AI Governance
Readiness Report

5 August 2026 Full Assessment Prepared for compliance@meridian-trust.example

SAFESPACE AI GOVERNANCE SCORE

33 / 100

Higher is better

MATURITY BAND

Emerging

PRIORITY DOMAIN

Shadow AI

Domain maturity

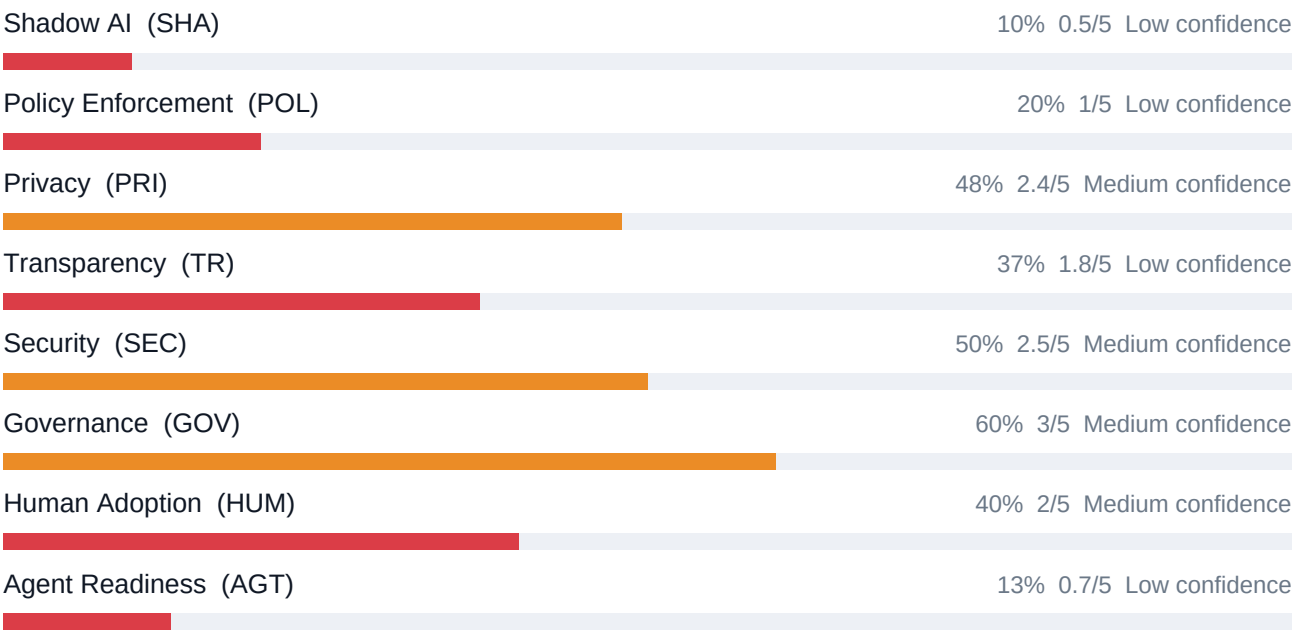

Recommended actions

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

Shadow AI - 10/100, 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 - 13/100, 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 - 20/100, Ad-hoc

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

Transparency - 37/100, 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

Human Adoption - 40/100, Emerging

Employees are adopting AI faster than the organization is teaching them how to do it safely. Deliver role-appropriate AI awareness training and give staff an approved tool for each common use case - bans without alternatives drive usage underground.

Privacy - 48/100, Defined

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

Security - 50/100, Defined

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 - 60/100, Defined

Governance structures are working. Tie them to evidence: link policy statements to the enforcement and audit records that prove they are being applied.

-> Audit Log

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-002 AI Inventory Business ownership

2. Is a named business owner assigned and recorded for each AI system in use?

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

TR-004 Usage Visibility Usage analytics

3. Can you report on who is using which AI tools, how often, and for what purpose?

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

TR-006 Usage Visibility Executive reporting

4. Do executives receive regular reporting on enterprise AI adoption and associated risk?

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

TR-009 Explainability Prompt traceability

5. Can you trace an individual AI interaction - prompt, model, and data used - after the fact?

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

TR-011 Explainability Article 50 disclosure

6. 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-001 Identity Security MFA for AI platforms

7. Is multi-factor authentication enforced for every corporate AI platform and workspace?

Managed and measured - monitored with defined owners (4/5)

SEC-004 API Security Secure API authentication

8. Are AI API keys and service credentials centrally issued, rotated, and scoped to least privilege?

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

SEC-007 Access Control Role-based access

9. Is access to AI systems granted by role, reviewed periodically, and revoked on leaver events?

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

SEC-009 Infrastructure Security Prompt injection protection

10. Do you have controls that detect or block prompt injection and model-abuse attempts?

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

SEC-012 Monitoring AI audit logging

11. 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

12. 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

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

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

PRI-004 Regulatory Compliance GDPR controls

14. Are GDPR obligations - lawful basis, DPIAs, data subject rights - applied to AI processing?

Managed and measured - monitored with defined owners (4/5)

PRI-006 Data Protection Data residency

15. Do you control and evidence where data processed by AI systems is stored and processed?

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

PRI-009 Privacy Governance Data retention

16. Are retention and deletion rules defined and enforced for AI prompts, outputs, and logs?

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

PRI-011 Data Protection Intellectual property protection

17. 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-001 Leadership Executive sponsor

18. Is there a named executive accountable for AI governance across the organization?

Managed and measured - monitored with defined owners (4/5)

GOV-003 Leadership AI governance committee

19. Does a cross-functional AI governance committee meet on a defined cadence?

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

GOV-007 Policy Management Enterprise AI policy

20. 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

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

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

GOV-018 Third-party Governance Vendor approval process

22. Do AI vendors and models go through a formal approval and due-diligence process before use?

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

HUMAN ADOPTION (HUM)

HUM-001 AI Literacy AI awareness training

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

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

HUM-004 Responsible AI Responsible AI certification

24. Are staff in higher-risk roles required to complete responsible-AI training or certification?

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

HUM-007 Business Adoption Department adoption

25. Is AI adoption tracked by department, with approved tools offered for common use cases?

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

HUM-010 Change Management Employee feedback

26. Is there a channel for employees to request AI tools or report AI-related concerns?

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

SHADOW AI (SHA)

SHA-001 Discovery Browser discovery

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

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

SHA-003 Risk Identification Unauthorized AI platforms

28. Are unauthorized AI platforms identified and risk-rated when they appear on the network?

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

SHA-007 Usage Monitoring Personal AI accounts

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

Don't know

SHA-010 Discovery AI API discovery

30. Can you detect direct AI API usage from applications, scripts, or developer tooling?

Not implemented (0/5)

SHA-013 Sensitive Prompt Detection Sensitive prompt detection

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

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

SHA-016 Usage Monitoring Continuous rediscovery

32. Is Shadow AI discovery continuous, rather than a one-off exercise?

Not implemented (0/5)

POLICY ENFORCEMENT (POL)

POL-001 Identity Enforcement AI allow/deny lists

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

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

POL-004 Content Protection Prompt inspection

34. Are outbound AI prompts inspected against policy before reaching an external model?

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

POL-006 Content Protection Tokenization

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

Not implemented (0/5)

POL-008 Content Protection Redaction

36. Is personal or confidential content redacted automatically when policy requires it?

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

POL-010 Monitoring DLP enforcement

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

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

POL-013 Remediation Automated policy remediation

38. When a policy violation occurs in AI usage, is remediation automated rather than manual?

Not implemented (0/5)

AGENT READINESS (AGT)

AGT-001 Agent Identity Agent registration

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

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

AGT-003 Human Oversight Human approval

40. Do consequential agent actions require documented human approval?

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

AGT-006 Delegated Authority Permission boundaries

41. Are agent permissions scoped and time-bound rather than inherited from a human account?

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

AGT-010 Agent Monitoring Agent audit trail

42. Is there a complete audit trail of actions taken by AI agents on your systems?

Not implemented (0/5)

AGT-014 Lifecycle Management Memory governance

43. Are agent memory and stored context governed, reviewed, and purged under defined rules?

Not implemented (0/5)

AGT-018 Multi-Agent Governance Cross-agent policy enforcement

44. Are the same policies enforced when agents call other agents or external tools?

Not implemented (0/5)

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.

ABOUT SAFESPACE AI

Say yes to AI - safely.

SafeSpace AI is a European, self-hosted privacy layer for enterprises that cannot afford to choose between productivity and control. It tokenizes sensitive data in every outbound prompt before it leaves your network and restores context in the response - and surfaces the Shadow AI already in use, so you can close the visibility gap before it becomes a breach.