



Sample deliverable

Sample: AI Governance Readiness Assessment

Sample document. All findings, scores, account identifiers, and company details are fictional demonstration data created to show report format and depth. This is not a real client engagement.

CLIENT	AI SYSTEMS	ASSESSED	42001 READINESS
Meridian Benefits Group (sample)	9 identified	Feb 2026	38%

Executive summary

Meridian has nine AI systems in use, of which **four were unknown to the security team** before this engagement. Two process member health information through a third-party summarization service with no data processing agreement covering AI use.

Against ISO/IEC 42001 Annex A, **38% of controls are in place**, concentrated in technical areas that AWS provides by default. Governance controls, where the standard places most of its weight, are largely absent: there is no AI policy, no approval path for new AI use cases, and no defined human oversight for the claims-triage model that influences benefit determinations.

Most urgent item. The claims-triage model affects benefit decisions without documented human review. This is both an ISO/IEC 42001 gap (A.9 human oversight) and a likely regulatory exposure. It is the one finding that warrants action before the policy work begins.

AI system inventory

SYSTEM	PURPOSE	PLATFORM	DATA	RISK
Claims triage model	Scores claims for routing and priority	Amazon SageMaker (fine-tuned)	Member health, claims history	Critical
Member chat assistant	Answers plan coverage questions	Amazon Bedrock, Claude	Plan documents, member profile	High
Document summarizer	Summarizes provider correspondence	Third-party SaaS	Member health information	High
Underwriting assistant	Drafts risk narratives	Amazon Bedrock	Group census data	Medium
Internal knowledge search	Policy and procedure lookup	Amazon Q	Internal documents	Medium
Marketing copy generation	Drafts campaign content	Third-party SaaS	Public content	Low
Developer code assistant	Code completion	IDE plugin	Source code	Medium
Call transcription	Transcribes member service calls	Amazon Transcribe	Call audio, member PII	High
Fraud pattern detection	Flags anomalous claims	Amazon SageMaker	Claims history	High

Systems 3, 6, 7, and 9 were not previously tracked by the security team.

ISO/IEC 42001 Annex A coverage

A.2 AI policy	no AI policy exists
A.3 Internal organization and roles	ownership informal, undocumented
A.4 Resources for AI systems	compute and data resources documented
A.5 AI system impact assessment	none performed
A.6 AI system life cycle	SageMaker registry used for two models
A.7 Data for AI systems	lineage tracked for in-house models only
A.8 Information for interested parties	no customer-facing AI disclosure
A.9 Use of AI systems	human oversight undefined for claims triage
A.10 Third-party and customer relationships	no AI terms in vendor agreements

NIST AI RMF function coverage

FUNCTION	COVERAGE	OBSERVATION
Govern	15%	No AI policy, approval path, or defined accountability. The largest gap and the prerequisite for the rest.
Map	45%	This engagement produced the first complete inventory. Context and impact documentation remains thin for third-party systems.
Measure	30%	Model evaluation performed at training time for in-house models; no ongoing monitoring of drift or output quality.
Manage	25%	No AI incident handling path. Bedrock Guardrails configured on the member assistant only.

Recommended sequence

PRIORITY	ACTION	ADDRESSES	EFFORT
1	Define and document human oversight for the claims triage model	A.9, Manage	1-2 weeks
2	Issue an AI policy with an approval path for new AI use cases	A.2, A.3, Govern	2-3 weeks
3	Add AI terms and a data processing agreement to the summarization vendor contract	A.10	Legal-dependent
4	Run impact assessments for the four highest-risk systems	A.5, Map	3 weeks
5	Enable Bedrock Guardrails across all Bedrock workloads and log invocations to CloudTrail	Measure, Manage	1 week
6	Publish customer-facing AI disclosure	A.8	2 weeks

Readiness, not certification. ISO/IEC 42001 certification must be performed by an accredited certification body. Model bias, fairness, and performance testing are separate technical engagements.