

Detection, Triage, and Containment: The First Hour

Fri, Sep 11, 6:00 PM EDT · <https://scottslab.io/posts/detection-triage-and-containment>



TL;DR

Detection surfaces from tooling (IDS/IPS, firewalls, auth systems, the SIEM, scanners, logs, netflow, anti-malware) and from people. NIST's own guidance notes that an organization can face thousands, even millions, of sensor alerts a day, which is why triage, not detection, is the actual bottleneck. NIST's real severity taxonomy scores functional impact, information impact, and recoverability effort, and it's that combination, not a gut call, that decides whether an incident gets a first responder or the whole team. Containment then escalates from segmentation to isolation to full removal, weighed against the six criteria from NIST's SP 800-61r2, the sharpest of which is the containment-versus-evidence tension from the IR plan made concrete.


Detection is where response starts, and the signals come from two directions. From your tooling: IDS/IPS, firewalls, authentication systems, the SIEM, vulnerability scanners, event logs, netflow, and anti-malware. The SIEM does the heavy lifting of aggregating all that log data and surfacing potential incidents through rules and correlation. And from people: plenty of incidents are first reported by a customer, an employee, or a third party who noticed something wrong before any alert

fired. A mature program watches both channels, because the tooling misses things and the human report is often the earliest warning you'll get.

The bottleneck isn't detection

NIST draws a useful line between two kinds of signal. A precursor is a sign an incident *might* happen: a vulnerability scan hitting your web server, a public exploit announcement for software you run, a threat actor's stated intent. Precursors are rare, and when you catch one you get something valuable: a chance to harden before anything happens. An indicator is a sign an incident is happening or already has: a failed-login spike, an antivirus alert, a filename with characters that shouldn't be there, a logged change to audit settings. Indicators are common, and that's the whole problem: NIST's own guidance points out that a single organization can face thousands, or even millions, of sensor alerts in a day. Almost none of them are an incident. The job of triage is compressing that flood down to the handful worth a human's sustained attention, and *that* compression, not the number of sensors you own, is what actually limits how fast you respond. A SIEM that ingests everything and a team that can investigate three things a day are the same organization with two very different bottlenecks, and only one of them is the one people usually spend money on.

Once something clears that bar, triage decides how hard you hit back, and NIST gives you a real taxonomy rather than a vibe. Functional impact is scored none, low, medium, or high, tied to concrete definitions: none means every service is available to every user; low means you're still delivering critical services but have lost efficiency; medium means a subset of users has lost a critical service; high means no one can get a critical service at all. Information impact is scored separately and isn't mutually exclusive: none, privacy breach, proprietary breach, integrity loss can all apply to the same incident. And recoverability effort (regular, supplemented, extended, or not recoverable) captures how much outside help and time the fix will cost. The reason this matters more than a simple low/moderate/high label: an incident with high functional impact but low recovery effort is exactly the kind of thing worth throwing immediate resources at, while a low-impact incident with an extended recovery effort might sit and wait. That combination is the actual decision engine underneath tiers like "a low-severity incident is one first responders resolve independently, no after-hours mobilization" up through "a high-severity incident means immediate full response and senior executive notification."

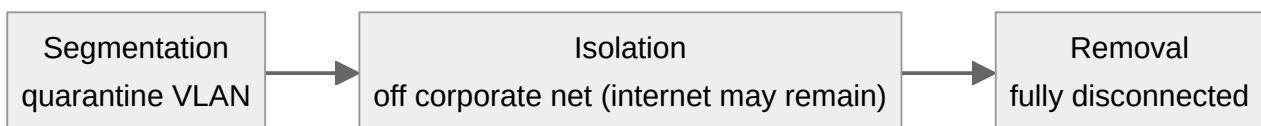


The bottleneck isn't detection

Getting the tier right matters in both directions. Over-escalate and you burn out the team on false alarms. Under-escalate and a real breach smolders because nobody treated it as urgent. That's also why declaration authority shouldn't be diffuse. Most organizations let any first responder or analyst declare and handle a low- or moderate-tier incident on their own judgment. Crossing into high is different: that declaration is the trigger for executive notification, possibly legal counsel and comms, and it should require someone with the authority to actually mean it: the incident commander or a duty manager, not whoever happened to be on shift when the alert fired.

Contain only as hard as the situation demands

The first responder's priority isn't to solve the incident. It's to stop the bleeding: containing the damage and isolating affected systems before it spreads. Containment escalates in three steps. Segmentation moves the compromised systems onto a quarantine VLAN: contained but still reachable for investigation. Isolation disconnects them from the corporate network, though internet access may remain, sometimes deliberately, to watch the attacker. Removal is full disconnection from all networks: the hard stop. You climb that ladder only as far as the situation demands: a single infected laptop and a domain controller with an active intrusion do not get the same first move.



Contain only as hard as the situation demands

Which rung you choose isn't a gut call. NIST's SP 800-61r2 gives you six criteria to weigh (the 2025 Rev 3 folds these into broader risk-based language, but the six are still the practical checklist): the damage potential if you don't act, evidence preservation, service availability, the time and

resources the strategy takes, its effectiveness, and its duration. That's the containment-versus-evidence tension from the IR plan, made concrete: a ransomware outbreak spreading laterally probably justifies fast isolation even at the cost of volatile memory you'd otherwise want to capture, while a quiet, contained reconnaissance intrusion you're trying to build a legal case around might justify watching a little longer to preserve evidence and attribution. The fastest containment and the best evidence are rarely the same move, and these six criteria are how you decide which one wins for *this* incident.

Written by Scott S. — Contact: scott@scottschlangen.com — scottslab.io — <https://scottslab.io/posts/detection-triage-and-containment>