AEGYS

FIELD NOTE

What a mid-market security program really needs first

Where a fractional CISO spends the first 90 days, and why it isn't another policy binder.

A field note from AEGYS
aegys.io · secure@aegys.io

When a mid-market company brings in security leadership for the first time, the temptation is to start with a policy binder. We start somewhere less glamorous and far more useful: knowing what you have, who can touch it, and whether you would notice if something went wrong.

Start by knowing what you own

You cannot protect an asset you have not counted. The first weeks are an inventory, not an audit: the systems that run the business, the data that matters, the accounts with privileged access, and the vendors holding your data. NIST's Cybersecurity Framework 2.0 puts two functions ahead of all the defensive work, **Govern** and **Identify**, for exactly this reason. You decide who owns security decisions, then you build the map. Everything after that is guesswork until the map exists.

Do the boring controls that stop most incidents

Most breaches do not require an exotic defense. They are stopped by a short list of fundamentals done consistently. The CIS Critical Security Controls call this baseline **Implementation Group 1**, the set every organization should have regardless of size: a known inventory of hardware and software, multi-factor authentication on everything that faces the internet, centralized logging, tested backups, and prompt patching. Get those real and verified and you have closed the doors attackers use most. A mid-market program that nails IG1 beats one with a beautiful policy and no MFA.

Make compliance a byproduct, not the goal

Plenty of mid-market companies need SOC 2 because a customer asked. That is a fine forcing function, but the order matters. Build the controls because they reduce real risk, then let the SOC 2 report describe controls that already exist. Programs that chase the audit first end up with evidence of things nobody actually does, and the Trust Services Criteria reward an honest program while punishing a theatrical one at the worst possible moment.

Then, and only then, talk about AI

Leadership will ask about AI early, often before the basics are in place. The honest answer is that AI safety rides on the same foundations: you cannot govern AI use without knowing your data, your identities, and your logging. Once IG1 is real, adding AI is a small extension, not a separate program. We cover that in the Practical AI Playbook.

The first 90 days, in one line

- Weeks 1 to 4: govern and identify. Own the decisions; inventory assets, identities, data, and vendors.

- Weeks 5 to 8: close the IG1 basics. Multi-factor authentication, logging, backups, patching.
- Weeks 9 to 12: pick the compliance target if there is one, and a first measurable risk to retire.

> **The point**
> A security program earns trust by reducing real risk in a visible order, not by producing a binder. The fundamentals are unglamorous, and they are what works.

This is the shape of a Security and AI Posture Assessment: where you stand, what to fix first, and a roadmap you can act on. Reach us at secure@aegys.io.

References

NIST Cybersecurity Framework (CSF) 2.0, NIST CSWP 29. https://doi.org/10.6028/NIST.CSWP.29

CIS Critical Security Controls v8, Implementation Group 1 (IG1). https://www.cisecurity.org/controls

AICPA SOC 2, Trust Services Criteria. https://www.aicpa-cima.com