AI Risk Snapshot

CONFIDENTIAL — PREPARED FOR MERIDIAN HEALTH PARTNERS

Prepared by SkyLabsAI · [Sibel Karabulut, Founder & Principal AI Governance Advisor] Date: June 23, 2026

---

1. Executive Summary

Overall AI Governance Score: 30 / 100 **Risk Tier:** Critical

Meridian Health Partners operates three AI systems in production — including a clinical intake triage assistant and a third-party clinical documentation platform — with no formal AI governance program in place. No one is accountable for AI risk, none of the three systems has been classified under applicable regulatory frameworks, and the company's primary AI vendor was contracted without any risk review. Enterprise customers are already asking questions Meridian cannot answer, and regulators are asking them next.

Areas of Existing Capability

1. **PRIVACY** — Meridian's HIPAA compliance program provides a meaningful foundation. Patient data rights, retention schedules, and a Business Associate Agreement with the clinical documentation vendor are already in place and can be extended to address AI-specific requirements without starting from scratch.
2. **SECURITY** — Existing security infrastructure is well-positioned. AI system activity is logged, model deployment is access-controlled, and the 2025 HIPAA Security Risk Assessment explicitly covered AI systems. The gap is AI-specific threat readiness, not foundational security maturity.
3. **COMPLIANCE** — Leadership is aware of the evolving AI regulatory environment and is actively tracking developments. The organization has the awareness and seniority to act quickly once compliance obligations are formally mapped.

Priority Risks

1. **UNOWNED AI RISK.** No one at Meridian holds documented accountability for AI risk. If a regulatory inquiry, enterprise procurement challenge, or AI-related incident occurs, there is no designated owner to respond — and no documented policy to stand behind.

2. **UNCLASSIFIED AI SYSTEMS.** None of Meridian's three AI systems has been evaluated against applicable regulatory risk frameworks. The clinical triage assistant — which influences patient prioritization decisions — carries material regulatory exposure that has not been formally assessed.

3. **UNREVIEWED AI VENDOR.** Scribe AI, which processes sensitive clinical documentation using a third-party large language model, was onboarded without a governance or security review. It remains unassessed over a year into its production use.

Recommended Next Step: SkyLabsAI recommends a focused remediation engagement to address Meridian's highest-priority governance gaps before the next enterprise sales cycle or regulatory touchpoint. The six foundational gaps identified in this assessment — ownership, regulatory mapping, vendor risk, policy, system classification, and AI-specific incident readiness — cannot be addressed in parallel with routine operations and require structured program design, not point-in-time fixes.

---

2. Engagement Scope

AI systems assessed: Clinical Intake Triage Assistant, Claims Fraud Detection Model, Clinical Note Summarization (third-party LLM)

Stakeholders engaged: Lena Ortiz (VP, Compliance & Risk), Marcus Webb (Director of Engineering), Priya Shah (Head of Claims Analytics)

Documentation reviewed: HIPAA Security Risk Assessment (2025), Scribe AI Inc. Master Services Agreement

Engagement period: June 15–23, 2026

All findings and assessments in this report are based on information provided by Meridian Health Partners' management and documentation made available during the engagement period. SkyLabsAI has not independently audited or verified the underlying systems, controls, or processes described. Findings reflect professional judgment informed by structured stakeholder discussions and document review, and should be read accordingly.

---

3. AI System Inventory

System	Business Owner	Built or Procured	Data Sensitivity	Status
Clinical Intake Triage Assistant	Marcus Webb	In-house	High	Production

System	Business Owner	Built or Procured	Data Sensitivity	Status
Claims Fraud Detection Model	Priya Shah	In-house	Medium	Production
Clinical Note Summarization	Marcus Webb	Scribe AI Inc.	High	Production

---

4. AI Governance Risk Scorecard

Governance Domain	Score	Rating
Governance & Accountability	1.7	Immediate Attention Required
Security	3.1	Developing
Regulatory Compliance	2.0	Immediate Attention Required
Privacy	3.2	Developing
Vendor Risk	1.5	Immediate Attention Required
AI Operations	1.8	Immediate Attention Required
Overall AI Governance Score	**30 / 100**	**Critical**

AI Governance Risk Profile

Governance & Accountability	**Immediate Attention Required**
Security	**Developing**
Regulatory Compliance	**Immediate Attention Required**
Privacy	**Developing**
Vendor Risk	**Immediate Attention Required**
AI Operations	**Immediate Attention Required**

Immediate Attention Required | Developing | Established

---

5. Findings by Governance Domain

Governance & Accountability — Immediate Attention Required

- **Current State:** No one at Meridian holds documented accountability for AI risk. The organization's closest approximation to an AI governance policy is an informational email distributed in 2024 that was never formalized, enforced, or updated.
- **Key Observations:** Leadership acknowledged during interviews that accountability for AI risk is informally assigned but undocumented. Issues are escalated informally with no defined path to leadership. Compliance was not involved in the decision to onboard the clinical note summarization vendor — learning of it only after the contract was signed. No formal review process exists for AI use cases prior to production deployment, and leadership has never received a structured AI risk briefing.
- **Gap:** The organization has no formal owner, no enforceable policy, no escalation structure, and no governance gate for new AI deployments.
- **Risk if Left Unaddressed:** As AI deployment decisions continue to be made without a consistent governance framework, Meridian's regulatory, reputational, and contractual exposure will grow with each new use case. When a regulator, enterprise customer, or auditor asks about AI governance, there is no one positioned to respond and no documentation to stand behind.

Security — Developing

- **Current State:** Meridian's security fundamentals are well-established, driven primarily by HIPAA compliance requirements. The organization has meaningful security infrastructure that has been extended to cover its AI systems at a baseline level.
- **Key Observations:** AI systems were explicitly scoped into the 2025 HIPAA Security Risk Assessment. Model deployment requires authorization and is access-controlled. AI system inputs and outputs are logged through the existing compliance-grade infrastructure. However, no adversarial testing — such as prompt injection testing or misuse simulation — has been conducted on any AI system, and the existing incident response plan does not address AI-specific threat scenarios.
- **Gap:** The organization's security posture is built for traditional infrastructure risk, not the distinct threat surface created by large language models and automated decision systems.
- **Risk if Left Unaddressed:** Meridian has no tested response plan for AI-specific security incidents, including prompt injection attacks, model-mediated data leakage, or malicious manipulation of AI-generated clinical outputs.

Regulatory Compliance — Immediate Attention Required

- **Current State:** Regulatory awareness exists at the leadership level, but no formal mapping of obligations to specific AI systems has been completed.
- **Key Observations:** Meridian has not assessed which AI-specific regulations apply to its operations or how its three production systems are classified under applicable frameworks. None of the systems have been evaluated under the EU AI Act's risk classification requirements. No designated owner monitors the AI regulatory landscape for developments relevant to Meridian's business. Legal has reviewed the organization's

AI systems for general HIPAA compliance — confirmed through documentation — but not through an AI-specific regulatory lens.
- **Gap:** Meridian's regulatory exposure across all three AI systems is currently unmapped, not merely unaddressed. The organization cannot currently articulate its compliance obligations or demonstrate readiness to an auditor, regulator, or enterprise customer.
- **Risk if Left Unaddressed:** Enterprise customers are already including AI governance requirements in procurement processes. Meridian risks losing or delaying deals before regulatory enforcement becomes a concern — and the EU AI Act's obligations for systems like the clinical triage assistant are approaching, not distant.

Privacy — Developing

- **Current State:** Meridian's HIPAA compliance program provides a strong foundation for AI privacy governance, with meaningful controls already in place for health data handling.
- **Key Observations:** Standard data retention schedules, patient rights processes, and a Business Associate Agreement with the clinical documentation vendor are all confirmed in place through documentation review. Health data processed by AI systems is subject to existing heightened safeguards. However, no privacy impact assessment specific to AI system use has been conducted, and the clinical triage assistant's role in influencing patient prioritization decisions has not been formally evaluated as a potential significant-effect automated decision under applicable privacy frameworks.
- **Gap:** HIPAA-level privacy controls are in place but have not been extended to address AI-specific privacy obligations, including automated decision-making assessments.
- **Risk if Left Unaddressed:** The clinical triage assistant's influence on patient callback prioritization represents the kind of automated decision-making that draws regulatory scrutiny once identified — and Meridian currently has no documentation demonstrating it was evaluated.

Vendor Risk — Immediate Attention Required

- **Current State:** Scribe AI, Meridian's only third-party AI vendor, processes sensitive clinical documentation and has operated without any formal risk assessment since it was onboarded.
- **Key Observations:** Scribe AI was selected and contracted without a governance or security risk review. The executed contract addresses data protection at a standard HIPAA level but contains no provisions governing AI-specific risks: model update notification, training data use restrictions, or liability allocation for AI-mediated errors. No reassessment has occurred since contract execution. Meridian has not confirmed whether Scribe AI uses client data to train or improve its models. No operational contingency plan exists for a Scribe AI outage or security incident.
- **Gap:** A vendor with access to high-sensitivity clinical data has zero ongoing governance oversight and no AI-specific contractual protections.
- **Risk if Left Unaddressed:** If Scribe AI experiences a security incident, is found to have used Meridian's clinical data without authorization, or discontinues service, Meridian has

no contractual recourse, no early warning mechanism, and no documented operational fallback.

AI Operations — Immediate Attention Required

- **Current State:** Both in-house AI models were deployed following functional testing but have received no AI-specific oversight since launch.
- **Key Observations:** Neither the clinical triage assistant nor the claims fraud model underwent formal bias, fairness, or accuracy testing prior to production deployment. No ongoing performance monitoring exists for either system — the fraud model has not been reviewed since its initial launch over a year ago. Human review of triage assistant outputs is practiced informally but is not a documented, auditable requirement. No process governs when or how underperforming models are retrained, updated, or retired. No testing or monitoring records exist for either system, which is itself a material governance finding.
- **Gap:** Two production AI systems — including one that influences clinical outcomes — have operated without validation or oversight for over a year.
- **Risk if Left Unaddressed:** An unmonitored claims fraud model may be systematically misclassifying claims without detection. A clinical triage assistant with informal-only human review carries direct patient-safety exposure if its decision quality degrades over time. Both conditions represent material risk that is currently invisible to the organization.

---

6. Regulatory Exposure

Framework	Exposure	Risk Level
EU AI Act (Annex III; Articles 6, 9)	The clinical triage assistant — which influences patient prioritization decisions — warrants classification as a high-risk AI system under Annex III healthcare provisions. None of Meridian's systems have been classified, and Article 9 risk management obligations remain unaddressed	High
NIST AI Risk Management Framework (GOVERN 1.1, MAP 1.1, MEASURE 2.5)	No governance function (GOVERN) is established, no risk identification process (MAP) exists, and no measurement or monitoring (MEASURE) is performed on production models	High
HIPAA / Health Data Regulation	Baseline program is solid and extends to AI systems at a general level; AI-specific obligations — including privacy impact assessment for automated triage prioritization — have not been addressed	Medium

Framework	Exposure	Risk Level
Third-Party AI Vendor Obligations	The Scribe AI contract lacks provisions governing model training use, update notification rights, and AI-specific liability — exposing Meridian to unmitigated contractual risk	High

---

7. Prioritized Action Roadmap

Priority	Recommendation	Business Risk Addressed	Owner	Timeframe
1	Designate a named AI risk owner with documented executive accountability	Governance; NIST AI RMF GOVERN 1.1	VP, Compliance & Risk	30 days
2	Conduct a regulatory applicability assessment to map obligations across all three AI systems	Compliance; EU AI Act Articles 6 and 9	VP, Compliance & Risk	30 days
3	Perform a retroactive governance and security review of Scribe AI; renegotiate AI-specific contractual terms	Vendor Risk; third-party liability	Director of Engineering	30 days
4	Establish a formal AI governance policy covering acceptable use, deployment approval, and executive escalation	Governance	VP, Compliance & Risk	30 days
5	Classify all three AI systems under applicable regulatory risk frameworks	Compliance; EU AI Act Annex III	VP, Compliance & Risk	60 days
6	Commission pre-deployment bias, fairness, and accuracy evaluations for both in-house models; apply retroactively to production systems	AI Operations	Director of Engineering	60 days
7	Develop an AI-specific security incident response plan addressing model misuse, adversarial attack, and data leakage scenarios	Security	Director of Engineering	60 days
8	Implement a formal pre-deployment review and approval process for new AI use cases	Governance	VP, Compliance & Risk	60 days

Priority	Recommendation	Business Risk Addressed	Owner	Timeframe
9	Establish ongoing performance and drift monitoring for production AI models with defined intervention thresholds	AI Operations	Director of Engineering	90 days
10	Commission a privacy impact assessment for the clinical triage assistant and the Scribe AI integration	Privacy	VP, Compliance & Risk	90 days

The Path Forward: The findings in this assessment reflect a common pattern in organizations deploying AI ahead of governance: capable teams, meaningful technical infrastructure, and a leadership team that wants to do the right thing — without the specialized framework, regulatory mapping, and governance design required to manage AI risk at an institutional level. Addressing these gaps requires more than internal effort. It requires structured program design, regulatory expertise, and ongoing oversight as the AI risk landscape continues to evolve.

SkyLabsAI recommends initiating a remediation engagement to address priorities 1 through 7 before Meridian's next enterprise procurement cycle or regulatory touchpoint. Following completion, an annual AI Risk Snapshot — or an accelerated reassessment upon the introduction of new AI systems or material regulatory developments — will ensure Meridian's governance posture remains calibrated to an environment that is changing faster than any single assessment can fully capture.

Appendix A — About This Assessment

The AI Risk Snapshot is SkyLabsAI's proprietary AI governance risk assessment, developed by practitioners with direct experience in enterprise AI risk management, regulatory compliance, and technology governance across regulated industries.

Each assessment evaluates an organization's AI governance posture across six domains — Governance & Accountability, Security, Regulatory Compliance, Privacy, Vendor Risk, and AI Operations — using SkyLabsAI's proprietary maturity framework, calibrated against the leading standards shaping AI governance globally: the NIST Artificial Intelligence Risk Management Framework (AI RMF), the EU Artificial Intelligence Act, and ISO/IEC 42001:2023.

Findings are developed through structured stakeholder engagement, documentation review, and the professional judgment of SkyLabsAI's assessment team. Scores reflect the organization's current governance maturity relative to the standard of care expected by enterprise customers, regulators, and auditors operating in the AI-enabled enterprise environment of 2026.

SkyLabsAI's assessment methodology is proprietary and is updated continuously as the AI regulatory landscape evolves.

Appendix B — Engagement Record

Participant / Document	Role	Date
Lena Ortiz	VP, Compliance & Risk	June 16, 2026
Marcus Webb	Director of Engineering	June 17, 2026
Priya Shah	Head of Claims Analytics	June 17, 2026
HIPAA Security Risk Assessment (2025)	Documentation	Reviewed during engagement
Scribe AI Inc. Master Services Agreement	Documentation	Reviewed during engagement

Appendix C — Engagement Terms and Limitations

This report constitutes the deliverable of an AI Risk Snapshot engagement, as specified in the executed Statement of Work between SkyLabsAI and Meridian Health Partners. It reflects a point-in-time assessment and does not constitute legal advice, a security audit, or certification of compliance with any law, regulation, or standard. SkyLabsAI's liability in connection with this engagement is governed exclusively by the Statement of Work.

Contact & Next Steps

Questions about this report? Lead Advisor · contact@skylabsai.com·https://calendly.com/sibel-karabulut-skylabsai/30min?month=2026-07 **Your Executive Readout Call** is scheduled for June 23, 2026. **If you would like to move forward:** Your AI Risk Snapshot fee is fully creditable toward a SkyLabsAI Consulting engagement or AI Governance Dashboard subscription if initiated within 30 days of this report's delivery date. SkyLabsAI will provide a focused remediation proposal at no additional charge — reach out to the contact above to request it.