where the incident ends and the closeout begins, and it stops the letters from being read as a word.

Start, work, finish, record

Activity	Starts when	Core work	Done when	Record produced
P. Preparation	No incident running	Policy, authority, playbooks, backups, tooling, exercises, hunting	Never closes. Reopened by every debrief.	Authority matrix, playbooks, asset and contact lists
I. Identification	Telemetry and reports flowing	Operate detections, work alerts, hunt on a cadence	An event of interest is raised	Alert or hunt finding with evidence references
V, T. Verification and triage	Event of interest raised	Confirm it requires a response, rank it against organizational objectives	Continue, stop, or defer, with a name against it	Incident record, triage brief
R. Response (scope)	Verified incident, one or more indicators	Derive indicators, enterprise hunt, progressive sweep, attack timeline	Breadth, systems involved and potential impact established	Scope findings, attack timeline
R. Response (contain)	Reach of the compromise clear enough to act	Stop attacker activity, preserve evidence by volatility	Three validation channels show activity stopped	Containment decision record
R. Response (eradicate)	Evidence collected and preserved	Short-form investigation, remove persistence, rotate credentials, patch	Rerunning the original enumeration comes back clean	Root cause, eradication log
R. Response (recover)	Eradication verified	Verification checks, validation testing, owner sign-off, restoration	Systems in production under enhanced monitoring	Pre-restoration sign-off, owner acceptance

Activity	Starts when	Core work	Done when	Record produced
D. Debrief	Loop exited and signed	Close temporary assets, consolidate records, metrics, after-action review	Improvement plan issued with owners and dates	Incident report, metrics, improvement register

2. Decide these before the incident

Settle these in advance and refer to them during the response. Working out who can authorize taking production offline while the file shares are encrypting spends the only time that matters on a phone tree.

Who authorizes what

Structure from p. 97. The role names and containment tiers are placeholders. Replace them with yours before anyone relies on this.

Decision	Authority	Delegate	Record required
Declare an incident	Incident commander	Senior analyst on duty	Ticket with declaration time and rationale
Isolate an endpoint	SOC or IR analyst	None needed	Containment decision record
Isolate a server or service	Service owner with incident commander	Incident commander alone if the owner is unreachable inside the response window	Containment decision record with impact assessment
Shut down production, or anything touching regulated services	Named executive	CISO after hours, ratified next business day	Containment decision record plus executive approval
Expand scope to other business units	Incident commander, CISO notified	None	Scope expansion entry with the trigger
Exit the loop, technical closure	Incident commander on eradication validation	Senior analyst with written approval	Signed technical closure record
Accept residual risk	CISO or named risk owner	Board-level executive for significant incidents	Signed risk acceptance naming the risks
Notify a regulator	Legal counsel with the CISO	General counsel	Decision log with timeline and regulator named
Speak to media or the public	Communications lead with legal sign-off	CEO for material incidents	Approved statement under version control
Pay a ransom	Written position agreed in advance, then named executive with counsel	None	Decision log naming counsel and carrier consulted

When containment cannot wait

Table 17, p. 227. Map each trigger to an approved action during preparation.

Trigger	Indicators	Approved action
Active destruction	Encryption spreading across shares, event logs wiped, database tables dropped	Isolate immediately

Trigger	Indicators	Approved action
Exfiltration	Large transfers to cloud storage, bulk queries against personal data, archives sent to external sites	Contain urgently to cap exposure
Safety or critical systems	ICS, medical devices, payment processing, other safety-critical infrastructure	Act immediately and accept the loss of intelligence
Regulatory clock	Data or services under HIPAA, PCI DSS, GDPR, NIS2, DORA, CRA, or a local equivalent	Contain rapidly to limit affected records

Which containment posture

Table 16, p. 223. Deceptive containment is drawn from the same chapter's step-by-step reference.

Situation	Posture
Destruction or ransomware underway	Active, required
Regulated data with a compliance clock	Active, required
Safety-critical or high-availability system at risk	Active, required
Attacker shows awareness of defenders	Active, strongly favored
Short dwell time, limited monitoring maturity	Active, strongly favored
Extended dwell time, mature monitoring	Adaptive, strongly favored
Unknown scope, uncertain sophistication	Adaptive, preferred
Commodity malware on a low-value asset	Passive, possible
Controlled deception environment	Passive or deceptive, possible

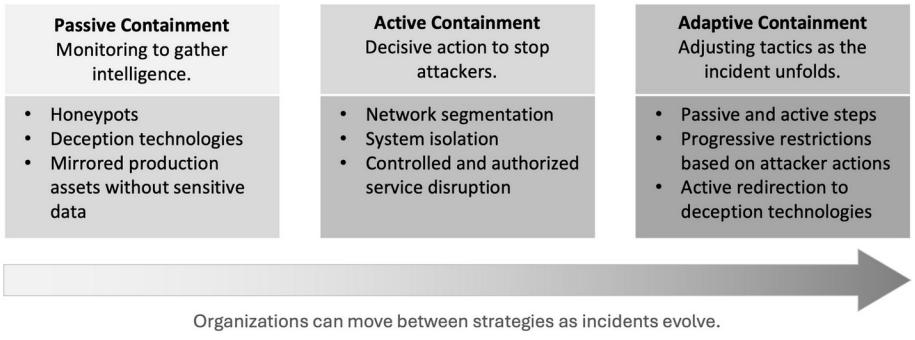

Figure 75. The arrow matters. Posture changes within a single incident as understanding improves.

Risk levels and what they trigger

Levels from p. 171. The trigger column is mine, and the thresholds should be yours.

Level	Meaning	What it triggers
Low	Minor anomaly, minimal impact	Analyst handles it, logged only
Medium	Worth investigating, operations unaffected	Incident commander informed, daily update
High	Confirmed incident affecting operations or data integrity	Team activated, decision makers briefed, update cadence set
Critical	Major incident	Executive escalation, immediate containment authority granted, counsel engaged

3. Phase playbooks, P to D

P. Preparation

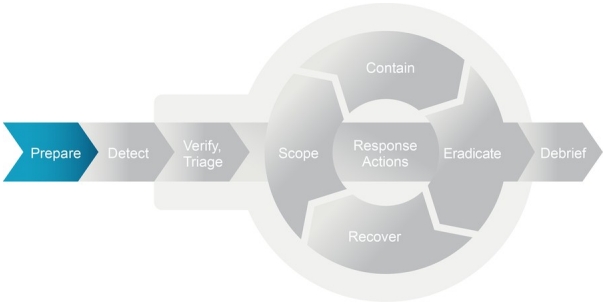

Figure 29

STARTS WHEN	DONE WHEN
No incident running. This is the default state.	Never closes. Every debrief reopens it with named owners.

Most of the team's hours belong here, and the work is hardest to fund when it is working, because nothing has gone wrong. Measurements make the case where a maturity model cannot: time to detect, how fast the last drill contained, what the last tabletop exposed.

Capability is what gets used. An intelligence feed nobody queries, a console nobody opens, and a playbook last edited by someone who has since left do not become useful when the incident starts.

I. Identification

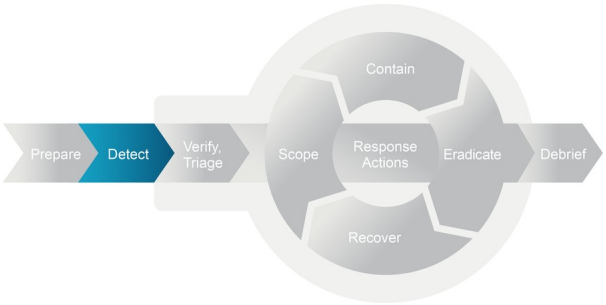

Figure 50

STARTS WHEN	DONE WHEN
Telemetry and reports flowing. Runs continuously.	An event of interest raised with evidence references attached.

Identification covers alerting, detection and hunting. Alerting is what reaches you, detection is the engineered logic behind it, and hunting is what finds whatever neither one caught.

Only hunting moves time to detect. Alerting and detection both mean something else found it first, and you inherit whatever dwell time that cost.

Layer the methods, and map detection logic to MITRE ATT&CK so the rules survive a change of tooling. Final classification stays with an analyst, because no model knows about the migration that started on Monday.

Method	Strong on	Weak on
Signature	Known malware and patterns at machine speed, few false positives	New variants, zero days, and any attacker who has read the signature
Behavioral	Mass encryption, off-hours access, living off the land, insider misuse	Depends on a tuned baseline. One company-wide baseline hides anomalies.
Machine learning	Subtle patterns, triage and prioritization	No organizational context, so the final call stays with a person
Hybrid	Layered coverage that survives tool changes	Cost and complexity of operating several platforms together

Spend hunting hours above the automation line, on the artifacts an adversary cannot cheaply swap out. Hashes and addresses burn on contact, and automation already covers them at machine speed.

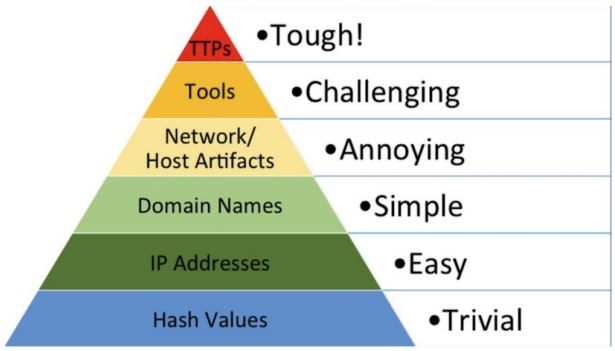

Figure 52. The Pyramid of Pain, David Bianco, 2013.

Layer	Cost to change	Who chases it
TTPs	Tough	Analysts. Methods the adversary cannot abandon.
Tools	Challenging	Analysts, with detections built on tool behavior
Network and host artifacts	Annoying	Analysts, on beaconing, process chains, lateral movement
Domain names	Simple	Automation, with analysts pivoting from the hits
IP addresses	Easy	Automation
Hash values	Trivial	Automation. High confidence, short shelf life.

Watch out

- A single baseline averaged across business units defines normal too broadly for anomalies to surface. Scope baselines per business unit or per function.
- Hunt from a catalog, with a hypothesis, a named data source, and a cadence. Findings from unstructured hunting are hard to repeat and harder to defend.

V and T. Verification and triage

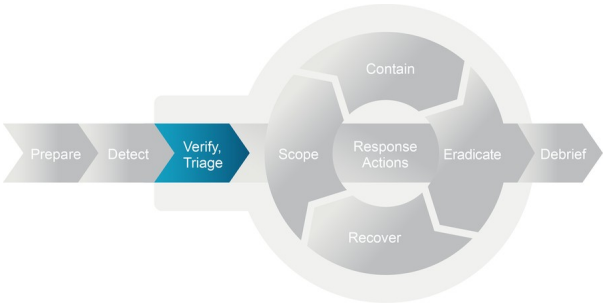

Figure 57

STARTS WHEN	DONE WHEN
An event of interest with evidence references.	Continue, stop, or defer, with a name and a time against it.

Verification establishes that the observed event is an incident, and one that requires a response. Triage ranks it against organizational objectives so the team works the most important incident first, which matters most when the impact is broad.

Verify before engaging the wider team. A declared incident has a cost of its own, pulling people off other work and slowing the systems under investigation, and that cost is worth paying only for real ones.

They run concurrently because both need the same insight and coordination from decision makers, and the outcome of both is the right resources on the right incident.

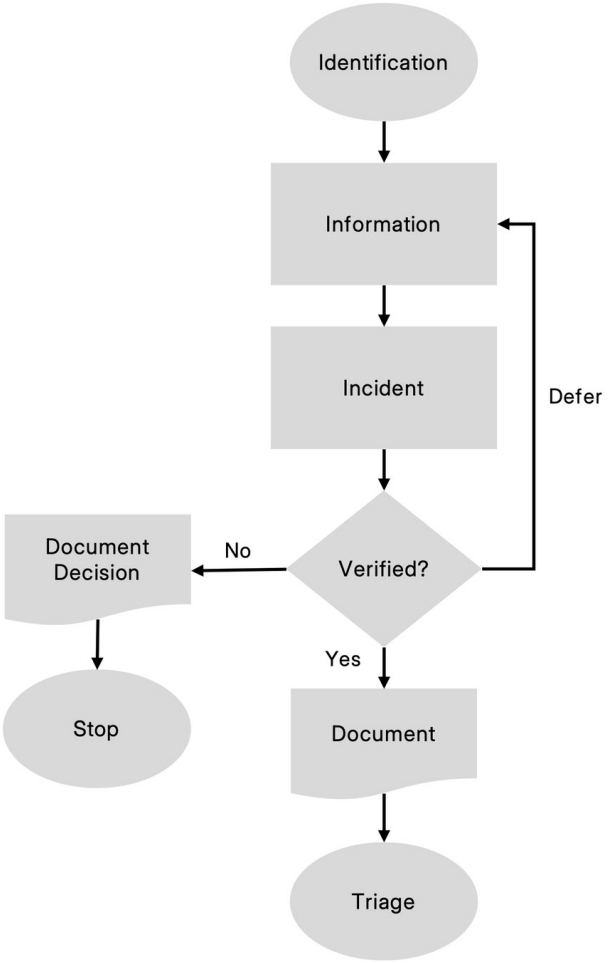

Figure 58. Deferring is a real outcome, and recording it is what keeps it from becoming a dropped ticket.

Outcome	Use when	Then
Continue	Real, and it warrants a response	Escalate, engage the team and decision makers, triage
Stop	Not real	Close it and document why, for the next analyst who sees the same thing
Defer	Available information cannot settle it	Request more from the reporter or other sources, then decide again

Watch out

- Documentation written after the fact is unreliable, and exact commands and timestamps are the first things to go. Open the record at verification and add to it as the response proceeds.

R. Response (scope)

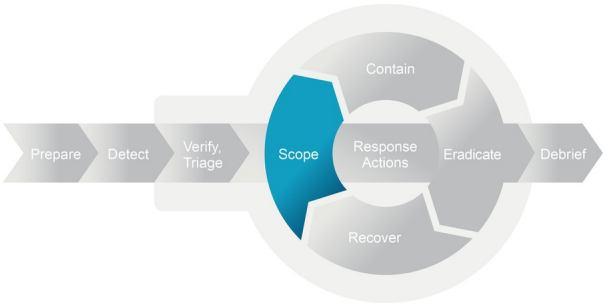

Figure 69

STARTS WHEN	DONE WHEN
Verified incident with at least one indicator.	Systems, accounts, and data enumerated, with visibility gaps named.

Scope establishes the breadth of the incident, the systems involved, and the potential impact on the organization. It starts from what identification and verification produced, works out the attacker's TTPs, derives indicators from them, and uses those indicators to find every system in scope.

Do it before cleaning anything. Cleaning only the systems you already know about leaves the attacker the rest, which is how the same incident gets declared twice in a quarter.

Scope owns the attack timeline. Eradication annotates it, the debrief finalizes it, nobody else redraws it.

Phase	Targets	Sources
Critical assets first	Domain controllers, file servers, databases, anything holding sensitive data, regardless of where the first indicator appeared	SIEM, EDR, authenticated scans
Lateral expansion	Same segment, systems authenticated to by compromised accounts, systems sharing administrative credentials	Authentication logs, EDR, network flow
Environmental sweep	Everything remaining. Days or weeks at scale.	EDR, inventory tools, scripted checks, compromise assessment tooling

Watch out

- In OT the visibility gap is the normal condition. Controllers do not log centrally, agents cannot be installed, and the maintenance window is next quarter. Plan on network flow as the primary evidence source, and write that into the response plan before it is needed.

R. Response (contain)

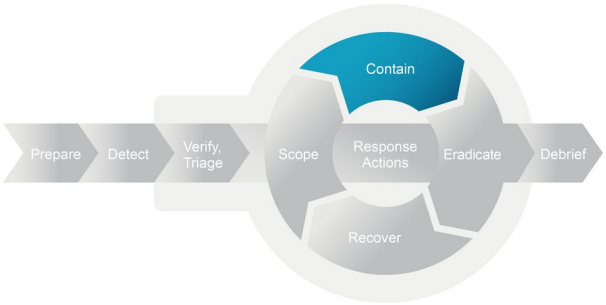

Figure 74

STARTS WHEN	DONE WHEN
Scope clear enough to act, posture chosen, authority in hand.	Network, process, and log channels all show attacker activity stopped.

Containment stops the attacker and preserves the evidence.

Cut the specific paths the attacker is using. Pulling the host stopped being a general strategy once estates went distributed. A compromised web server isolated from internal systems while it keeps serving the internet is contained, and the site is still up.

Network	• VLAN Isolation • Firewall Rules • DNS Sinkholes
Host	• EDR Isolation • Local Firewall • Process Restrictions
Application	• Web Application Firewall • App-specific Controls • Revoking Automations
Identity	• Credential Revocation • Session Termination • Conditional Access

Figure 77. Containment layers.

Layer	Use it for
Network	Broad, fast control over many systems at once
Host	Precision on a known-compromised endpoint
Application	Attacker tooling that network controls do not see
Identity	SaaS and cloud paths, where it is frequently the only layer that reaches

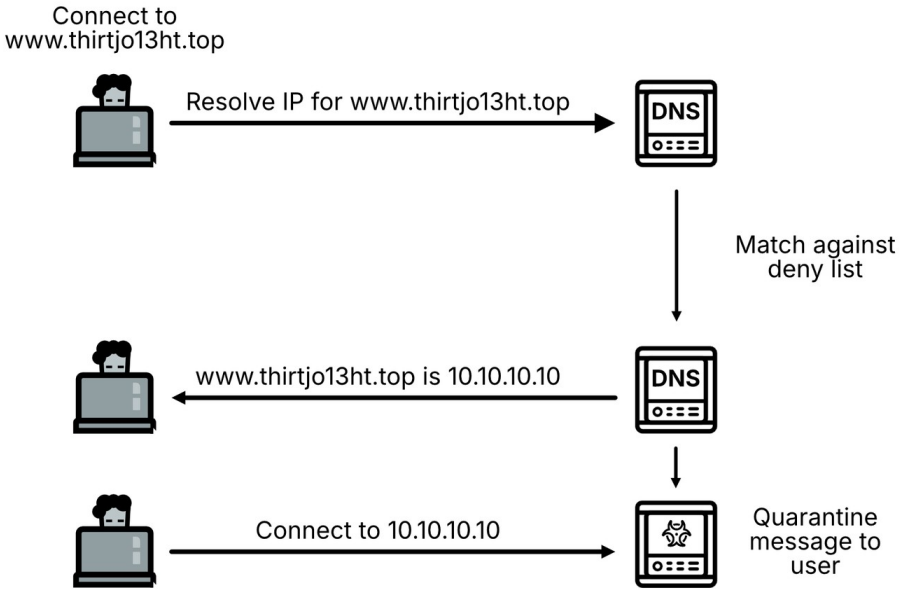

Figure 79. DNS sinkholing redirects command and control traffic to a controlled destination.

The best-known sinkhole in the field

Reporting by WIRED, TechCrunch and the Associated Press; sinkhole data published by Kryptos Logic

WannaCry, 12 May 2017. The worm was spreading across Europe and the NHS was turning patients away. Marcus Hutchins, working for Kryptos Logic from England, found an unregistered domain inside a WannaCry sample and bought it on Namecheap for $10.69, intending to sinkhole the traffic and count infected hosts. The malware checked whether that domain resolved before encrypting anything and halted when it did. Another researcher reverse-engineered the check and confirmed the behavior. Spread of that variant stopped the same afternoon.

The technique was ordinary DNS sinkholing, the same mechanism in Figure 79, applied at internet scale by accident. Hutchins told the Associated Press at the time that he was not a hero and was just doing his part to stop botnets.

It eradicated nothing. Variants with different killswitch domains appeared within days, unpatched hosts were still reaching the sinkhole years later, and Boeing, Honda, Connecticut state agencies and Victoria state police all reported infections during 2018. A sinkhole buys time to patch. The patching still has to happen.

Watch out

- Human-operated ransomware groups have accelerated encryption after detecting defensive action, and watchdog processes respawn malware or trigger destructive payloads. Move quietly until the adversary is understood.
- Sequential isolation signals the response and gives the adversary time to escalate. Execute simultaneously across the network, endpoint, identity, and application teams.
- Write the rollback procedure before changing firewall rules or network topology. It will be requested at the worst possible time.

R. Response (eradicate)

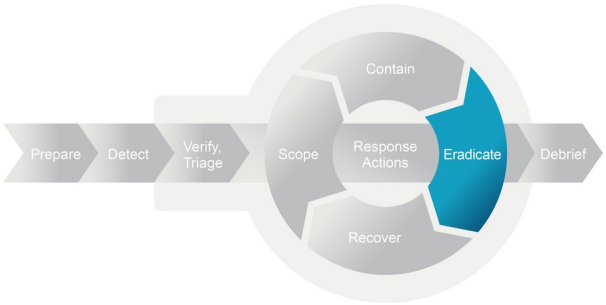

Figure 91

STARTS WHEN	DONE WHEN
Evidence collected and preserved.	Rerunning the enumeration that found the problem comes back clean.

Contained systems cost money every hour they stay down, and reading disk and memory properly takes days or weeks, so the investigation runs on two tracks at once. Short-form work answers enough to act on. Long-form work carries on for the full TTPs, regulatory reporting, and anything that ends up in front of a court.

The short-form track establishes what the attacker did, how they got in, where else they might be, and what has to go so they cannot come back the same way. The output is rough by design. The source book's tech editor calls it a dirty kebab, which is a fair description of a legitimate deliverable.

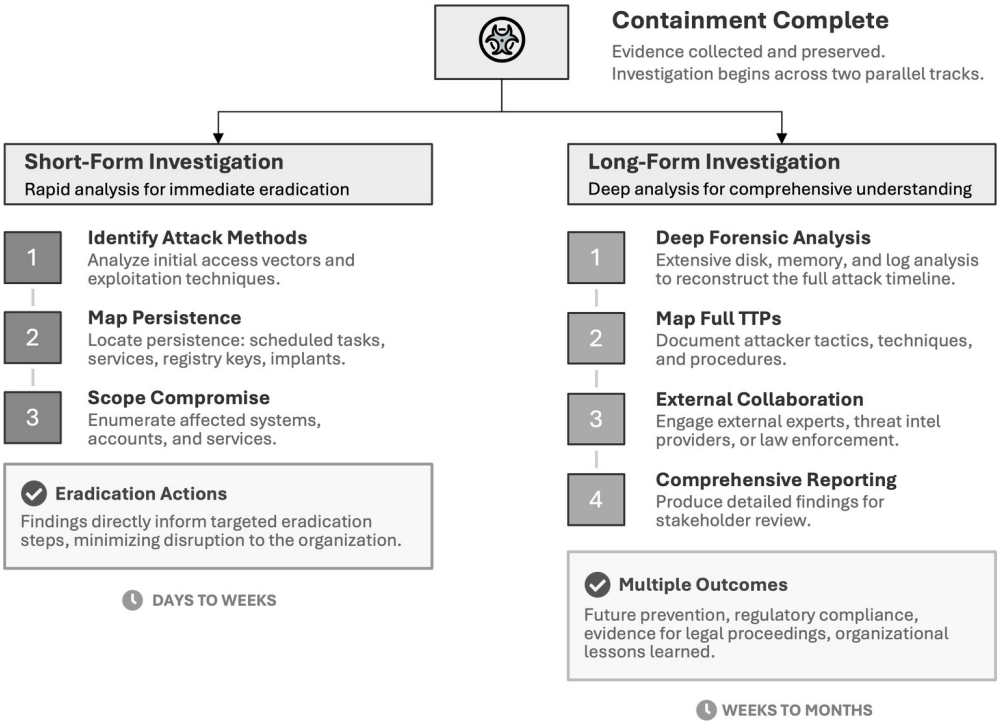

Figure 92. Agree the scope of both tracks with decision makers, in writing.

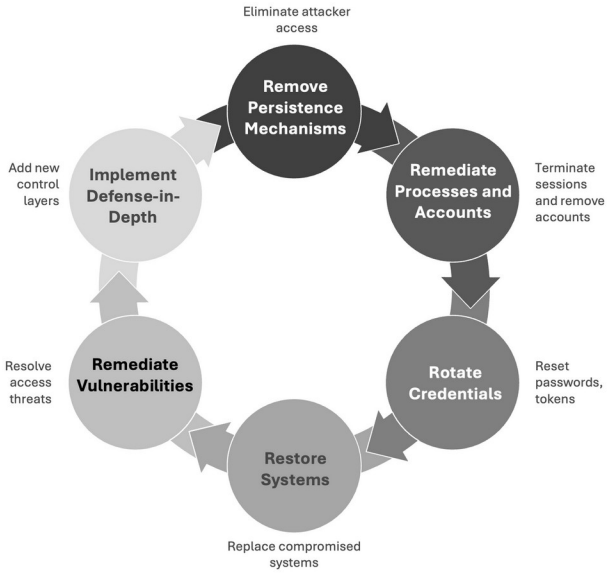

Figure 108. The six categories of removal work.

Credential reset phasing

Structure from Table 31, p. 369. The windows come from the book's worked example and need replacing with yours.

Phase	Scope	Window in the example
1	KRBTGT, reset twice, then Tier 0 accounts	Friday 11 PM to Saturday 3 AM, scheduled maintenance
2	Tier 1 accounts and critical service accounts	Saturday 6 AM to noon, reduced trading hours
3	Tier 2 accounts and remaining service accounts	Saturday 2 PM to Sunday 6 PM

Watch out

- The double KRBTGT reset depends on the interval between resets. Reset once, confirm replication across all domain controllers, then wait longer than the maximum ticket lifetime before the second reset so any existing golden ticket expires. The example allows ten hours.
- Track every service account with its dependent applications and owner, and have the application team test the credential update in staging before the window. Expect the backup service account to be the one that breaks.
- Keep emergency access available. Issue short-validity temporary accounts to on-call staff whose credentials are being reset.

R. Response (recover)

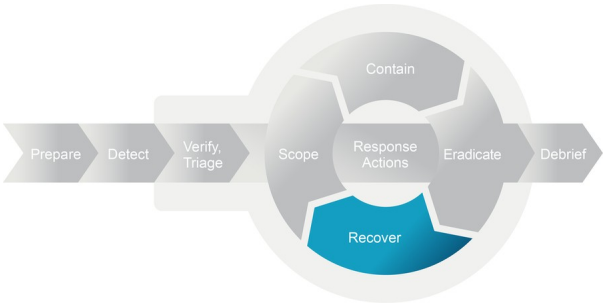

Figure 122

STARTS WHEN	DONE WHEN
Eradication verified.	Systems in production under enhanced monitoring, owner sign-off on file.

Work the checks in order before anything returns to production. The indicator scan comes last, because scanning a half-patched box tells you nothing about the build that ships.

D. Debrief

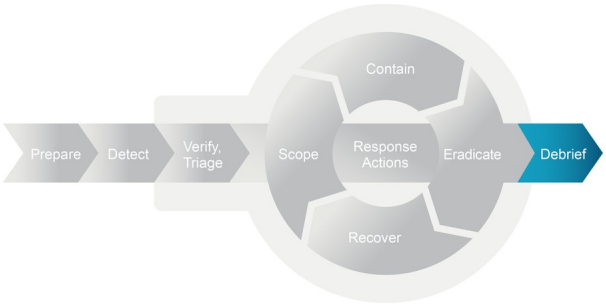

Figure 131

STARTS WHEN	DONE WHEN
Loop exited, technical closure and risk acceptance both signed.	Improvement plan issued with named owners and milestone dates.

This is the phase other traditions call lessons learned, the post-mortem, or root cause analysis. Same activity, whichever name your organization uses.

Closing the incident and turning it into changes are one activity. Funding stays available for about a month after closure, while the board still remembers being frightened, so arrive with the improvement plan already costed and an owner and a date on every line.

Running the response loop

Put the same question to every new fact: does this change scope, containment, eradication or recovery? If it does, the team is on another pass. Say so on the bridge, because silence reads as progress and stakeholders start planning the all-clear.

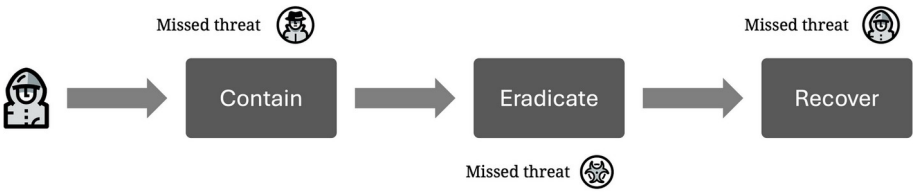

Figure 67. Three threats pass through unnoticed.

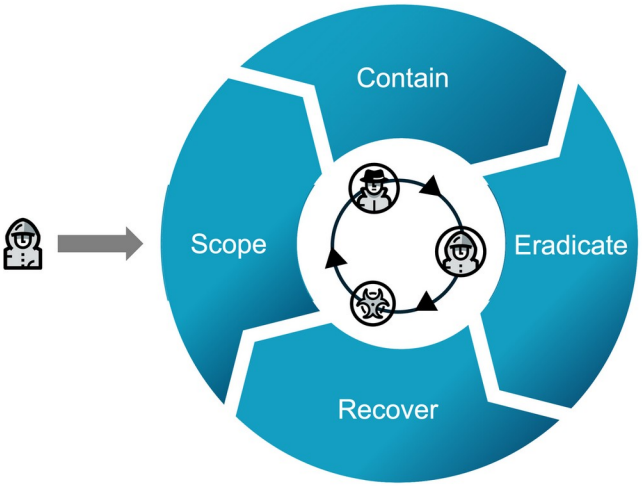

Figure 68. The same three threats, caught across successive passes.

A straight line through contain, eradicate and recover misses whatever only becomes visible once the evidence has been read. That is the difference the two figures are drawing.

Trigger	Looks like	Next move
New indicators	A PowerShell script in a registry key, traffic to unknown C2	Back to scoping with the new indicator
Scope expansion	A known compromised service account also active in a regulated system	Another scoping pass, with new stakeholders and new obligations
Incomplete eradication	A system that looked clean shows reinfection	Hunt the missed persistence or the unidentified access vector
Analysis revelations	Memory forensics finds a second malware family	Rescope, then revisit containment and eradication
Business change	A new obligation or a shift in priorities	Re-triage and reset the plan with decision makers

Watch out

- Rotate people and schedule rest. Fatigue accumulates across iterations and rarely appears in any report.
- Documentation has to survive shift handover, or incoming analysts repeat work or skip it.
- Name the person who authorizes the exit. An attributed written decision holds up if the incident is reopened or reviewed, and people commit more carefully when their name is on it.

4. Templates

Each record copies to the clipboard.

5. Metrics

Definitions from pp. 445 to 446. Iteration count is mine.

Metric	Measured from	This incident
Mean time to detect	Incident start to detection	
Mean time to respond	Detection to resolution	
Time to containment	Detection to attacker activity stopped	
Time to eradication	Containment to persistence removal complete	
Time to full recovery	Detection to all systems restored	
Total lifecycle	Initial compromise to verified resolution	
Systems affected	Servers, workstations, cloud resources	
Accounts affected	Compromised or requiring reset	
Data exposure	Record count, classification, regulatory categories	
Business impact	Disruption duration, revenue effect	
Effort and cost	Personnel hours, external services, tooling, any ransom paid	
Iterations	Passes through the response actions loop	

Record these the same way every time. Comparison across incidents is what shows whether the team is improving, and the numbers are what fund the improvement plan. Iteration count is here because it shows whether scoping caught things early or late, and nothing else in the set does.

6. Attribution

The model is adapted from *Dynamic Incident Response* by Joshua Wright, published by the SANS Institute in 2026 under CC BY 4.0. The numbered figures are reproduced from it. Changes were made.