

The *Drift* Audit

A standalone 2-week diagnostic engagement. Fixed-scope. Fixed-fee. Five reconciliation pairs. Usable without adopting any larger framework.

— 10-DAY EXECUTION SCHEDULE

10 days.

No wasted time.

Two weeks is the precise point where all five pairs can be examined with enough depth to produce actionable findings without over-scoping.

WEEK 1 Discovery & Investigation

Day 1

Kickoff & Inventory

Access provisioning, artefact inventory, scope confirmed in writing

Day 2

Hypothesis Formation

Spec read-through, "suspicious areas" shortlist

Day 3

Pair A Sweep

Code ↔ SSOT — counts, signatures, locations

Day 4

Pair A Deep + C Begin

Pair A frozen; Pair C (Schema ↔ Code) draft

Day 5

Status Call

Pair C complete; interim memo to client

WEEK 2 Analysis & Reporting

Day 6

Pair B Sweep

Decision Log ↔ SSOT — traceable decisions

Day 7

Pairs D + E

Immutable Contracts + Docs ↔ Reality

Day 8

Root Cause Analysis

RCA across all findings, severity classification

Day 9

Report Drafting

Remediation plan, Prevention Shortlist, exec summary

Day 10

Delivery

Final report + readout call with leadership

Five pairs. Drift is always *between two things*.

Click a pair to see exactly what to check during the audit.

A

B

C

D

E

Pair A: Code ↔ SSOT

- ✓ Function, table, and endpoint counts — verify actual vs SSOT claim
- ✓ API signatures — parameters, return types, endpoint shapes
- ✓ File locations — does code live where the SSOT says?
- ✓ Existence — does every referenced entity actually exist in code?

Every finding gets a *non-negotiable* severity.

Consistent classification makes the Drift Audit a credible diagnostic. When in doubt, classify up.

CRITICAL

Could cause a production incident, security breach, compliance failure, or silent data loss if shipped.

Code references an unfixed vulnerability the SSOT claims is resolved; auth flow calls a function that doesn't exist.

HIGH

Could cause a visible bug, failed deployment, or incorrect business decision in the next 30 days.

SSOT claims 19 endpoints; code has 23. Schema has a column used by 3 queries that was recently renamed.

MEDIUM

Creates confusion, slows onboarding, or increases risk of future drift without immediate failure.

Decision Log entries superseded but not marked. README setup instructions reference a deprecated package.

LOW

Cosmetic, stylistic, or informational inconsistencies with no operational impact.

Function names in docs use camelCase; code uses snake_case. Both work — docs just look inconsistent.

Cumulative rule: Five Medium findings in the same subsystem = one High finding about that subsystem. Novelty is not severity — a boring count mismatch that could ship wrong data outranks

an exotic drift pattern that cannot ship.

Fixed-scope. *Fixed-fee. Always.*

Hourly pricing destroys the value proposition. The Drift Audit's selling point is that it is bounded, predictable, and decision-ready in 2 weeks.

TIER 1

Starter

5 working days

One reconciliation pair, one subsystem, up to 5 findings.

Small teams, single-product SaaS, proof-of-value

MOST COMMON

TIER 2

Standard

10 working days

All five pairs, one product/subsystem, unlimited findings.

The default engagement

TIER 3

Extended

15 working days

All five pairs, multiple products/subsystems, interview-heavy.

Larger organisations

FOLLOW-ON

Retainer

1 day / month

Monthly Reconciliation Sweep using PL-08 methodology.

Post-audit, high-velocity teams

Starter \approx 50% of Standard · Extended \approx 150% of Standard · Retainer \approx 15% of Standard / month

— QUALIFYING CONVERSATION

When to say *no*.

These disqualifiers exist to be used. Decline engagements that match them — they produce bad outcomes and damage the methodology's reputation.

- ✘ Client expects working code changes, not a plan — that is a Remediation Sprint, priced differently
- ✘ Client expects the audit to exonerate or blame specific individuals
- ✘ No specification exists and none is planned — nothing to reconcile code against
- ✘ Client will not grant read access to code, SSOT, and decisions together
- ✘ Timeline demand is under 5 working days — cannot compress further without loss of rigour

Conversion rule of thumb: An honest audit converts roughly 30–50% of clients into at least one follow-on. A client who takes no follow-on is not a failure. Pushing too hard on follow-on sales cheapens the diagnostic.

Email Suneesh

· MiVA Education · suneesh.sidharthan@gmail.com · +91 9986898076

Released under [CC BY 4.0](https://creativecommons.org/licenses/by/4.0/)