

Building and Staffing the Incident Response Team

Fri, Sep 11, 10:00 AM EDT · <https://scottslab.io/posts/building-and-staffing-the-ir-team>




TL;DR

An IR team has to be reachable 24/7 with primary and backup coverage, because incidents don't wait for business hours. It's not just security people: management, cybersecurity, legal, HR, public affairs, technical SMEs, and physical security all belong, and NIST gives you real structural choices: central versus distributed teams, in-house versus outsourced staffing. One role deserves its own line: whoever commands the incident must not be the same person elbows-deep in fixing it. Train regularly, watch for burnout, and if you lack forensic capability in-house, pre-negotiate an external retainer before you need it.

A plan without people to run it is paper. The first requirement for an IR team is availability: it has to be reachable 24/7, with primary and backup coverage for every role, because attackers keep worse hours than you do and "we'll deal with it Monday" is how a contained incident becomes a breach. Backup coverage matters as much as primary. The incident that lands while your one firewall expert is on a plane is exactly the one you planned for by naming a second.

The composition surprises people who think IR is a security-team activity. It's a cross-functional group. Management, to make and own the hard calls. Cybersecurity personnel, obviously. Legal, because evidence handling, disclosure obligations, and liability are live from minute one. HR, when an insider or an employee is involved. Public affairs, because how you communicate a breach can

matter as much as how you fix it. Technical subject-matter experts: DBAs, developers, sysadmins, virtualization specialists, because the person who knows the compromised system is the person who can actually reason about it. And physical security, because plenty of "cyber" incidents have a badge, a door, or a stolen laptop somewhere in the story.



Building and Staffing the Incident Response Team

Central, distributed, or borrowed

NIST lays out real structural choices here, and picking the wrong one is a slow-motion mistake you only notice mid-incident. A central team, one group handling every incident across the whole organization, works well for a small organization or one without much geographic spread. A distributed model gives large or multi-site organizations a team per division or region, closer to the systems they're defending, but NIST is explicit that those teams still have to function as one coordinated entity: shared process, shared information, because more than one of those teams will eventually see pieces of the same incident and needs to recognize it as the same incident.

Staffing follows a similar fork. Fully in-house means your own people do everything, with contractors filling technical or administrative gaps. Fully outsourced puts an onsite contractor in the seat full-time, which only really works when you have employees senior enough to supervise the work. You can't outsource accountability. In between sits the arrangement most organizations actually run: partially outsourced, where 24/7 monitoring of the sensors, firewalls, and other security tooling goes to a managed security services provider that watches the noise around the clock and escalates the real incidents to your internal team. That's not a lesser option. It's the standard pattern for organizations that can't justify a full internal SOC but still need eyes on the environment at 3am.

The person in charge isn't the person with their hands on the keyboard

One role belongs on every IR team regardless of size, and it's easy to skip because nobody explicitly assigns it: an incident commander. That's the person who decides what gets communicated, whether a system comes offline, when legal or the executives get pulled in, and when the incident is declared over. It cannot be the same person who's actively working the compromised host, and not because that person isn't capable of both. Command and hands-on remediation are two different cognitive jobs, and doing them at the same time means doing both worse at the exact moment neither can afford it. The engineer buried in a shell on a compromised server should not also be the one

deciding whether to notify the board. Name the commander in advance, the same way the plan names who can approve taking a revenue system offline.

Train before it's real, and watch for burnout

Two habits separate teams that perform from teams that panic. First, train and exercise regularly. Run the tabletop and the drill when nothing is wrong, so the real incident isn't the first time these people have worked together. For organizations that can't justify a full-time team, NIST's own comparison is a volunteer fire department: a virtual team, mostly on other duties, that gets contacted and assembles rapidly when something happens, with an existing group like the help desk trained to do first-pass triage before escalating. That model works, but only if the drills happen anyway. A volunteer team that has never rehearsed is just a group of strangers under stress with a slower callout.

Second, be honest that incident response work is genuinely stressful, and stacking on-call rotations on top of a day job burns people out, which is exactly how organizations end up short-staffed on the one skill they can't improvise. Segregate roles and rotate them deliberately rather than leaning on the same two people every time. And be honest about your gaps: if you don't have real forensic capability in-house, retain an external IR provider ahead of time. Pre-negotiate the contract *before* an incident occurs. The moment you're breached, your negotiating leverage evaporates, and the retainer you signed calmly last quarter is worth far more than the emergency engagement you're begging for at premium rates while the clock runs. Line up the help while you don't need it, so it's there the instant you do.