From *Solo* to Production

The actual 6-month architecture-to-beta build that produced the methodology. Four drift incidents. Five model transitions. What broke, what worked, and what we'd do differently from day one.

THE PLATFORM

Three surfaces. One *PostgreSQL* backend.

MiVA is an education platform for Indian Class 10 students. The core thesis: 3–4 hours of daily Instagram Reel consumption redirected toward syllabus content in the same format.

📱

Android Application

CBSE Class 10 Science delivered as vertical reels — same swipe, same loop, same voluntary consumption as social media.

⚙️

Admin Console

Browser-based content production with automated subtitle generation via speech-to-text alignment. Built in a single evening.

🌐

MiVA Web (SCERT)

Free, static bilingual (English/Malayalam) interactive learning platform for Kerala SCERT ICT Standard X students.

CODEBASE

~80+ Kotlin files

DATABASE

40+ tables · 7 schemas

SECURITY

85+ RLS policies

SEALED INTERFACES

~20 frozen contracts

AI MODEL FAMILIES

Claude · ChatGPT · Gemini

BUILD TIMELINE

The messy reality. Not the *clean SDLC.*

The build didn't follow the clean SDLC the framework now prescribes — that mess is precisely what the SDLC is designed to prevent next time.

OCT – NOV 2025

Early video content prompts, experiment visuals, initial content transformation pipeline established.

→ FRAME phase – "video IS the product" as the core bet

DECEMBER 2025

Architecture foundations — sealed module concept, device management, B2C pricing, anti-drift-by-design philosophy formalised.

→ SEAL phase – the first sealed contracts declared

DEC 2025 – JAN 2026

Android build grows. Database migration. RLS across 40+ tables with 85+ policies deployed.

→ SPEC phase – SSOT born; extension pattern established

FEBRUARY 2026

Content production pipeline for Chapter 1. First NCERT accuracy audit reveals hallucination in chemistry equations.

→ VERIFY – scientific accuracy requires highest-reasoning model

FEB – MAR 2026

Ship Battle Plan. Feature freeze. Admin console built in a single evening using one-chat-per-unit pattern.

→ One-chat-per-unit rule discovered organically

MARCH 2026

v40 source validation — first major drift discovered. CD1 formally logged. Authority Chain introduced.

→ RECONCILE phase – Drift Incident process formalised

MARCH 2026

v41/beta6 — vulnerability fix confirmation drift (CD4). CancellationException guards. Sealed-baseline hash regenerated.

→ Reconciliation Sweep as first-class process

APRIL 2026

MiVA Web pivot — SCERT ICT Standard X. One-chat-per-chapter with Sonnet Extended. SDLC v1.0 written.

→ The framework closes on itself

THE EVIDENCE BASE

Four incidents. Every rule *traces to one.*

Click each incident to see what happened, the root cause, and what rule it produced in the SDLC.

HIGH

→ RECONCILE phase · Pair A sweep · weekly cadence

CD2

CRITICAL

→ Artefact Update Policy · PL-03 Surgical Edit prompt

CD3

CRITICAL

→ Source-First Rule · PL-05 Source-First Enforcement

HIGH

→ Authority Chain · Decision Log ↔ SSOT sweep

MODEL TRANSITIONS

Five transitions. Artefacts *unchanged.*

TRANSITION	WHAT BROKE	WHAT HELD
Opus 4.5 → 4.6	Prompts tuned for 4.5 were slightly under-specified; more concise outputs	All artefacts — SSOT, Decision Log, sealed contracts — unchanged
Opus 4.6 → 4.7	Context packaging needed tightening; 4.7 caught more subtle drift	Phase structure; prompts needed minor tuning only
Sonnet 4.5 → 4.6 Extended	Attachment budgets and output window changed; Model Selection Matrix updated	The one-chat-per-unit pattern scaled unchanged
Claude ↔ ChatGPT	More explicit role instructions needed; formatting conventions differ	Review-prompt structure (PL-07) generalised with minor adaptation
Claude ↔ Gemini	Tool use patterns differ substantially	The principle of cross-family review held

CLOSING OBSERVATIONS

Three things the build kept *proving*.

9.1

Verification Is Non-Optional

Every drift incident — CD1 through CD4 — traces to a missing verification step. Not weak; missing. The SDLC makes verification non-optional at phase gates, with "Ready for external review" as the load-bearing trigger.

9.2

Plausibility Is Dangerous

AI produces plausible work at a pace humans can't match. "It looks right" is not evidence of correctness. Viewing the actual source is. Every prompt in the MBG is designed to force this distinction.

9.3

Durable Artefacts Outlast Models

When Opus 4.5 became 4.7, the SSOT did not change. When Claude was replaced by ChatGPT for a review, the sealed contracts did not change. This is the result of deliberately designing artefacts to be model-agnostic.

"The SDLC framework is not the MiVA build. It is what the MiVA build taught us the process should have been from day one. The next project gets to start there."

· MiVA Education