SHADOW AI LABS

DELIVERABLE 04 · ACCEPTABLE USE POLICY

Acceptable Use Policy for Artificial Intelligence Tools.

A policy drafted to operational reality. Sanctioned tools enumerated, prohibited use defined, no-blame reporting path established.

PREPARED FOR

Northgate Family Medicine

ENGAGEMENT

SAL-2026-014

EFFECTIVE DATE

April 22, 2026

VERSION · NEXT REVIEW

1.0 · April 2027

CONTENTS

Eight sections, one acknowledgment form. Twelve pages total.

HOW TO READ THIS DOCUMENT

The substance for most employees lives in §3 (sanctioned tools) and §4 (prohibited use). The substance for compliance teams lives in §5 (data handling) and §8 (review cadence). The substance for HR lives in §7 (training) and the acknowledgment form. The CIO, Security Officer, and IT Steering Committee should read the whole thing.

§1 · PURPOSE AND SCOPE

What this policy governs, who it applies to, and under what authority.

This policy governs the use of Artificial Intelligence (AI) tools by all employees, contractors, and authorized agents of Northgate Family Medicine ("Northgate"). It establishes the standards under which AI tools may be evaluated, sanctioned, and used in the course of clinical, administrative, billing, and operational work.

§1.1 PURPOSE

The purpose of this policy is to **(a)** protect Protected Health Information (PHI) and other Confidential data from unauthorized disclosure to AI vendors; **(b)** maintain compliance with applicable regulatory frameworks including HIPAA (45 CFR Parts 160 and 164), the cyber insurance AI Security Rider issued by Northgate's carrier, and the NIST AI Risk Management Framework; and **(c)** enable employees to use sanctioned AI tools confidently and productively, with clear boundaries around what is permitted.

§1.2 APPLICABILITY

This policy applies to all use of AI tools that occurs **(i)** on Northgate-provided devices or networks; **(ii)** using Northgate-provided credentials or accounts; or **(iii)** in the course of work performed on behalf of Northgate, regardless of the device or network used.

Personal use of AI tools on personal devices and personal time, outside the scope of Northgate work, is not governed by this policy except where personal AI use intersects with Northgate data — see §4.1(c).

§1.3 AUTHORITY

This policy is issued under the authority of the Chief Information Officer (CIO), with oversight by the IT Steering Committee. The CIO has documented stop-authority over any AI tool use that creates material risk to PHI, regulatory compliance, or business continuity.

§1.4 RELATIONSHIP TO OTHER POLICIES

This policy supersedes the AI-related provisions of the Employee Handbook §7.4 ("Approved Software Use") as of the effective date. It does not modify the HIPAA Privacy and Security policies, the Acceptable Use Policy for general IT resources, or any provisions of executed employment agreements. Where this policy and a more specific clinical workflow policy appear to conflict, the more specific policy controls — escalate questions to the Security Officer.

§2 · DEFINITIONS

The terms used throughout this policy, defined for unambiguous application.

For purposes of this policy, the following definitions apply. Terms are listed in the order they first appear in the policy body.

"AI Tool" — any software application, service, browser extension, embedded feature, or API that uses machine learning, large language models, generative AI, or similar technologies to produce outputs based on input data. Includes chat-based assistants (e.g., ChatGPT, Claude, Gemini, Microsoft Copilot), AI scribes and transcription tools, AI-enhanced features within otherwise non-AI software (e.g., Outlook AI summarization), and AI-powered browser extensions.

"Sanctioned Tool" — an AI tool that appears on the Sanctioned Tools List (§3) and may be used within the conditions stated for that tool.

"Consumer AI" — an AI tool accessed via a consumer account (free-tier or paid-personal-subscription) rather than an enterprise contract held by Northgate. Examples: a personal ChatGPT Plus subscription; a free Otter.ai account.

"PHI" — Protected Health Information as defined at 45 CFR 160.103. Includes any individually identifiable health information transmitted or maintained in any form — patient names, medical record numbers, diagnosis codes, treatment plans, billing information, and clinical narratives that could identify a patient.

"Confidential Data" — any non-public information whose disclosure could harm Northgate or its patients, including but not limited to PHI, Personally Identifiable Information (PII) of patients or employees, payroll data, vendor contracts, strategic plans, security configurations, and credentials.

"BAA" — Business Associate Agreement as required under 45 CFR 164.504(e). Required between Northgate and any AI vendor that processes PHI on Northgate's behalf.

"DPA" — Data Processing Agreement. Required between Northgate and any AI vendor that processes Confidential Data other than PHI on Northgate's behalf.

"Sanctioned Tools List" — the enumeration in §3 of AI tools approved for use within Northgate, maintained by the CIO and reviewed quarterly by the IT Steering Committee.

§3 · SANCTIONED TOOLS LIST

AI tools approved for use, with the conditions that apply to each.

Inclusion on this list reflects that (a) a Business Associate Agreement or Data Processing Agreement is on file with the vendor where required; (b) the vendor has been subjected to a security review per §5; and (c) the use case described is governed by appropriate operational controls.

§3.1 CLINICAL WORKFLOW TOOLS

TOOL	APPROVED USE	CONDITIONS
EHR-embedded AI scribe (Epic)	Draft clinical documentation from in-visit dictation	Clinician must review and edit before signing into chart. BAA on file. No raw audio retained beyond 30 days.
Dragon Medical One AI	Voice-to-text for clinical documentation	Same clinician-review requirement. BAA on file.
Otter.ai Enterprise+ (post-BAA)	Transcription of clinical and intake calls	Permitted only after BAA execution (expected May 2026). Patient consent required per existing intake protocol.

§3.2 ADMINISTRATIVE AND PRODUCTIVITY TOOLS

TOOL	APPROVED USE	CONDITIONS
Microsoft Copilot E5	General productivity: email drafting, document summarization, meeting notes, ad-hoc Q&A	BAA via Microsoft Cloud Agreement. May process PHI within M365 boundaries. Do NOT paste PHI into Copilot Chat from external sources.
Grammarly Business	Writing assistance for non-clinical text	No PHI. Limited to administrative, marketing, and external communication where no patient information is present.
Calendly AI scheduling	Patient and staff scheduling assistance	No PHI in event titles or descriptions. Patient names allowed; clinical details prohibited.

§3.3 BILLING AND OPERATIONS TOOLS

TOOL	APPROVED USE	CONDITIONS
Sanctioned Custom GPT for billing (post-remediation)	Drafting appeals letters from de-identified case summaries	BAA-covered API key. Audit log of prompts retained. Use anonymized case references; no raw patient identifiers in prompts.
Excel AI formula generation	Formula and analysis assistance in finance workbooks	No PHI in workbook context shared with AI features.

§3.4 — Adding a new tool to this list requires CIO approval and security review per §5. Submit requests via the IT ticketing system; average review cycle is 8 business days.

§4 · PROHIBITED USE

What employees may not do with AI tools, and why each prohibition exists.

Employees may not enter Protected Health Information (PHI), Personally Identifiable Information (PII) beyond the employee's own, or any data classified as Confidential under §2 into any AI tool that does not appear on the Sanctioned Tools List (§3). This prohibition specifically includes, but is not limited to:

(a) Consumer accounts of OpenAI ChatGPT, Anthropic Claude, Google Gemini, Microsoft Copilot (non-E5 tier), or any other consumer LLM accessed via web or mobile, including free-tier and paid-personal-subscription accounts;

(b) Browser extensions that submit data to AI APIs not contracted by Northgate, including but not limited to Bardeen.ai, AnyText.ai, and any extension installed from the Chrome Web Store, Edge Add-ons store, or Firefox Add-ons without prior CIO approval;

(c) AI features embedded in tools whose Northgate-contracted version does not include those features (e.g., the AI summarization feature in personal Otter.ai accounts is prohibited; the AI scribe feature in Otter.ai Enterprise+ — post-BAA execution — is sanctioned per §3.1);

(d) AI tools accessed via personal devices when those tools are used for Northgate work, even if the employee believes the personal account provides a more convenient experience;

(e) Any AI tool that has been disabled or blocked at the network or endpoint level by IT, even if the employee discovers a means to bypass the block (e.g., personal hotspot, VPN to a personal network).

§4 · PROHIBITED USE (CONTINUED)

Specific data types that may never be entered, even into sanctioned tools.

§4.1 DATA TYPES PROHIBITED IN ANY AI TOOL

Even within sanctioned tools, the following data types may NEVER be entered as prompts or input:

- Patient names in combination with any clinical information, when the combination could re-identify the patient
- Medical record numbers (MRNs) in any context
- Social Security Numbers, driver's license numbers, or other government identifiers
- Full credit card numbers, bank account numbers, or other payment credentials
- Internal credentials including passwords, API keys, or authentication tokens
- Information about ongoing legal matters, regulatory investigations, or pending business transactions, except as explicitly authorized by the CEO

§4.2 CONSEQUENCES OF VIOLATION

Violation of this section may result in disciplinary action up to and including termination of employment, and may trigger mandatory reporting obligations under 45 CFR 164.408 (HIPAA breach notification) if PHI is found to have been transmitted to an unsanctioned AI tool. Consequences will be applied with reference to: (a) whether the violation was inadvertent or willful; (b) whether the violation was self-reported per §6; (c) the nature and quantity of data involved; and (d) the employee's history of prior violations.

MITIGATING FACTORS

"I didn't know" is not a defense, but it is a **mitigating factor** when the employee has completed the required training (§7) and the violation does not involve consumer tools explicitly named in §4 above. **Self-reporting under §6 is always a significant mitigating factor** and is explicitly encouraged.

§5 · DATA HANDLING STANDARDS

The four-tier data classification, vendor review checklist, and ongoing monitoring.

§5.1 DATA CLASSIFICATION

Northgate data is classified into four tiers for purposes of AI use:

- **Restricted** — PHI, authentication credentials. May only be entered into AI tools where a BAA is in place AND the specific workflow is described in §3.
- **Confidential** — employee PII, vendor contracts, strategic plans. May be entered into sanctioned AI tools where a DPA is in place.
- **Internal** — operational documents not intended for public distribution. May be entered into sanctioned AI tools at the employee's discretion.
- **Public** — marketing materials, patient education content, published policies. No AI restriction.

§5.2 VENDOR PROCUREMENT REVIEW

All proposed AI vendors must be reviewed via the IT Security Review checklist before any procurement contract is signed. The checklist requires verification of:

(a) Business Associate Agreement or equivalent for any vendor processing PHI;

(b) SOC 2 Type II or equivalent attestation;

(c) Documented data residency (must be US-based for PHI-processing vendors);

(d) Explicit prohibition on training the vendor's models on Northgate data;

(e) Breach notification SLA of 72 hours or less.

The checklist owner is the CIO; the average review cycle is 8 business days. No exceptions may be granted without written CEO approval.

§5.3 VENDOR CHANGE MONITORING

Sanctioned AI vendors that materially change their data handling practices (e.g., changing data residency, adding training-on-customer-data clauses, modifying breach notification SLA, or subprocessing to new vendors) must notify Northgate per the BAA/DPA terms. The CIO is responsible for evaluating such changes and may move a vendor from the Sanctioned Tools List pending re-review.

§6 · REPORTING AND ESCALATION

How to report suspected violations. The no-blame standard. Response timeline.

Employees who suspect they may have inadvertently transmitted PHI or other Confidential data to an unsanctioned AI tool must report the incident to the Security Officer within 24 hours of discovery. Reports may be made by email (**security@northgatefm.example**) or via the Incident Hotline (1-555-xxx-xxxx).

§6.1 NO-BLAME REPORTING

No disciplinary action will be taken solely for the act of reporting a suspected violation. This policy explicitly favors disclosure over concealment to enable timely breach assessment under HIPAA §164.402. Reports made in good faith do not constitute admission of intent and will not be used as the sole basis for disciplinary action.

Disciplinary action remains possible where (a) the underlying behavior was willful or grossly negligent, (b) the employee has a documented history of similar violations, or (c) the reporting was untimely (significantly more than 24 hours from discovery) without justification.

§6.2 WHAT TO REPORT

- Inadvertent pasting of PHI or Confidential data into an unsanctioned AI tool
- Discovery that an AI tool in use turns out not to be on the Sanctioned Tools List
- Suspicious behavior from a sanctioned AI tool (unexpected outputs, prompts for unusual permissions, performance changes)
- Communications from AI vendors that suggest a breach, security incident, or policy change
- Receipt of phishing attempts or social engineering specifically targeting AI tool credentials

§6.3 RESPONSE TIMELINE

On receipt of a report, the Security Officer will: **(a)** acknowledge within 4 business hours; **(b)** conduct initial assessment within 24 hours to determine whether the incident triggers HIPAA breach notification requirements; **(c)** escalate to the CIO and Compliance Officer as appropriate; and **(d)** communicate the outcome to the reporting employee within 5 business days, including any required follow-up actions.

§7 · TRAINING AND ACKNOWLEDGMENT

Four modules, annual cycle, what counts as completion.

All employees must complete the AI Acceptable Use Training program within 14 days of policy distribution (or 14 days of hire for new employees). The program consists of four modules totaling approximately 70 minutes.

§7.1 REQUIRED MODULES

01 Why this policy exists — 15 MIN · ALL EMPLOYEES · WEEKS 1–2

Frames the policy around HIPAA, the carrier rider, and the realities of consumer AI tools. Includes the carrier renewal context so employees understand the business stake. Quiz: 5 questions, 80% pass required.

02 The sanctioned tools and how to use them — 25 MIN · ROLE-SPECIFIC · WEEKS 3–6

Three role tracks (clinical / office / billing) with 3–4 worked examples each drawn from real workflows.

03 The line that triggers reporting — 20 MIN · ALL EMPLOYEES · WEEKS 7–10

Concrete examples: what counts as PHI in the AI context, what to do if you accidentally paste it, no-blame reporting path. Six scenario-based decision exercises with feedback.

04 Annual refresh and quiz — 10 MIN · ANNUAL · YEAR 2+

Annual re-acknowledgment cycle. 10-question scenario quiz; 80% pass required. Questions rotate annually to prevent gaming. Failure auto-routes to a 25-minute remedial module.

§7.2 ACKNOWLEDGMENT

Completion of all four modules (Modules 1–3 for new employees in the first year) results in an acknowledgment of this policy retained in the personnel file for the duration of employment plus six years per HIPAA records retention. Acknowledgment is captured electronically via TalentLMS. Employees who decline to acknowledge must meet with the Compliance Officer before continuing to use any AI tools in the course of work.

§8 · POLICY REVIEW AND REVISION

Annual review cadence, interim revision triggers, version history.

This policy is reviewed annually by the IT Steering Committee. The annual review occurs in the first quarter and is timed to feed into the cyber insurance renewal cycle.

§8.1 INTERIM REVISION TRIGGERS

Interim revisions may occur in response to:

- Material changes in the AI tool landscape (new vendor categories, major regulatory action against an AI vendor)
- Regulatory updates (HIPAA amendments, state law changes, EU AI Act implementation milestones)
- Carrier rider updates
- Lessons learned from incident reports
- Changes in Northgate's sanctioned tool roster

§8.2 REVISION HISTORY

VERSION	DATE	SUMMARY OF CHANGES	APPROVED BY
1.0	April 22, 2026	Initial policy issuance; supersedes "Employee Handbook §7.4 — Approved Software Use" insofar as it relates to AI tools.	CIO and IT Steering Committee

§8.3 QUESTIONS AND FEEDBACK

Questions about the application of this policy in specific situations should be directed to the Security Officer (security@northgatefm.example) or the CIO. Feedback on the policy itself, including suggestions for improvement, is welcomed at any time and will be considered during the next scheduled annual review.

ABOUT THIS DOCUMENT

This is a **sample AUP** distributed as part of the Shadow AI Labs AI Risk Sprint engagement bundle. The full policy that a real client receives is tailored to their specific tool inventory, regulatory context, carrier rider language, and operational realities. The structure, citations, and policy patterns shown here are representative of what we deliver.

APPENDIX · EMPLOYEE ACKNOWLEDGMENT FORM

Print, sign, return to Human Resources — or acknowledge electronically via TalentLMS.

This page may be printed and signed in lieu of electronic acknowledgment, where electronic acknowledgment is not feasible. Submit completed forms to Human Resources for inclusion in the employee's personnel file.

EMPLOYEE ACKNOWLEDGMENT

I, the undersigned employee of Northgate Family Medicine, acknowledge that I have:

1. Read this AI Acceptable Use Policy in its entirety;
2. Completed the required training modules (Modules 1–3 minimum);
3. Understand the prohibitions described in §4, including the prohibition on entering Protected Health Information into unsanctioned AI tools;
4. Understand my obligation to report suspected violations to the Security Officer within 24 hours of discovery per §6;
5. Will re-acknowledge this policy annually on the anniversary of this acknowledgment.

EMPLOYEE NAME (PRINTED)	EMPLOYEE ID
_____	_____
SIGNATURE	**DATE**
_____	_____

— End of Policy —