AI Governance
Readiness Report

5 August 2026 Quick Scan Prepared for compliance@meridian-trust.example

SAFESPACE AI GOVERNANCE SCORE

32 / 100

Higher is better

MATURITY BAND

Emerging

PRIORITY DOMAIN

Shadow AI

Domain maturity

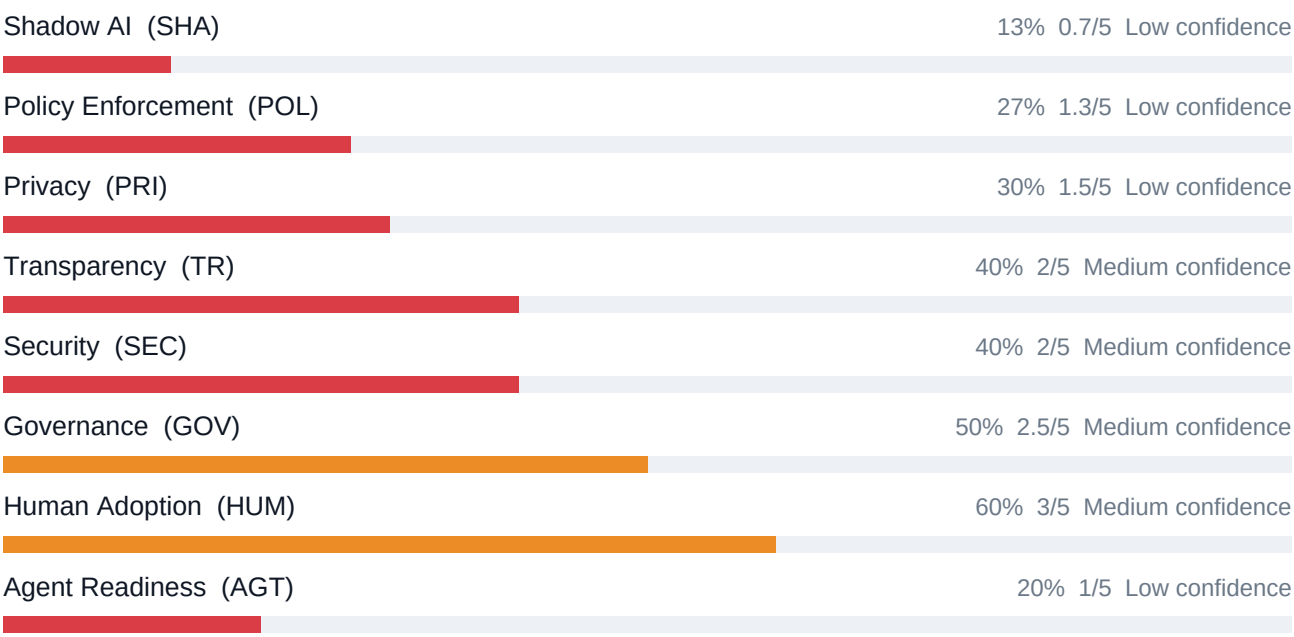

Recommended actions

One tailored next step per domain, weakest first. Directional guidance only.

Shadow AI - 13/100, Ad-hoc

Unmanaged AI usage is happening faster than you can see it. Shadow AI Discovery builds a complete inventory of AI tools in use - including personal-account and direct API usage - within 48 hours, and keeps rediscovering continuously.

-> Shadow AI Discovery

Agent Readiness - 20/100, Ad-hoc

Autonomous agents will inherit every gap above and act on them at machine speed. Register agents with distinct identities, scope their permissions and time-bound them, and require human approval for consequential actions before agent use scales.

-> Agent Governance

Policy Enforcement - 27/100, Emerging

Policy exists on paper but is not enforced at the network layer. A self-hosted tokenization proxy inspects outbound prompts and applies allow/deny, redaction, and tokenization rules before anything reaches an external model.

-> Tokenization Proxy

Privacy - 30/100, Emerging

Personal and confidential data is likely leaving through AI prompts, which conventional DLP does not inspect. Stateless tokenization replaces PII, source code, and contract content with reversible tokens before the prompt leaves your tenant.

-> Tokenization Proxy

Transparency - 40/100, Emerging

You cannot currently show what AI is in use, by whom, or produce Article 50 disclosure evidence. A hash-chained audit log records every AI interaction automatically and turns transparency into a queryable record.

-> Audit Log

Security - 40/100, Emerging

AI systems are not yet protected to the standard of the rest of your estate. Prioritise MFA on every AI workspace, centrally issued and rotated API credentials, and tamper-evident logging of AI interactions.

-> Audit Log

Governance - 50/100, Defined

Accountability for AI is unclear. Name an executive owner, stand up a cross-functional governance committee on a fixed cadence, and put AI risks into the enterprise risk register with named mitigation owners.

Human Adoption - 60/100, Defined

Adoption is being managed deliberately. Keep the request-and-feedback channel visible so demand for new tools reaches governance before it becomes Shadow AI.

Your responses

Every control assessed and the maturity level submitted, grouped by domain.

TRANSPARENCY (TR)

TR-001 AI Inventory AI asset inventory

1. Does the organization maintain a centralized, current inventory of every AI system and tool in use?

Partially implemented - some coverage, inconsistent (2/5)

TR-011 Explainability Article 50 disclosure

2. Do AI systems that interact with people disclose that they are AI, as EU AI Act Article 50 requires?

Partially implemented - some coverage, inconsistent (2/5)

SECURITY (SEC)

SEC-012 Monitoring AI audit logging

3. Are AI interactions logged in a tamper-evident way and retained for audit purposes?

Partially implemented - some coverage, inconsistent (2/5)

SEC-015 Incident Response AI incident response

4. Does your incident response plan explicitly cover AI and Shadow AI related incidents?

Partially implemented - some coverage, inconsistent (2/5)

PRIVACY (PRI)

PRI-001 Data Classification Sensitive data detection

5. Can you detect when personal or confidential data is sent to an AI system?

Partially implemented - some coverage, inconsistent (2/5)

PRI-011 Data Protection Intellectual property protection

6. Are source code, contracts, and strategic documents protected from being pasted into external AI tools?

Initial awareness - recognised, nothing in place yet (1/5)

GOVERNANCE (GOV)

GOV-007 Policy Management Enterprise AI policy

7. Is there an enterprise AI usage policy that is documented, communicated, and kept current?

Operational - documented and running in production (3/5)

GOV-011 Risk Management AI risk register

8. Are AI-related risks captured in a risk register with owners and mitigation plans?

Partially implemented - some coverage, inconsistent (2/5)

HUMAN ADOPTION (HUM)

HUM-001 AI Literacy AI awareness training

9. Do employees receive AI awareness training covering safe and unsafe use of AI tools?

Operational - documented and running in production (3/5)

SHADOW AI (SHA)

SHA-001 Discovery Browser discovery

10. Can you discover AI tools accessed from managed browsers and endpoints across the organization?

Initial awareness - recognised, nothing in place yet (1/5)

SHA-007 Usage Monitoring Personal AI accounts

11. Do you know what share of AI usage happens through personal, unmanaged accounts?

Don't know

SHA-013 Sensitive Prompt Detection Sensitive prompt detection

12. Can you detect prompts that contain sensitive or regulated data before they leave your network?

Initial awareness - recognised, nothing in place yet (1/5)

POLICY ENFORCEMENT (POL)

POL-001 Identity Enforcement AI allow/deny lists

13. Are approved and blocked AI tools enforced technically, not just documented in policy?

Partially implemented - some coverage, inconsistent (2/5)

POL-006 Content Protection Tokenization

14. Is sensitive data tokenized or pseudonymized before it is sent to an AI model?

Not implemented (0/5)

POL-010 Monitoring DLP enforcement

15. Does your DLP tooling cover conversational AI traffic, not only email and file transfer?

Partially implemented - some coverage, inconsistent (2/5)

AGENT READINESS (AGT)

AGT-001 Agent Identity Agent registration

16. Are AI agents and autonomous workflows registered with a distinct identity before deployment?

Initial awareness - recognised, nothing in place yet (1/5)

YOUR APPROACH

What tools or vendors are you currently considering to address this?

> Currently evaluating a mix of native data-loss prevention controls and third-party AI gateway proxies. Piloted an enterprise assistant with a small team; no organisation-wide decision yet.

Which of the following approaches would you consider?

> Middle security layer | Blocking all AI tools and providing our own LLM

How this score is calculated

This report follows the SafeSpace AI Governance Assessment Model. Each control is rated on a 0-5 maturity scale, from Not implemented to Optimized. Control scores are averaged into eight domains - Transparency, Security, Privacy, Governance, Human Adoption, Shadow AI, Policy Enforcement and Agent Readiness - and combined into the SafeSpace AI Governance Score (AGS) using weights that reflect enterprise risk: Shadow AI 18%, Policy Enforcement 18%, Privacy 15%, Transparency 14%, Security 13%, Governance 10%, Human Adoption 6% and Agent Readiness 6%. Higher is better. Maturity bands: Ad-hoc (0-20), Emerging (21-40), Defined (41-60), Managed (61-80), Optimized (81-100). Each domain also carries a confidence rating based on the evidence quality implied by the answers given. This assessment is self-reported, indicative only, and is not a certification, audit, or legal opinion.

Next step

A working session with our team turns this snapshot into a prioritized remediation plan mapped to your environment. Reply to the delivery email to schedule.