HealthEquity AI Governance Fact Sheet

For Client, Partner, and Prospect Inquiries

Last updated: June 09, 2026 | v 1.4.0

Purpose. This fact sheet summarizes how HealthEquity governs, secures, and operates AI across our products and services. It is intended as a single reference for client, partner, and prospect inquiries about our applied AI practices. For a quick reference, see Sections I, II, and VI.

Bottom Line. HealthEquity applies AI to defined use cases under a formal governance structure, with security, privacy, legal, risk, and data engineering controls mapped to applicable laws, the NIST AI Risk Management Framework, the Open Worldwide Application Security Project (OWASP) LLM and Agentic Top 10, and HealthEquity-specific internal standards. Governance is proportionate to risk: controls, oversight, and testing are scaled based on the risk tier assigned to each use case.

I. Responsible AI Principles

AI use cases at HealthEquity are reviewed against seven principles before approval and re-evaluated over their lifecycle.

Principle	What It Means in Practice
Transparency	HealthEquity provides notice of AI involvement in member- or client-facing interactions where appropriate to the use case and where required by law, contract, or company policy.
Explainability	AI-assisted decisions are designed to be explainable in context. HealthEquity designs AI systems so that their outputs and decision factors can be understood at a level appropriate to the use case and audience.
Human Oversight	Consequential decisions — such as claims adjudication, fraud determinations, or eligibility decisions — retain human oversight. Human-in-the-loop requirements are defined by the risk tier assigned to each use case, and members can transfer to a human agent in conversational channels.
Fairness	When appropriate, AI use cases are evaluated for potential bias across training data, model design, and outputs. Controls are applied proportionate to the risk tier of the use case to promote

Principle	What It Means in Practice
	fair and non-discriminatory outcomes.
Robustness & Safety	AI systems are designed to operate reliably under expected conditions and to fail safely under unexpected ones. Models and data are secured, empirically validated, and monitored for drift, degradation, and adversarial manipulation.
Validity & Reliability	AI systems are tested to confirm they perform as intended before deployment. Ongoing monitoring validates that outputs remain accurate, consistent, and fit for purpose over time.
Accountability	Named owners are accountable for each use case across compliance, ethics, and non-discrimination requirements.

II. Governance Structure

HealthEquity operates a formal AI Governance Council with executive, legal, compliance, security, risk, privacy, and technology representation. The Council:

- Reviews and approves AI use cases before deployment, with the depth of review scaled to the use case's risk profile.
- Maintains a risk-tiered inventory of AI use cases.
- Assigns reviewed use cases an internal risk tier, which drives the required depth of controls, testing cadence, and human oversight.
- Conducts ongoing monitoring and performance review, with reporting to executive leadership.
- Oversees AI training and education programs to foster AI literacy and ensure personnel are informed about safe, authorized use of AI across the organization.
- Governs AI use cases across their full lifecycle — from design and approval through deployment, monitoring, and retirement.
- Monitors emerging AI technologies, regulatory developments, and industry practices to inform ongoing program evolution.
- Incorporates feedback from clients, members, and internal stakeholders to continuously improve AI governance practices.

Oversight model

HealthEquity's AI governance operates within a three-lines-of-defense model:

- **First line:** Business and technology teams that build, deploy, and operate AI systems are responsible for adhering to governance requirements and applying controls appropriate to the risk tier.
- **Second line:** The AI Governance Council provides independent oversight, sets policy, and reviews higher-risk and external-facing use cases. Roles and responsibilities are defined across Council functions.
- **Third line:** Internal audit provides independent assurance that governance controls are operating as intended.

Bottom Line. AI governance at HealthEquity is proportionate to risk. The depth of Council review, the controls applied, and the testing cadence are determined by the risk tier assigned to each use case.

III. Control Domains

HealthEquity's AI governance program applies controls across five domains. Each has a named executive owner and a defined set of criteria applied proportionate to risk. The commitments below describe HealthEquity's posture; the supporting practices illustrate how these commitments are operationalized.

Security

Commitment. HealthEquity secures AI systems against unauthorized access, adversarial manipulation, and data leakage, extending our enterprise Cybersecurity program to address AI-specific threats.

Supporting practices

- **AI-specific threat modeling** covering prompt injection, data poisoning, model extraction, evasion attacks, and risks specific to agentic architectures, retrieval-augmented generation (RAG), and model output accuracy.
- Input validation and output screening to prevent sensitive data leakage and harmful content from AI outputs.
- **Authentication and authorization** of model endpoints using role- and attribute-based access controls. AI agents and automated workflows are treated as scoped non-human identities with least-privilege access.

- **Encryption** of data in transit (TLS 1.2+) and at rest (AES-256), with documented key management.
- Logging and monitoring of AI interactions with user identity, model version, and timestamp.
- **Adversarial testing** of applicable production AI at least quarterly via our internal red-team program.
- **Supply chain validation** for third-party models and libraries, including provenance and AI Bill of Materials (AI-BOM) requirements.
- **AI-specific incident response** procedures integrated with the enterprise incident response plan.
- **Continuous monitoring** for unauthorized AI usage (shadow AI) and data loss prevention, supported by ongoing AI asset discovery across the environment.

Privacy

Commitment. HealthEquity embeds privacy into AI system design and enforces data protection throughout the AI lifecycle, ensuring personal data is processed lawfully, minimally, and transparently.

Supporting practices

- **Data inventory and classification** for personal data processed by AI systems, categorized by sensitivity (PII, PHI, PCI).
- **Privacy Impact Assessments** completed for appropriate AI use cases, documenting data flows, purposes, and privacy risks.
- **Data minimization and purpose limitation** — AI systems process only the minimum necessary data for a clearly stated purpose.
- **Deidentification** of training and inference data where feasible; synthetic data is evaluated for quality and bias impact when used.
- **Consent, notice, and disclosure mechanisms** aligned to applicable legal requirements and industry best practices.
- **An alternative service path** is available where required by law, contract, or the applicable use case design for members or clients who prefer not to engage with AI-assisted interactions.
- **Sub-processor governance** with Data Processing Agreements and, for PHI, BAAs in place.

Legal & Regulatory

Commitment. HealthEquity reviews AI use cases against applicable law and contractual obligations before deployment, and maintains the legal and contractual infrastructure to support responsible AI use.

Supporting practices

- **Regulatory assessment** against HIPAA, GLBA, ERISA, IRS, ACA, and applicable U.S. state AI and privacy laws.
- **IP and licensing review** for model inputs, outputs, and training data.
- **Vendor contracts** include applicable AI-specific terms, data processing agreements, and appropriate allocation of liability for AI errors.
- **Disclosures** for AI involvement in member-facing interactions and any consequential decisions, with human-in-the-loop requirements defined by risk tier.
- **Retention schedules** for AI inputs, outputs, and training data, aligned with legal and regulatory requirements.

Risk Management

Commitment. HealthEquity identifies, assesses, and manages AI-related risks within the broader enterprise risk framework, with controls proportionate to the risk profile of each use case.

Supporting practices

- **Risk assessment** — Impacts to enterprise risks and the company risk profile are assessed holistically.
- **Risk appetite alignment** — Risks are evaluated against organizational appetites before approval.
- **AI risk controls** documented for model failure, degradation, and unexpected behavior.
- **Business continuity** and fallback procedures if an AI system becomes unavailable.
- **Drift and degradation monitoring** with defined response triggers and key risk indicators that may include prompt updates, control changes, rollback, retraining, or escalation.
- **AI/model registry** entries for tracked AI use cases, with risk tiers and controls.

Data Governance & Engineering

Commitment. HealthEquity governs the data that feeds AI with the same rigor applied to the models themselves, ensuring data quality, lineage, and bias controls are in place throughout the

data lifecycle.

Supporting practices

- **Data quality assessment** with documented completeness, accuracy, and consistency metrics for both training and inference data.
- **Documented pipeline architecture** with end-to-end schema, transformation, and lineage documentation.
- **Bias detection in training data** using statistical analysis and review procedures to identify and mitigate bias prior to deployment.
- **Alignment with enterprise data governance** policies and standards across retention, archival, and reproducibility (documented seeds, versions, and configurations).
- **Automated validation and scalability testing** to ensure pipelines meet quality and load requirements.

IV. Framework Alignment

HealthEquity's AI control set aligns with and is informed by the following:

Framework	How We Align
HIPAA Privacy and Security Rules	Minimum-necessary, deidentified PHI access by AI systems; DLP scanning of AI outputs for PHI/PII leakage; audit logging of AI interactions with user identity; documented risk assessments for AI processing deidentified PHI; BAAs with third-party AI vendors handling PHI; deidentified or synthetic training data; AI-specific incident response procedures.
NIST AI Risk Management Framework (AI RMF 1.0)	Formal AI governance body with defined roles; inventory of AI use cases with risk tiers; risk categorization using the NIST AI RMF taxonomy; continuous performance and drift monitoring; established incident management procedures for AI-related events.
HealthEquity Internal AI Governance Standard	AI Governance Council review scaled to risk; prohibition on unapproved (shadow) AI with continuous detection; guard-proxy enforcement for production AI traffic; quarterly red-team testing of applicable production AI systems; risk tiering of models; validated AI supply chain (AI-BOM); and defined remediation SLAs for critical findings.

Framework	How We Align
OWASP LLM and Agentic Top 10	HealthEquity implements controls aligned with the OWASP LLM and Agentic Top 10, including prompt injection protection, robust input validation, monitoring for model abuse, and documented procedures for agentic system oversight — integrated across both development and production environments.

V. Applied AI in Our Products

The following are representative examples of AI capabilities currently in production at HealthEquity. Each operates under the governance framework described above and is approved, inventoried, and monitored per the controls in Sections I–III.

Conversational Voice and Chat Support

What it does. Members can speak or chat naturally to check account balances, get transaction decline explanations, report lost or stolen cards, and check reimbursement claim status, instead of navigating touch-tone menus or waiting for an agent. Rollout is phased throughout 2026.

How it is built. Conversational AI is delivered through an enterprise AI agent platform designed for auditable, traceable interactions, with the underlying chat infrastructure running on AWS Connect. The platform maintains ISO 27001 and SOC 2 compliance and uses encrypted tokens for member data retrieval.

How members are protected.

- **Active authentication:** Passkey (fingerprint, Face ID, or PIN) may be required before a member can access AI-powered self-service in IVR or Chat.
- **Passive fraud detection:** A specialized voice security and authentication platform runs transparently in the background, analyzing voice patterns to detect fraud hubs, deepfakes, and voice spoofing.
- **Human handoff with context:** Members can transfer to a human agent at any time; when chatting, the agent receives the full conversation transcript.
- **Scope of data:** The IVR and Chat AI accesses member-specific account information, including member ID, which is why passkey authentication is required up front.

Expedited Claims (EZ Receipts)

What it does. Expedited Claims accelerates reimbursement claim processing for FSA, HRA, and Limited Purpose FSA accounts. A member photographs a receipt; the system uses OCR to extract

text, NLP to identify merchant, date, item, and amount, and rules-based AI to determine eligibility against plan rules. Claims that previously took days now are resolved in under two minutes.

Privacy by design. The Expedited Claims AI evaluates item eligibility without relying on member identity as a decision input. Personal data exposure to the AI component is minimized and governed under HealthEquity's standard privacy and security controls. Authentication occurs at the app login layer, not within the AI decisioning layer.

VI. Quick-Answer FAQ

Answers to the questions most frequently raised by clients and partners.

Question	Answer
Does HealthEquity have a formal AI governance program?	Yes. An executive-level, cross-functional AI Governance Council oversees AI use cases, which are inventoried, risk-tiered, and monitored. The depth of review and controls applied are proportionate to the risk profile of each use case.
Is identifiable PHI used to train AI models?	No. Training data is deidentified, synthetic, or aggregated. Any AI use case that processes PHI requires a documented risk assessment, a BAA with the relevant vendor, and minimum-necessary access controls.
How is AI output prevented from leaking sensitive data?	AI outputs are screened for PHI, PII, and confidential content before leaving the model boundary.
How often is AI tested for security weaknesses?	Applicable production AI systems undergo internal adversarial (red-team) testing at least quarterly, covering prompt injection, data poisoning, model extraction, and evasion.
Do humans remain in the loop?	Yes, for consequential decisions. Human-in-the-loop requirements are defined by the risk tier assigned to each use case; members can transfer to a human agent in conversational channels.
How do you prevent unapproved or shadow AI?	Unapproved AI use is prohibited by policy and continuously monitored. Production AI traffic is routed through our AI guard proxy.
What frameworks do you align to?	HIPAA Privacy and Security Rules, the NIST AI Risk Management Framework 1.0, the OWASP LLM and Agentic Top 10, and HealthEquity's internal AI governance standard.

Question	Answer
Do third-party AI vendors have BAAs?	Yes. Any third-party AI vendor processing PHI is under a Business Associate Agreement, and contracts include AI-specific terms, data processing agreements, and liability provisions.
How is model drift detected?	Models are continuously monitored for performance degradation and drift, with defined retraining triggers and escalation procedures.
Is there an AI-specific incident response plan?	Yes. AI-specific incident response procedures are documented, tested, and integrated with the enterprise incident response program.
How does HealthEquity address AI-related employee training?	Personnel using AI complete mandatory training on safe, authorized use prior to access. The AI Governance Council oversees ongoing education programs to foster AI literacy and ensure awareness of governance requirements, acceptable use, and emerging risks.
Can members opt out of AI-assisted interactions?	Yes. Members can transfer to a human agent at any time in conversational channels. Where applicable, an alternative service path is available in accordance with law and contractual requirements.
How does HealthEquity manage AI across the full lifecycle?	AI use cases are governed from design and approval through deployment, ongoing monitoring, and retirement. Risk tiers, controls, and performance requirements are re-evaluated at each stage.

VII. Glossary

Acronyms and abbreviations used in this document.

Acronym	Definition
ACA	Affordable Care Act
AI-BOM	AI Bill of Materials
BAA	Business Associate Agreement
DLP	Data Loss Prevention

Acronym	Definition
ERISA	Employee Retirement Income Security Act
FSA	Flexible Spending Account
GLBA	Gramm-Leach-Bliley Act
HIPAA	Health Insurance Portability and Accountability Act
HRA	Health Reimbursement Arrangement
IRS	Internal Revenue Service
IVR	Interactive Voice Response
NLP	Natural Language Processing
NIST	National Institute of Standards and Technology
OCR	Optical Character Recognition
OWASP	Open Worldwide Application Security Project
PCI	Payment Card Industry
PHI	Protected Health Information
PII	Personally Identifiable Information
RAG	Retrieval-Augmented Generation
RBAC / ABAC	Role-Based / Attribute-Based Access Control
SLA	Service Level Agreement

About this document. *This fact sheet is maintained by HealthEquity for use in client, partner, and prospect communications regarding our applied AI practices. It summarizes current program posture at a point in time and is not intended to modify contractual commitments. For questions or tailored inquiries, contact your HealthEquity relationship manager or email security@healthequity.com. For privacy practices, see healthequity.com/privacy.*