Executive Toolkit: Runtime AI Governance Readiness & Alignment

Every organization adopting AI faces the same fundamental question:

Can we demonstrate that our AI systems are governed; not merely documented?

Most organizations have begun developing AI policies, inventories, and governance processes. Far fewer have evaluated whether those controls operate consistently in practice, produce defensible evidence, or provide meaningful executive oversight as AI systems become increasingly autonomous.

This Executive Toolkit provides a structured assessment of AI governance readiness using widely recognized governance practices while introducing additional executive considerations for organizations preparing for higher levels of AI autonomy.

Who This Is For

Ideal for CTOs, CIOs, CISOs, CROs, compliance leaders, and founders scaling AI initiatives.

This toolkit is designed as a practical governance readiness diagnostic that helps organizations assess foundational, operational, and executive runtime governance maturity. It supports executive discussions, governance planning, and audit preparedness by identifying strengths, gaps, and areas requiring additional governance capability.

It is intended as a governance readiness assessment and planning resource and does not constitute legal, regulatory, or compliance advice.

Why Runtime Governance Matters Now

- Rising autonomy in AI agents and workflows
- Increasing regulatory pressure (EU AI Act, etc.)
- Growing risk of “silent failures” where systems act without current legitimacy
- The gap between documented policies and actual runtime behavior

---

This toolkit helps executives evaluate whether governance controls operate at runtime and provides a structured path from discovery to validation.

1. Executive AI Governance Readiness Checklist

Foundational Readiness

- Comprehensive inventory of AI systems and agentic workflows (with named business + technical owners)
- Consequential or high-risk use cases identified and classified
- Clear accountability assigned for each system
- AI-specific incident response plan documented and tested

Operational Readiness

- Runtime activity logging and decision trails in place
- Human oversight and escalation paths defined for high-consequence decisions
- Third-party AI vendors assessed for risk and contractual controls

Runtime Governance Readiness

- Execution boundaries identified and documented
- Refusal / pause / escalation mechanisms enforced when legitimacy fails
- Authority and admissibility verified at bind time
- Governance decisions replay-verifiable with custody proof

---

2. Executive Runtime Governance Implementation Framework

Phase 1 — Discovery (Weeks 1–2)

- Map all consequential AI use cases
- Identify execution boundaries
- Document current authority, oversight, and evidence mechanisms

Phase 2 — Gap Assessment (Weeks 3–4)

- Compare static controls vs. runtime requirements
- Evaluate refusal rail effectiveness and continuity risks
- Assess replay and reconstruction capabilities

Phase 3 — Governance Design (Weeks 5–6)

- Define runtime authorization and custody proof requirements
- Design refusal, escalation, and intervention pathways
- Develop governance artifact templates and board reporting

Phase 4 — Validation

- Conduct shadow-mode testing
- Verify refusal / escalation mechanisms
- Confirm regulatory alignment and replay capability

---

3. How to Select a Runtime Governance Partner

Must-Have Capabilities

- Deep expertise in runtime enforcement (not just policy)
- Ability to demonstrate custody proof and refusal rails in practice
- Structured methodology for execution-time governance
- Clear, fixed-scope diagnostic offering

Strong Positive Signals

- Uses terms like *execution boundary*, *admissibility*, and *consequence formation*
- Emphasizes evidence at bind time rather than documentation alone
- Provides operational examples, not just frameworks

Red Flags

- Heavy focus only on policies, ethics, or high-level frameworks
- Vague deliverables and timelines
- No clear plan for runtime controls or verifiable artifacts

Test Questions for Vendors

- *"Can you provide custody proof for a decision made six months ago?"*
- *"How does your system block execution when governance conditions are not met?"*

---

4. EU AI Act Readiness Check

For organizations operating within or serving the European Union.

Governance & Accountability

☐ High-risk AI systems have been identified and documented.

☐ Roles and responsibilities for EU AI Act compliance have been assigned.

☐ AI systems are classified according to applicable EU AI Act risk categories.

Risk & Control Requirements

☐ A documented risk management process exists for high-risk AI systems.

☐ Human oversight requirements have been defined and operationalized.

☐ Accuracy, robustness, and cybersecurity requirements have been assessed.

☐ Data governance and data quality controls are documented.

Transparency & Evidence

☐ Technical documentation is maintained and available for review.

☐ Logging mechanisms support traceability and post-decision reconstruction.

☐ Governance decisions can be independently reconstructed if challenged.

☐ Significant decisions and interventions are supported by preserved evidence.

Runtime Governance Considerations

☐ Consequential actions are evaluated at execution time, not solely at design time.

☐ Authority and oversight conditions can be revalidated before execution.

☐ Refusal, escalation, and intervention mechanisms exist when governance conditions fail.

☐ Runtime evidence supports demonstration of ongoing compliance and accountability.

Advisory: The EU AI Act establishes obligations for certain AI systems throughout their lifecycle. Organizations should assess whether their AI systems fall within prohibited, high-risk, transparency, or general-purpose AI categories and obtain appropriate legal and regulatory advice where necessary.

EU AI Act Question	Executive Consideration
Have high-risk systems been identified?	☐ Yes ☐ No
Is human oversight operationalized?	☐ Yes ☐ No
Can decisions be reconstructed?	☐ Yes ☐ No
Can consequential actions be refused or paused?	☐ Yes ☐ No
Can governance decisions be demonstrated to regulators?	☐ Yes ☐ No

---

5. NIST AI RMF Executive Alignment Check

NIST AI RMF provides process maturity; runtime governance extends it to execution legitimacy

GOVERN

☐ Organizational accountability for AI risk has been formally assigned.

☐ AI governance policies, standards, and procedures have been documented.

☐ AI risk management roles and decision rights are clearly defined.

☐ AI risks are regularly reviewed by executive leadership.

Executive Question:

Who owns AI risk, and how is accountability demonstrated?

MAP

☐ AI use cases, stakeholders, and affected populations have been identified.

☐ Consequential and high-risk AI systems have been classified.

☐ Intended use, limitations, and potential impacts have been documented.

☐ Cross-system dependencies and third-party relationships are understood.

Executive Question:

Do we understand where AI creates consequence and who may be affected?

MEASURE

☐ AI systems are monitored for performance, reliability, and drift.

☐ Risks, incidents, and control effectiveness are regularly assessed.

☐ Testing includes failure scenarios and degraded operating conditions.

☐ Governance evidence can be independently reconstructed and reviewed.

Executive Question:

Can we demonstrate that controls are operating effectively?

MANAGE

☐ Risk responses are documented and periodically reviewed.

☐ Escalation, intervention, and incident response mechanisms exist.

☐ Human oversight responsibilities are operationalized.

☐ Governance mechanisms can pause, refuse, or retire consequential actions when conditions require.

Executive Question:

Can the organization intervene effectively before unacceptable consequences occur?

Runtime Governance Extension (Codex Sovereign)

While NIST AI RMF establishes foundational governance expectations, increasingly autonomous systems may require additional runtime capabilities:

☐ Authority can be revalidated at execution time.

☐ Consequential actions are evaluated for admissibility before execution.

☐ Refusal and escalation mechanisms operate under degraded conditions.

☐ Governance decisions produce replay-verifiable evidence.

☐ Legitimacy can be demonstrated when consequence forms.

Executive Summary

NIST AI RMF Function	**Status**
GOVERN	☐ Green ☐ Yellow ☐ Red
MAP	☐ Green ☐ Yellow ☐ Red
MEASURE	☐ Green ☐ Yellow ☐ Red
MANAGE	☐ Green ☐ Yellow ☐ Red

5. Executive Summary Scorecard

Use this scorecard to self-assess governance maturity across domains:

Domain	Guidance for Rating	Status
Foundational Readiness	**Green:** Core governance foundations are established and documented. **Yellow:** Partial coverage exists, but important gaps remain. **Red:** Significant foundational governance gaps exist.	☐ Green ☐ Yellow ☐ Red
Operational Readiness	**Green:** Runtime logging, oversight, and escalation mechanisms operate consistently. **Yellow:** Controls exist but are inconsistently applied. **Red:** Key operational controls are absent or largely informal.	☐ Green ☐ Yellow ☐ Red
Runtime Governance Readiness	**Green:** Refusal rails, custody proof, and replay capabilities are operational and enforced. **Yellow:** Runtime controls exist but are only partially implemented. **Red:** Runtime governance mechanisms are absent or purely manual.	☐ Green ☐ Yellow ☐ Red
Overall Governance Maturity	**Green:** Governance capabilities are mature and integrated across domains. **Yellow:** Governance capabilities are uneven or inconsistently implemented. **Red:** Governance maturity is weak, fragmented, or largely undocumented.	☐ Green ☐ Yellow ☐ Red

Recommended Next Steps

If your assessment identified gaps:

Foundational Governance

- Build an AI inventory
- Establish governance ownership
- Develop AI policies
- Classify AI use cases

Operational Governance

- Define oversight processes
- Establish monitoring
- Improve audit evidence
- Validate vendor governance

Runtime Governance

- Identify consequence-bearing AI
- Evaluate execution boundaries
- Assess intervention mechanisms
- Evaluate runtime enforcement capabilities

Runtime governance transforms AI oversight from policy compliance into operational legitimacy - ensuring every consequential action remains governed before execution.

Continue Building Your AI Governance Program

This Executive Toolkit establishes a governance baseline. As organizations mature, governance expands beyond policies and documentation toward operational implementation and executive runtime governance.

Upcoming Resources

- Executive AI Inventory Toolkit
- AI Governance Policy Library
- AI Discovery & Interview Scripts

- Executive Governance Templates
- Executive Runtime Governance Readiness Assessment
- Executive Advisory Services

Whether you're establishing your first AI governance program or preparing for increasingly autonomous AI systems, these resources are designed to help organizations move from governance planning to governance in practice.

Governance is not demonstrated by the policies an organization writes. It is demonstrated by the authority it can exercise before consequential AI actions occur.

Learn more at:
CodexSovereign.org