

Orchestra AI

A · I

A I · D R I V E S · H U M A N S · N A V I G A T E

Navigator *Lite*.

The Navigator Framework for AI Native Development — for teams of 2-15.

Navigator Lite keeps the three things that matter most: named human accountability at every gate, a defined HITL moment, and evidence you can defend. What it trades away is redundancy — not rigour. Same five gates. Minimum viable governance surface.

RISK CLASSES

Standard & Elevated only

MINIMUM ROLES

3 distinct humans

CONTROLS

15 critical (vs 25 full)

DISTRIBUTION

Open · For Adoption

I.

Why a *lite framework*?

The full Navigator Framework targets enterprise teams with seven named roles, five risk classes, and 25 quantitative controls per project. Startups cannot staff that overhead — but they face the same core risk: AI generates work faster than humans can verify it, and code ships that nobody has actually read.

Navigator Lite applies the same operating principle with the minimum viable governance surface. It is appropriate for Standard and Elevated risk classes. If your project is High or Critical — regulated financial systems, safety-of-life, or irreversible AI decisions — use the full framework.

The same five gates. The same three-beat cycle. The same non-negotiable: AI cannot close a gate.

FULL FRAMEWORK VS NAVIGATOR LITE

FULL NAVIGATOR FRAMEWORK	NAVIGATOR LITE
7 named roles	3 minimum roles
4 risk classes (Standard → Critical)	2 classes (Standard, Elevated)
5 controls per gate · 25 total	3 critical controls per gate · 15 total
Full attestation templates	Single-paragraph sign-off
3 approvers for risk classification	Tech Lead + Founder sufficient
Quarterly calibration protocol	Quarterly 5-question self-assessment

II.

Three humans, *not three titles*.

In a startup, one person often wears many hats. Navigator Lite requires three distinct humans — not three distinct job titles. The constraint is independence, not headcount.

Tech Lead

Owens engineering decisions. Closes Build and Design gates. Can be CTO, senior engineer, or founding engineer.

Product / Founder

Owens requirements and acceptance criteria. Closes Define and Verify gates. Can be CEO, CPO, or product lead.

Reviewer

A named person who did not author the AI work under review. Can be another engineer, a trusted advisor, or a fractional principal. Must rotate — cannot review their own output.

Two-person teams: You need at least one external reviewer — an advisor, fractional CTO, or trusted peer — to close gates. The same person who authored the spec or code cannot be the Reviewer. Self-review is not independence.

III.

Two classes. *Set once per project*.

Class is set at project charter time by Tech Lead + Founder, documented, and revisited only if scope changes materially. If your work hits High or Critical criteria, Navigator Lite does not apply.

STANDARD

Internal / reversible

No customer PII. No revenue path. Fully reversible. Low blast radius.

Admin dashboards · internal scripts · dev tooling · prototypes

ELEVATED

Customer-facing / PII

Revenue-affecting, customer-facing, or handles GDPR-grade data.

Public APIs · payment flows · user auth · customer features

High / Critical: Regulated workloads, AI in the decision loop, safety-of-life, or irreversible decisions require the full Navigator Framework.

IV.

The three-beat *cycle*.

Every gate runs the same three beats. This is the non-negotiable core — unchanged from the full version. AI can run individual checks; it cannot decide which checks apply, and it cannot close a gate.

01 · SPEC

Define what to build
Before AI starts. Functionality, acceptance criteria, what is explicitly out of scope. A named human authors and owns each item.

02 · GUARDRAILS

Set up the checks
Run tests, scans, and sampling probes that prove the work is correct. Humans sample AI output and write down what they find.


03 · ATTEST

Approve and sign
A named human signs a dated statement: spec was followed, guardrails applied, work passed. AI cannot attest. An approval you cannot defend was not honestly given.

V.

The five phase *gates*.

Each gate is a mandatory checkpoint. A named human reviews artefacts and provides recorded approval. A gate either closes or it does not — there is no informal "good enough" path.



GATE 1

Define

ACCOUNTABILITY · COMPLIANCE · QUALITY DRIFT

Certifies what to build is clear, human-authored, and traceable before AI starts.

CONTROL	THRESHOLD	WHAT IT PREVENTS
D1 · AI-story adjudication	100%	Hallucinated requirements shipping as features
D2 · Reverse traceability	100%	Stories with no business justification
D3 · Open ambiguities	0 High/Critical open	Building the wrong thing

DEFINE ATTESTATION

"I have read and ratified every AI-generated story. Each traces to a business outcome I authored. No High or Critical ambiguity remains open. I can defend what we are building without consulting the AI agent."

Name · Date · Role

GATE 2

Design

ACCOUNTABILITY · COMPETENCY · QUALITY DRIFT · COMPLIANCE

Certifies how to build it is decided, alternatives considered, and design traces to the locked spec. AI may not make architecture decisions by default.

CONTROL	THRESHOLD	WHAT IT PREVENTS
DS1 · Architecture Risk Index	≤70 Std / ≤60 Elev	Runaway architectural risk undetected until production
DS2 · AI-design ratification	100%	AI deciding architecture by default
DS3 · ADR coverage	100%	Untraceable design choices

DESIGN ATTESTATION

"Every significant design decision has a human-authored ADR with rejection reasons I wrote. AI-generated threats have been assessed by me. The design traces to the locked spec. I can defend this architecture without consulting the AI agent."

Name · Date · Role

Who closes it: Tech Lead signs both · Reviewer (not the author) reviews ADRs

GATE 3

Build

ALL FIVE FAILURE MODES

Certifies code builds, has been read by a qualified human, and AI-generated tests are not the sole validators of AI-generated code.

CONTROL	THRESHOLD	WHAT IT PREVENTS
B1 · Reviewed-and-probed rate	100% attested	Code no human has read reaching production
B2 · Test Independence Ratio	≥30% Std / ≥40% Elev	AI grading its own work
B3 · Mutation kill rate	≥50% Std / ≥65% Elev	Tests that catch no real defects

Model independence: The AI reviewing code must be a different model from the one that wrote it. Same-model self-review is not independence.

BUILD ATTESTATION

"Every change has been read by a named reviewer who can explain what it does, why this approach, and what would fail first. AI-generated tests have not been the sole validators of AI-generated code."

Name · Date · Role

GATE 4

Verify

ACCOUNTABILITY · COMPLIANCE · QUALITY DRIFT · COMPETENCY

Certifies the system passes the acceptance criteria humans declared in Define, independently of how the code was authored.

CONTROL	THRESHOLD	WHAT IT PREVENTS
V1 · Human Verification Depth	100%; ≥10% spot-checked	Trusting AI summaries without reading evidence
V2 · AC execution	100%; 0 unresolved failures	Shipping without meeting declared criteria
V3 · Verification independence	Tests not from code's agent; different model	Self-confirming AI verification loops

VERIFY ATTESTATION

"Every acceptance criterion we declared has been executed. Every failure resolved or formally accepted. I have audited a sample of AI-summarised results against raw output. We can defend this verification record."

Name · Date · Role

Who closes it: Product/Founder + Tech Lead co-sign

GATE 5

Run

OPERATIONAL RESILIENCE · QUALITY DRIFT · ACCOUNTABILITY

Certifies the system is operable by humans — including when AI tools are absent, slow, wrong, or unavailable.

CONTROL	THRESHOLD	WHAT IT PREVENTS
R1 · AI-Absent operability	Drill passed; ≥70% Std / ≥80% Elev	On-call collapse when AI tools go down
R2 · Rollback readiness	Drill within N days; time ≤ RT0	Inability to recover from a bad deploy
R5 · Accountability transfer	Named on-call + Founder sign-off	Nobody owning production

RUN ATTESTATION

"Named on-call can operate, diagnose, and recover this system without AI assistance for the declared window. We have rehearsed a real rollback. We retain accountability for production operation."

Name · Date · Role

Who closes it: Tech Lead + Founder co-sign

VI.

Gate failure & *exceptions*.

A gate either closes or it does not. When a Completeness or Sign-off artefact is incomplete, the work is rejected to the prior phase — not waved through. There is no informal exception path.

- 1 **Named approver** — Tech Lead or Founder, not the person requesting the exception.
- 2 **Hard expiry date** — written down before the exception is granted.
- 3 **Compensating control** — a documented mitigation that stands in for the missed artefact.
- 4 **Calendar entry for retrospective review** — the exception is re-examined when it lapses. It does not silently become permanent.

Why this matters: Governance frameworks die through accumulated informal waivers. The exception process is where discipline is most at risk under delivery pressure. Making every exception named, bounded, and reviewed keeps the gate meaningful.

VII.

The five-question *self-assessment*.

Every 90 days, Tech Lead and Founder answer these questions together. If the answer to questions 2, 3, or 5 is yes — stop and remediate before the next increment.

- ☐ Can our engineers explain what the AI built, without re-running it?
- ☐ **Have we shipped anything no human has actually read?** → stop if yes
- ☐ **Did any gate close on "trust" rather than evidence?** → stop if yes
- ☐ Can our on-call diagnose and recover without AI tools for the declared window?
- ☐ **Have any exceptions become silently permanent?** → stop if yes

VIII.

When to upgrade to the *full framework*.

Navigator Lite is appropriate for Standard and Elevated risk classes. Move to the full framework when any of the following apply:

- You handle regulated data — financial services, legal, or HIPAA/DPDP-grade PII

- AI is in a decision loop that affects individuals (lending, hiring, clinical triage)
- Your blast radius becomes irreversible — safety, compliance exposure, critical infrastructure
- Your team grows past ~15 and role collapse is no longer reasonable
- A regulator asks "show me human review" and Lite evidence doesn't satisfy them

IX.

One page. *Five gates.*

The complete Navigator Lite governance model — who closes each gate, and the single failure it guards against.

GATE	HUMAN WHO CLOSES IT	THE ONE THING IT PREVENTS
1 · Define	Product/Founder + Tech Lead	Building what the AI hallucinated, not what the business needs
2 · Design	Tech Lead + Reviewer	AI making architecture decisions by default
3 · Build	Tech Lead + Reviewer	Code no human has read reaching production
4 · Verify	Product/Founder + Tech Lead	Shipping without meeting declared acceptance criteria
5 · Run	Tech Lead + Founder	On-call collapse when AI tools are unavailable

NAVIGATOR FRAMEWORK FAMILY

- Sub-framework 01 · Identify** — Navigator Evaluator for Technical Hiring. A live 60–90 min evaluation identifying engineers who direct AI agents vs. those directed by them.
- Sub-framework 02 · Structure** — Navigator Operator for AI Native Development. Named human approval at five SDLC phase gates — Define, Design, Build, Verify, Run — scaled by EU AI Act-aligned risk class. Navigator Lite is the startup variant of this sub-framework.
- Sub-framework 03 · Sustain** — Sustaining Human Oversight. Three operational signals (CSI, HOI, ATDI) that make oversight sustainability measurable as AI throughput scales.

AI does the work. *A named human owns it.*

Navigator Lite provides assurance, not proof of correctness. The question at every gate: who can defend the decision without re-running the agent. For the full framework, operational detail, and adoption support — reach out.

Navigator Lite · Sub-Framework 02 Lite · v1.0
rajesh@orquestraai.io
orquestraai.io