

Hardening an Application: The Boring Controls That Actually Hold

Mon, Sep 28, 10:00 AM EDT · <https://scottslab.io/posts/application-hardening>

TL;DR

Hardening isn't one clever move, it's a short list of unglamorous controls done consistently: real authentication and least privilege, encryption at rest and in transit, server side validation on everything a user can touch, patches applied promptly from verified sources, and a configuration baseline so drift is visible instead of discovered during an incident. Insecure design gets fixed earlier than any of that, at the architecture stage. Obfuscation doesn't stop anyone determined, it just raises the toll.

Why this list and not a bigger one

Every hardening checklist I've read eventually converges on the same handful of items, and the ones that survive contact with a real audit are the boring ones. Not because nobody thought of anything cleverer, but because the clever controls tend to be brittle, expensive, or both, and the boring ones are the load bearing walls. Skip one of these and you don't get a slightly weaker app, you get a specific, well understood way in.


Identity, access, and what travels

Strong authentication and least privilege are the same control looked at from two directions: prove who someone is with enough confidence, then give that identity only what its job actually requires. An account with more access than its function needs is a standing liability whether or not it's ever misused, because the day it's compromised, everything it could touch is now in scope, not everything it should have needed.

Encryption at rest and in transit protects the same data at two different moments. In transit, TLS keeps a network observer from reading credentials, tokens, and payloads off the wire. At rest, encrypting the storage layer means a stolen disk, a misconfigured backup, or a database dump doesn't hand over plaintext for free. Neither substitutes for the other; a database encrypted at rest that talks to its application over plaintext, or an API served entirely over TLS but backed by an unencrypted disk, has only closed half the door.

Server side input validation is the one people think they've done because they added a form field pattern in the browser. Client side validation is a usability feature. It runs on a device the attacker fully controls, so it stops nothing determined. The check that actually matters happens on the server, on every field, on every request, treating anything that arrived from outside the trust boundary as

hostile until proven otherwise, because that's exactly what most injection and logic abuse turns out to be: input that should never have been trusted, trusted anyway.



Identity, access, and what travels

Patching, and trusting what you install

Patching promptly closes the gap between a vulnerability becoming public and it becoming someone else's foothold, and that gap is where most opportunistic exploitation happens: not novel zero days, known bugs against systems that hadn't gotten around to the fix yet. The part that's easy to skip is verifying what you're patching with. An update should be digitally signed and pulled from a source you trust, because an attacker who can slip an unsigned or tampered package into your update path doesn't need to find a vulnerability at all, they get to hand you one directly, wearing the shape of routine maintenance. Signature verification turns "we applied the update" into "we applied the update we actually meant to."

Baselines make drift visible

A configuration baseline is a documented, known good state for how a system should be set up, and its value isn't the document itself, it's what the document makes possible: comparison. Without a baseline, a setting that quietly changed six months ago looks identical to one that's always been that way, and you find out the difference during an incident, reconstructing history from logs if you're lucky. With a baseline, drift shows up as a diff, something concrete you can review and remediate quickly, before it's the reason something went wrong rather than a footnote in the postmortem explaining why.

Insecure design is a different layer entirely

Everything above hardens a system that's already been built a certain way. Insecure design is a problem hardening cannot reach, because it's baked into the shape of the thing before any of these controls get applied. Threat modelling, secure design patterns, and security reference architectures are how that gets addressed, at the point where the cost of changing course is a conversation instead of a migration. I've written separately about threat modelling and STRIDE in enough depth that I won't repeat it here, but the short version is worth restating: the earlier a design flaw is caught, the cheaper it is to fix, and hardening a bad design just makes it a more expensively defended bad design.

Obfuscation buys time, not safety

Obfuscation and camouflage, making source code and internal structure harder to read and reverse engineer, don't prevent an attack. They raise its cost. A determined, well resourced adversary gets through obfuscated code eventually; what changes is how long it takes and how much tooling it requires, which matters more than it sounds like it should. Time and cost are real deterrents against opportunistic attackers who move to an easier target rather than grind through deliberately obscured logic, even though they do nothing against someone with the resources and motivation to persist. Treat obfuscation as a delay tactic layered on top of real controls, never as a substitute for one.

What actually holds

None of this is exotic, and that's the point. The controls that hold under real pressure are the ones that get applied everywhere, all the time, without exception for the one endpoint someone forgot about or the one service account that got created in a hurry. Hardening fails less often because a control was unknown and more often because it was known, half applied, and assumed to be everywhere it wasn't.