

The Three Manifestos

Orienting, Internal, and Public—Mapped to Controls, Release Gates, and Governance

Manifestos are not one thing. They operate at different moments, for different audiences, and under different risk conditions. Confusing them—or treating them all as timeless truth—creates drift, capture, or theater.

This companion piece gives a **clean operational map**: when to use each manifesto type, what it's for, how it fails, and how to bind it to **controls** so it survives incentives.

1) Orienting Manifesto

When: Before a project “begins” (the preparation phase)

Audience: The team doing the work

Function: Direction-setting, not decision-making

What it is

An Orienting Manifesto sets **headings**, not outcomes. It aligns attention and language before concrete artifacts exist. It answers: *What kind of work are we trying to do together?*

What it's good at

- Creating shared vocabulary
- Signaling values and posture
- Encouraging exploration without premature lock-in

How it fails

- Becomes mistaken for a plan
- Hardens into ideology
- Gets cited after-the-fact as justification

Bind it to controls

Orienting manifestos should **never ship products**. They should ship **questions** and **guardrails**.

Control binding (lightweight):

- **Decision log requirement:** any major choice must reference which orienting principle it tested
- **Exploration budget:** time-boxed experiments tied to principles

- **Review cadence:** scheduled check-in to retire or refine principles

Release gate: *None.*

Orienting manifestos do not authorize scale.

2) Internal Manifesto

When: Novel differences begin to emerge

Audience: The builders/operators closest to the system

Function: Protection and refusal (“No”)

What it is

An Internal Manifesto exists to **keep difference alive**—to prevent default incentives from reverting the work to familiar, harmful, or easier patterns.

It is the most powerful—and most dangerous—manifesto type.

What it's good at

- Blocking known failure modes
- Creating productive constraints
- Establishing non-negotiables during scale pressure

How it fails

- Turns into dogma
- Becomes unrevisable
- Loses enforcement and becomes symbolic

Bind it to controls

Every Internal Manifesto line must be **executable**.

Required bindings (non-negotiable):

- **Control:** a concrete constraint that changes behavior
- **Verification:** how you prove the control works
- **Owner:** named accountability
- **Stop condition:** what forces a pause or rollback

Release gate: *Fail-closed.*

If a manifesto “No” is violated without mitigation, the system does not scale.

3) Public Manifesto

When: The system affects people beyond the builders

Audience: Users, regulators, partners, the public

Function: Legitimacy and accountability

What it is

A Public Manifesto communicates commitments outward. It invites scrutiny. It creates expectations others can hold you to.

What it's good at

- Establishing trust through clarity
- Making boundaries explicit
- Aligning external stakeholders

How it fails

- Becomes marketing
- Incentivizes reputation over truth
- Creates liability without capability

Bind it to controls

A Public Manifesto must be a **projection** of existing controls—not a promise of future virtue.

Required bindings:

- **Evidence links:** audits, reports, metrics, or summaries
- **Enforcement clarity:** who investigates and remedies breaches
- **Revision protocol:** how commitments change with context

Release gate: *Disclosure-ready*.

If you can't evidence it, don't publish it.

One Page: Manifesto-to-Control Map (Use This)

Manifesto Type	Primary Risk	Control Weight	Release Authority
Orienting	Drift, vagueness	Light	None
Internal	Incentive collapse	Heavy (fail-closed)	Required
Public	Performative ethics	Evidence-backed	Conditional

The Governance Spine (What Keeps This Sane)

To prevent capture or decay, all three manifesto types share a governance backbone:

1. **Ownership:** named role + escalation path
2. **Revision cadence:** scheduled, evidence-driven updates
3. **Evidence standard:** what counts as proof
4. **Separation of layers:** orienting ≠ internal ≠ public
5. **Sunset clauses:** retire what no longer serves

This is how manifestos stay **alive** without becoming **authoritarian**.

Do This Next (30-minute exercise)

1. Write **three short manifestos** for your current initiative:

- Orienting (5 lines)
- Internal (5 “We will not...” lines)
- Public (5 commitments)

2. For the **Internal** lines only, attach:

- one control,
- one verification step,
- one owner,
- one stop condition.

3. Label anything without controls as **aspirational**—not enforceable.

That labeling step is the difference between honesty and theater.

Bottom Line

Manifestos are not relics. They are **tools**.

But tools only work when paired with mechanisms that hold under stress. In modern systems, the order matters:

Orientation → Protection → Accountability

Poetry can inspire.

Controls decide outcomes.

If you want your values to survive scale, you must design for the moment when words are no longer enough.