

SAI RAI Solutions

# SERVE™ Framework

*A Practical Guide to Responsible AI Governance*

Scan · Evaluate · Remediate · Verify · Embed

GROUNDED IN

**NIST AI RMF · ISO 42001**  
**EU AI Act · OECD Principles**

LICENSED UNDER

**Creative Commons BY 4.0**  
Free to use with attribution

PUBLISHED BY

**Vijay Reddy**  
sairaisolutions.com · April 2026**SECTION 01**

## Why Responsible AI Governance Fails

Most organizations deploy AI faster than they govern it. The result is a growing gap between what AI systems do and what executives, auditors, and regulators need to see. Boards face audit exposure. Legal teams face regulatory uncertainty. Program managers face pressure to deliver — without clear guidance on what responsible AI compliance actually requires of them.

Three patterns drive this failure:

- Governance lives in policy, not delivery. RAI principles are written at the executive level but never translated into program-level action. Teams don't know which of their AI use cases are in scope, at risk, or already non-compliant.
- Standards are intimidating. NIST AI RMF, ISO 42001, and the EU AI Act are rigorous — and essential. But without a bridge framework, practitioners struggle to connect those standards to their day-to-day work.
- Compliance is treated as a one-time event. Real governance is continuous. It requires monitoring, adjustment, and embedding into delivery cadences — not just a checkbox before launch.

### The governance gap in numbers

- 74% of organizations have deployed AI but lack a formal governance program (Gartner, 2025)
- EU AI Act fines reach up to €35M or 7% of global revenue for high-risk AI violations
- ISO 42001 is now the baseline expectation for enterprise AI management systems globally
- NIST AI RMF has been adopted by the US federal government and most regulated industries

## SECTION 02

# The SERVE™ Framework

SERVE™ is a five-stage responsible AI governance methodology designed for organizations deploying AI in real delivery environments. Unlike frameworks built for policy teams, SERVE™ is built for practitioners — program managers, agile coaches, product owners, and technology leaders who need governance that works alongside delivery, not against it.

Each stage produces concrete deliverables: documented artifacts that satisfy auditors, inform boards, and give regulators what they need to see. SERVE™ maps directly to NIST AI RMF, ISO 42001, and the EU AI Act — so your governance work is automatically aligned to the standards that matter.

## STAGE 01

S

SCAN

## Know what AI you have — and what risk it carries

Before you can govern AI, you have to find it. Most organizations have more AI in production than their governance teams know about — including shadow AI, vendor-embedded AI, and AI features in existing tools. Scan builds a complete picture.

- Build a comprehensive AI system inventory — including vendor, shadow, and embedded AI
- Classify each system by risk tier (minimal, limited, high, unacceptable under EU AI Act)
- Identify use cases touching regulated activities, personal data, or consequential decisions
- Assign named owners to each system in the inventory

**NIST:** Map 1.1, Map 2.1    **ISO 42001:** Clause 6.1, A.6.1.1

## STAGE 02

E

EVALUATE

## Measure maturity honestly against the standards that matter

Knowing what AI you have is not enough. You need to know where your governance practices stand relative to the standards your auditors and regulators are checking against. Evaluate produces an honest, evidence-based maturity picture across six responsible AI pillars.

- Assess maturity across Fairness, Transparency, Accountability, Privacy, Safety, and Oversight
- Map current controls to NIST AI RMF Measure functions and ISO 42001 Annex A
- Identify the highest-priority gaps — ranked by risk tier and regulatory exposure
- Produce a maturity report suitable for board and audit committee presentation

**NIST:** Measure 2.1–2.11    **ISO 42001:** Clause 6.1.2, A.6.2

## STAGE 03

## R

REMEDiate

## Close the gaps that matter most — in priority order

A gap analysis without a remediation plan is just a list of problems. Remediate turns the Evaluate findings into a prioritized action plan, matched to your organization's risk appetite, timeline, and delivery capacity.

- Build a prioritized remediation roadmap — high-impact, low-effort items first
- Design and implement controls for your highest-risk AI systems
- Develop policies: data governance, model lifecycle, incident response, bias testing
- Integrate governance checkpoints into existing delivery ceremonies (PI Planning, Sprint Review)

**NIST:** Manage 1.1-4.2    **ISO 42001:** Clause 8.1, A.7

## STAGE 04

## V

VERIFy

## Build the evidence your auditors and board actually need

Governance without documentation is invisible. Verify assembles the audit trail — the policies, test results, review records, and control evidence that demonstrate your governance program is real, operating, and effective.

- Conduct internal audit against ISO 42001 control domains and NIST AI RMF
- Assemble evidence packages for each high-risk AI system
- Prepare board-ready reporting: governance dashboards, risk registers, incident summaries
- Conduct red-team and adversarial testing for generative and high-risk systems

**NIST:** Measure 2.5, 4.1    **ISO 42001:** Clause 9.1, 9.2, A.6.2.8

## STAGE 05

## E

EMBED

## Make governance sustainable — not a one-time project

The goal of SERVE™ is not a governance program that exists on paper. It is governance that is alive inside your organization — embedded in delivery, owned by practitioners, and continuously improving. Embed institutionalizes everything the prior stages built.

- Establish an AI Governance Office or Center of Excellence with clear RACI
- Embed RAI review checkpoints into PI Planning, Sprint Reviews, and release gates
- Launch practitioner training — so every team member understands their RAI role
- Build continuous monitoring: drift detection, performance metrics, harm signals

**NIST:** Govern 1.1-6.2    **ISO 42001:** Clause 10.1, A.9

## SECTION 03

# How SERVE™ Maps to Global Standards

SERVE™ was designed from the ground up to align with the four frameworks that regulators, auditors, and enterprise governance teams rely on. This means that work done inside SERVE™ automatically produces evidence and artifacts that satisfy external standards — without duplicate effort.

SERVE™ Stage	NIST AI RMF Functions	ISO 42001 Clauses	EU AI Act Articles	OECD Principles
<b>Scan</b>	Map 1.1, 2.1, 2.3	Cl. 6.1, A.6.1.1	Art. 10, Annex III	1.1 – Sustainable AI
<b>Evaluate</b>	Measure 2.1-2.11	Cl. 6.1.2, A.6.2	Art. 9, 10, 13	1.3 – Transparency
<b>Remediate</b>	Manage 1.1-4.2	Cl. 8.1, A.7	Art. 14, 15, 16	1.4 – Robustness
<b>Verify</b>	Measure 4.1, 2.5	Cl. 9.1, 9.2	Art. 43, 61, 72	1.5 – Accountability
<b>Embed</b>	Govern 1.1-6.2	Cl. 10.1, A.9	Art. 17, 26, 27	1.2 – Human oversight

## SECTION 04

# Getting Started with SERVE™

SERVE™ is designed to be adopted incrementally. You do not need to implement all five stages at once. Most organizations start with Scan and Evaluate — understanding what they have and where their gaps are — before moving into active remediation and governance office setup.

## Option A — Self-Directed Start

If your organization wants to begin with a self-assessment before engaging external support, start here:

- Take the free RAI Maturity Assessment at [sairaisolutions.com/rai-assessment](https://sairaisolutions.com/rai-assessment) — 10 questions, 5 minutes, instant report across 6 pillars
- Download the ISO 42001 Readiness Checklist (coming soon) - gap analysis against the key control domains
- Review the SERVE™ stage descriptions above and identify which stage your organization is currently operating at
- Use the free resources in our Learn Library — NIST AI RMF, OECD AI Principles, EU AI Act, UNESCO AI Ethics

## Option B — Guided Workshop

The fastest way to get your program managers and agile teams aligned on RAI governance obligations:

### "Responsible AI for Agile Teams" Workshop

- Half-day, virtual or on-site — designed for program managers, RTEs, and agile practitioners
- Maps your actual AI use cases to SERVE™ stages and identifies priority governance gaps
- Participants leave with a starter AI risk register and RAI compliance checklist
- SAI RAI Practitioner badge awarded upon completion
- Price: \$799/person · Group rates available · Contact us to schedule

## Option C — Consulting Engagement

For organizations ready to build or mature a formal AI governance program:

- AI Governance Jumpstart: 2-week sprint delivering a risk register, governance charter, and SERVE™ Stage 1 implementation. Starting at \$8,500.
- AI Risk Assessment: 4-week NIST AI RMF mapping across your AI portfolio. Starting at \$15,000.
- ISO 42001 Readiness Program: 8–12 week engagement producing audit-ready documentation. Starting at \$25,000.
- Fractional AI Governance Lead: embedded 2–3 days/week. Starting at \$15,000/month.

## SECTION 05

# Standing on the Shoulders of Giants

SERVE™ does not claim to have invented responsible AI governance. It is a synthesis framework — designed to make the work of remarkable institutions accessible to practitioners and organizations that might not otherwise reach it.

We are grateful to, and explicitly acknowledge, the following organizations whose work forms the foundation of SERVE™:

## Standards & Regulatory Bodies

- NIST — AI Risk Management Framework (AI RMF 1.0) · [airc.nist.gov](https://airc.nist.gov)
- ISO/IEC JTC 1/SC 42 — ISO/IEC 42001:2023 AI Management Systems Standard
- European Commission — EU Artificial Intelligence Act · [artificialintelligenceact.eu](https://artificialintelligenceact.eu)
- OECD — AI Principles, adopted by 46 countries · [oecd.ai](https://oecd.ai)
- UNESCO — Recommendation on the Ethics of AI · [unesco.org](https://unesco.org)

## Research & Practitioner Community

- Center for AI Safety: Dan Hendrycks · "Introduction to AI Safety, Ethics and Society" (free textbook, slides, videos) · [aisafetybook.com](https://aisafetybook.com) · <https://safe.ai/>
- RAI Institute: Independent responsible AI verification and frameworks · [responsible.ai](https://responsible.ai)
- Partnership on AI: Multi-stakeholder governance research · [partnershiponai.org](https://partnershiponai.org)
- All Tech Is Human: Responsible Technology Practitioner community. [alltechishuman.org](https://alltechishuman.org)
- AI Governance International — Independent advisory on multi-jurisdictional AI regulation. <https://www.aigovernanceinternational.org/>
- Montreal AI Ethics Institute — Research and practitioner newsletter · [montrealethics.ai](https://montrealethics.ai)
- AIGL Blog — AI Safety, Ethics, and Society resources. <https://www.aigl.blog/>

SERVE™ is licensed under Creative Commons BY 4.0. You are free to use, share, and adapt it — with attribution. The pip-installable Python library is open-sourced under Apache 2.0 via the SAI4RAI nonprofit: `pip install sai4rai-serve`

## Ready to close your governance gap?

Take the free 5-minute AI Maturity Assessment, then schedule a 30-minute discovery call. No pitch — we start with your situation.

→ [sairaisolutions.com/rai-assessment](https://sairaisolutions.com/rai-assessment)

→ [sairaisolutions.com/contact](https://sairaisolutions.com/contact)

Vijay Reddy · Founder, SAI RAI Solutions · AI Governance Lead & Principal SAFe SPC