



Alignment Before Autonomy

How legal, compliance,
security, product, and engineering
align around AI governance

**GOVERNANCE FAILS WHEN THE ORGANIZATION IS NOT ALIGNED ON
THE DETAILS.**

Audience: executive sponsors, product leaders, platform/security
teams, legal/compliance stakeholders.

LSAS Stack by Inference Stack LLC
Enterprise AI Execution Authority™

ADDRESS
Inference Stack LLC
5365 W. 73rd Pl.
Arvada, CO 80003

CONTACT
+1 720.334.3838
www.lsas-stack.com
www.lsas-stack.com/contact

Alignment Before Autonomy

Alignment is not consensus theater.

Organizations often say they want responsible AI, but the details remain unresolved: which workflows are high-stakes, what counts as acceptable evidence, who owns exception policy, who signs off on release, and what telemetry must be reviewed.

Autonomy scales only after those details are aligned across functions.

Why alignment is now a governance requirement

28 %

of AI-using organizations say their CEO oversees AI governance.

McKinsey, 2025

<25 %

of companies have board-approved, structured AI policies, according to McKinsey citing NACD research.

McKinsey, 2025

53 %

of corporate legal departments report having a formalized technology roadmap.

WEF, 2025

Governance drift rarely begins with bad intent. It usually begins with ambiguity. Product teams optimize for velocity, legal teams optimize for defensibility, security teams optimize for blast radius, and business leaders optimize for outcomes. If those groups are not aligned on the operating details, the enterprise gets one of two bad outcomes: either AI is slowed by endless handoffs, or it ships with unclear authority and inconsistent controls.

The first order of business is to align on what posture the business should take with AI.

McKinsey's 2025 board-governance guidance puts posture first for good reason. Posture determines what the organization is actually trying to optimize for: aggressive scale, bounded experimentation, highly controlled deployment, or something in between. Without that alignment, every control conversation becomes a local negotiation.

Alignment Before Autonomy

The seven alignment decisions

Decision area	What must be agreed
AI posture	What level of aggressiveness, caution, and business ambition the enterprise is willing to take.
Boundary map	Which apps, agents, APIs, systems of record, and data flows are in scope first.
Authority model	What AI may say, do, and rely on for each critical workflow.
Evidence model	What counts as sufficient proof for release, recommendation, or action.
Escalation model	When humans must review, who can approve, and what timeline/attestation applies.
Telemetry model	Which stakeholders review which signals, at what cadence, and with what thresholds.
Change control	How policy packs, validators, integrations, and exceptions are updated and approved.

Cross-functional role design

Illustrative RACI-style split

Function	Primary concern	Owns / decides	Reviews / advises
Executive Sponsor	Risk appetite and business value	Posture, priority workflows, funding, escalation expectations	Major incidents, expansion decisions
Legal / compliance	Obligations, defensibility, evidence standards	Release rules, exception policy, approved evidence requirements	Policy-pack changes, customer/board materials
Security / privacy	Data exposure, blast radius, operational risk	Boundary controls, logging posture, secrets handling, residency constraints	Incident review, control drift
Product	Workflow design and user impact	Where controls appear in product behavior and how governed outcomes surface to users	Success metrics, remediation UX

Alignment Before Autonomy

Engineering / platform	Runtime implementation and reliability	Gateway placement, integrations, telemetry plumbing, deployment topology	Policy execution mechanics
Domain SMEs	Professional acceptability	Clinical, legal, or financial correctness expectations	Escalation criteria and evidence sufficiency

A practical operating cadence

- Monthly AI posture and control council: executive sponsor, product, security, legal/compliance, platform lead.
- Release-gate review for material workflow changes: new tool use, new evidence source, new outbound boundary, or new automation path.
- Incident and exception review: recurring inspection of escalations, overrides, blocked flows, and drift signals.
- Quarterly expansion decision: determine whether the current boundary is stable enough to add the next workflow or environment.

What alignment looks like in practice with LSAS

LSAS already frames implementation around boundaries, control objectives, policy packs, workflow configurations, telemetry, and executive checkpoints. That is useful because it gives the organization concrete artifacts to align around instead of abstract principles. A boundary map, an evidence rule, a policy tier, and an escalation flow are much easier to govern than the broad directive to “use AI responsibly.”

Misaligned organization

- Each team defines risk differently
- Product ships drafts that legal later rejects
- Security finds boundary issues late in the cycle
- Exceptions are made informally and lost in chat threads

Aligned organization

- Critical workflows have named control owners
- Evidence thresholds are agreed before launch
- Runtime telemetry becomes a shared review surface
- Exceptions follow an explicit authority and audit path

A 90-day alignment agenda

Function	Primary concern	Owens / decides
Days 0-15	Lock scope and posture	In-scope boundary map, executive owners, control objective register.
Days 16-45	Implement and observe	Configured policies, telemetry setup, baseline-versus-governed comparisons.
Days 46-90	Operationalize and harden	Exception model, audit review flow, expansion recommendation.

Alignment Before Autonomy

Bottom line

Governance is not just an issue of controls. It is an issue of organizational coherence. If the enterprise wants AI systems that operate under deliberate authority, the organization must align on posture, evidence, escalation, telemetry, and change. Only then does autonomy become something the institution owns rather than something the model improvises.

Selected references

1. **McKinsey – The AI reckoning: How boards can evolve (2025).** Board posture, ownership, and governance policy guidance for AI oversight. <https://www.mckinsey.com/capabilities/mckinsey-technology/our-insights/the-ai-reckoning-how-boards-can-evolve>
2. **McKinsey – The State of AI: How organizations are rewiring to capture value (2025).** Covers AI adoption rates, workflow redesign, and governance oversight patterns. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-how-organizations-are-rewiring-to-capture-value>
3. **LSAS Stack – Pricing / implementation approach.** Implementation phases, inputs, first-90-day cadence, executive telemetry, and delivery outputs. <https://www.lsas-stack.com/pricing>
4. **LSAS Stack – Compliance posture.** Controls mapping language, deployment/residency posture, data minimization, auditability, and operational hooks. <https://www.lsas-stack.com/compliance-posture>
5. **FTI Consulting / Relativity – The General Counsel Report, AI adoption in legal departments (2026).** Corporate legal department AI adoption, roadmap maturity, and common task areas. <https://www.fticonsulting.com/about/newsroom/press-releases/ai-adoption-in-corporate-legal-departments-doubles-according-to-the-general-counsel-report>
6. **Inference Stack – Contact.** Inference Stack contact page; based in Denver, Colorado; serving clients worldwide. <https://www.inference-stack.com/contact>

