SHADOW AI LABS

CONFIDENTIAL · CLIENT EYES ONLY

AI RISK SPRINT · FIELD REPORT

Two weeks of AI risk discovery, classification, and remediation planning.

A productized engagement covering 118 employees, 47 AI tools in active use, and a cyber insurance renewal landing in 47 days.

PREPARED FOR

Northgate Family Medicine

ENGAGEMENT

SAL-2026-014

ENGAGEMENT WINDOW

April 7 – April 18, 2026

PREPARED BY

Peter Kwidzinski · shadowailabs.com

EXECUTIVE SUMMARY · FINDINGS AT A GLANCE

Forty-seven tools. Nine create material exposure. Five close inside the Sprint window.

47

AI tools, sanctioned and unsanctioned, in active use across 118 employees at four clinical locations.

04	05	38
Consumer tools transmitting PHI without a BAA. Immediate remediation required.	Enterprise tools with BAA gaps that the carrier will flag at renewal.	Tools eligible for sanctioned use once policy and DPA review is in place.

> *"The billing clerk's productivity hack put PHI on three consumer LLMs in fourteen months. She didn't know. Nobody asked. The carrier is going to ask in six weeks."*
>
> PETER KWIDZINSKI ENGAGEMENT OBSERVATION

Northgate's exposure is not unusual for a multi-specialty practice at this scale. Sanctioned tools (Microsoft Copilot E5, EHR-embedded scribe) have BAA coverage and aligned governance. Consumer tools (free-tier ChatGPT, three Chrome AI extensions, Otter.ai on intake calls) entered the workflow through individual employee initiative and have neither contract coverage nor policy authorization. The cyber insurance AI Security Rider that Northgate's carrier (Travelers) is introducing at the May 24 renewal questionnaire will require documented coverage of all ten control areas — five of which are currently unmet.

EXECUTIVE SUMMARY · ACTION THIS WEEK

Three actions should be in motion before this report leaves the leadership table.

Two are remediation moves you can take internally. One requires vendor coordination and should start today. None require new spend beyond effort already authorized in your existing Microsoft 365 E5 agreement.

01 **Block consumer ChatGPT and three Chrome AI extensions at the network egress layer.**

Cisco Umbrella DLP rules already exist as configured-but-not-enforcing categories. Enabling enforcement is a 20-minute change in the admin console. Replacement: Microsoft Copilot E5 (already licensed under your M365 agreement, BAA on file) covers ~80% of the use cases currently driving employees to consumer ChatGPT.

45 CFR 164.308(a)(1)(ii)(D) · NIST AI RMF MS-2.7 · carrier rider §4.7

02 **Execute the Otter.ai Enterprise+ BAA repaper.**

Otter.ai offers BAA coverage on Enterprise+ tier; your current Enterprise contract does not include it. Upgrade pricing differential is ~$140/user/year. Seven providers currently using Otter on clinical and intake calls would be covered. Otter's legal will turn the BAA in 5-7 business days; start now to land before May 24.

45 CFR 164.504(e) · HIPAA Business Associate Rule

03 **Distribute the AUP (Deliverable 04) with formal acknowledgment.**

Carrier rider §4.1 requires distributed, acknowledged AUP. Your existing employee handbook references "approved tools only" without enumerating AI. A complete AUP draft is included as Deliverable 04 of this report. Distribution via your LMS with required acknowledgment satisfies the rider language and creates an evidence trail.

Carrier rider §4.1 · NIST AI RMF GV-1.3

IF YOU DO NOTHING ELSE THIS WEEK

Action **01** alone reduces your highest-severity exposure by approximately 80%. The technical change is minutes; the policy change (distributing the new AUP) is a leadership communication, not a procurement decision. Both close before any external party sees the renewal questionnaire.

DELIVERABLE 01 · AI TOOL DISCOVERY · METHODOLOGY

How we surfaced every AI tool in active use across your organization.

Discovery combines three independent data sources. The triangulation matters: any single method misses something, but the overlap surfaces the tools your IT inventory doesn't know about and your employees don't volunteer.

118	14 days	47
EMPLOYEES IN SCOPE	BROWSER TELEMETRY WINDOW	DISTINCT AI TOOLS SURFACED

Method 1 · Browser telemetry analysis

Fourteen days of egress traffic from your Cisco Umbrella deployment, filtered for known AI service domains and emerging-LLM endpoints. Surfaces unsanctioned use that doesn't appear in any procurement record. Captured ChatGPT consumer accounts, three Chrome extensions invoking AI APIs, and two browser-embedded summarizers nobody on the IT team knew existed.

Method 2 · Employee survey

Anonymous survey distributed to all 118 employees with 89% completion rate. Surfaced tool use that's intentionally outside the formal procurement process — "productivity hacks" employees adopt because the official path is slow. Eight tools came from this method that did not appear in telemetry, because they're used only on personal devices or via mobile apps.

Method 3 · Procurement and SSO log audit

Twelve months of procurement records and 90 days of Okta SSO logs cross-referenced against AI vendor SKUs. Surfaces the sanctioned-but-undocumented use: tools your finance team paid for but your security team never reviewed. Examples include the Otter.ai Enterprise contract (no BAA on file), the EHR-embedded AI scribe (acquired through a vendor upgrade without separate governance review), and a 6-seat Gamma account that finance approved as "presentations software."

Triangulation summary

Twenty-six tools appeared in at least two methods. Twenty-one appeared in only one — these are the tools an inventory built on any single source would have missed. The full enumeration is in **Appendix A**; the risk-relevant nine tools are detailed in Deliverables 02 and 03.

DELIVERABLE 01 · AI TOOL DISCOVERY · SAMPLE DATA

What the discovery data actually looked like, with PII removed.

The raw discovery dataset is held in encrypted client storage. Below is a representative subset of egress telemetry, employee survey signals, and procurement record cross-referencing that drove the findings on the previous page.

Browser telemetry sample · top 8 AI domains observed, 14-day window

DOMAIN	CATEGORY	HITS	USERS
microsoft.com/copilot	Sanctioned LLM	14,328	87
chat.openai.com	**Consumer LLM**	**2,847**	**12**
otter.ai	Transcription AI	1,203	9
notion.so/ai	Sanctioned AI	891	14
claude.ai	**Consumer LLM**	**414**	**3**
grammarly.com/api	Sanctioned AI	312	22
bardeen.ai	**Chrome ext · AI**	**287**	**2**
gamma.app	Sanctioned AI	156	6

Employee survey · representative free-text responses

Q11 · "Do you use AI tools at work? If yes, which ones?"

"I use ChatGPT for patient education materials — much faster than writing from scratch." — Clinical, anonymized

"Otter.ai records all my intake calls and sends a summary. I didn't know I needed approval." — Front desk, anonymized

"Built a custom GPT to help write appeals letters — saves me 6 hours a week." — Billing manager, anonymized

Procurement cross-reference · vendor contracts containing AI clauses

Audit of 12 months of procurement (143 vendor agreements) found 9 vendor contracts with AI-feature clauses. Three of nine had no corresponding BAA or DPA review on file. Two were procured outside the IT security review path (Gamma via marketing budget code; Otter.ai Enterprise renewed by office manager). One — the EHR vendor's AI scribe upgrade — was bundled into a renewal and never broken out as a separate security review item.

DELIVERABLE 02 · RISK CLASSIFICATION MATRIX

Each tool mapped against business criticality and risk severity, NIST AI RMF–aligned.

Risk severity uses the NIST AI Risk Management Framework measure-and-manage taxonomy. Business criticality is your organization's own operating definition, validated in the Day 1 kickoff. The combination tells leadership where to spend remediation budget first.

	NEGLIGIBLE	LOW	MEDIUM	HIGH	CRITICAL
CRITICAL		EHR-embedded AI scribe	MS Copilot E5	Otter.ai clinical	**ChatGPT · PHI** **3× Chrome AI**
HIGH	Grammarly free	Notion AI · no PHI	Claude · Gemini	Custom GPT billing Otter intake	**ChatGPT billing**
MEDIUM	Calendly AI	Loom · Zoom AI	Copilot · Gamma	Fathom sales	
LOW	Browser autocomplete	Canva · Descript	Grammarly Business		
TRIVIAL	Reply suggestions				

↑ Business criticality · Risk severity →

TOP 3 CELLS BY PRIORITY

- Critical × Critical (top right): four consumer tools transmitting PHI without a BAA. Section 03 details carrier consequences.
- High × High (band): Custom-built GPT used by billing department and Otter.ai on intake calls. Both have remediation paths; neither is solved by policy alone.
- Critical × Low (left of critical column): EHR-embedded AI scribe is in this cell because vendor has BAA, configuration is locked, output is reviewed. This is what "good" looks like for a critical workflow — keep it.

DELIVERABLE 02 · PER-TOOL SCORING RUBRIC

The methodology behind every cell on the previous page.

Risk severity for each tool is scored against five questions derived from the NIST AI RMF Measure function. A tool's score is the highest single-question result, not an average — one critical exposure outweighs four well-managed ones.

CRITICAL	**PHI / PII transmission without contract.** Tool receives Protected Health Information, Personally Identifiable Information, or other Confidential data, AND no Business Associate Agreement or Data Processing Agreement is in place with the vendor. NIST AI RMF MS-2.7 (data privacy). Triggers immediate remediation. *Example:* consumer ChatGPT used for patient communication drafts.
HIGH	**Contractual gap with material data exposure.** BAA or DPA exists but does not cover the specific AI workflow in use; OR vendor has BAA but tier in use is not BAA-covered. NIST AI RMF MS-2.6 (data governance). Remediation typically: contract amendment or tier upgrade. *Example:* Otter.ai Enterprise (no BAA on this tier) used for clinical recording.
MEDIUM	**Sanctioned tool with policy or training gap.** Vendor relationship is properly contracted but internal use is not policy-governed, training is missing, or audit trail is incomplete. NIST AI RMF GV-1.3 (governance). Remediation typically: policy update and training rollout, no vendor action. *Example:* Microsoft Copilot E5 with no employee training on PHI-handling boundaries.
LOW	**Sanctioned tool with full governance.** Vendor contracted, BAA/DPA in place, internal policy enumerates approved use, employees trained, audit trail retained. NIST AI RMF aligned across all five functions. *Example:* EHR-embedded AI scribe with locked configuration, vendor BAA, and clinician-review-before-sign-off workflow.

Scoring application example · Custom GPT used by billing

The billing manager's custom GPT scored **High** on questions 2 (no DPA for the custom API key in use) and **High** on question 4 (no audit log of prompts containing patient data). It scored **Low** on question 1 (no PHI in observable use) and **Low** on question 5 (output is human-reviewed before any external action). Final tool score: **High** — driven by the contractual gap on the API key, not by direct PHI exposure. Remediation: BAA-covered API key, prompt logging, documented owner. See Deliverable 06 for the full sanction recommendation.

DELIVERABLE 03 · CYBER INSURANCE RIDER GAP ANALYSIS

Five of ten carrier requirements are gaps. Three close as Sprint deliverables. Two need execution.

Carrier: Travelers Cyber · 2026 AI Security Rider (Form CY-AISR-26-001). Rider language reviewed in full: 14 underwriting questions, 10 documented-control requirements. Renewal questionnaire arrives 47 days from today; submission deadline May 24.

#	CARRIER REQUIRES	CURRENT STATE AT NORTHGATE	STATUS
01	**Written AI Acceptable Use Policy** §4.1 · Distributed, acknowledged	Employee handbook mentions "approved tools only" without enumerating AI. Not distributed as AI-specific policy. Deliverable 04 closes.	GAP
02	**Inventory of AI tools in use** §4.2 · Annual review	No inventory existed pre-Sprint. Deliverable 01 surfaces 47 tools. Closes on report acceptance.	GAP
03	**Risk classification methodology** §4.3 · NIST/ISO equivalent	No methodology in place. Deliverable 02 closes with NIST AI RMF mapping.	GAP
04	**Annual AI risk assessment** §4.4 · Documented, retained	This Sprint satisfies 2026. Recurrence required for 2027 renewal — flag for calendar.	PARTIAL
05	**Vendor security review for AI vendors** §4.5 · DPAs/BAAs on file	Otter.ai BAA missing (upgrade required). ChatGPT consumer not contracted. Microsoft DPA covers majority.	GAP
06	**AI-specific training records** §4.6 · Annual, 3yr retention	Existing HIPAA training references AI use in clinical workflow. LMS retains records. Adequate.	MET
07	**DLP coverage of AI domains** §4.7 · Block or monitor egress	Cisco Umbrella has 6 AI domains in default category but enforcement is off. Configuration change closes (Week 1).	PARTIAL
08	**Incident response runbook · AI scenarios** §4.8 · Tested annually	Existing IR runbook (HIPAA scope) does not include AI scenarios. AI runbook delivered Day 12.	GAP
09	**Governance with AI oversight authority** §4.9 · Documented stop-authority	CIO has documented stop-authority via existing IT Steering Committee charter. Carrier will accept.	MET
10	**BAA coverage for AI processing of PHI** §4.10 · HIPAA cross-reference	Microsoft Copilot E5 BAA covers approved use. Others sanctioned/blocked per Deliverable 05.	MET

UNDERWRITER-READY CONCLUSION

Three of five gaps close as Sprint report deliverables (documents). Two require execution: **Otter.ai BAA repaper** (vendor lead time 5–7 business days) and **DLP enforcement enablement** (internal IT, ~20 minutes). Total remediation effort to underwriter-ready: **approximately 4 weeks**. See **Appendix B** for the full sub-control mapping behind each row.

DELIVERABLE 03 · CARRIER RESPONSE LANGUAGE DRAFTS

The actual language you write back to the underwriter on the four hardest rider questions.

Underwriters score answers, not assessments. These drafts answer the four rider questions most likely to be scored negatively without specific language. Use as-is or adapt with your broker; both your Sprint report and these drafts attach as exhibits.

RIDER Q5 · "HAS THE APPLICANT CONDUCTED A DOCUMENTED AI RISK ASSESSMENT IN THE LAST 12 MONTHS?"

Yes. Northgate completed a structured AI risk assessment in April 2026, conducted by Shadow AI Labs under engagement SAL-2026-014. Methodology aligns with NIST AI Risk Management Framework (NIST AI 100-1) and includes tool discovery (47 tools enumerated), risk classification (NIST AI RMF–aligned matrix), and remediation roadmap. Full assessment report attached as Exhibit A. Next assessment scheduled for Q1 2027.

RIDER Q7 · "DOES THE APPLICANT MAINTAIN A WRITTEN, DISTRIBUTED AI ACCEPTABLE USE POLICY?"

Yes. Northgate's AI Acceptable Use Policy was distributed via the corporate LMS (TalentLMS) on April 22, 2026, with required employee acknowledgment. Current acknowledgment rate: 96% (114/118 employees, with remaining 4 on leave and scheduled for return-from-leave training). Policy attached as Exhibit B. Annual review and re-acknowledgment scheduled for April 2027.

RIDER Q9 · "DESCRIBE THE APPLICANT'S PROCESS FOR EVALUATING NEW AI VENDORS BEFORE PROCUREMENT."

All proposed AI vendors are routed through the IT Security Review checklist (revised April 2026 to add AI-specific items) before any procurement contract is signed. The checklist requires: (a) Business Associate Agreement or equivalent for any vendor processing PHI; (b) SOC 2 Type II or equivalent attestation; (c) documented data residency and retention; (d) explicit prohibition on training the vendor's models on Northgate data. Checklist attached as Exhibit C. Owner: CIO. Average review cycle: 8 business days.

RIDER Q12 · "WHAT CONTROLS PREVENT EMPLOYEES FROM USING CONSUMER AI TOOLS WITH SENSITIVE DATA?"

Layered controls: (a) Cisco Umbrella DLP enforcement blocks egress to 18 enumerated consumer AI domains (effective April 18, 2026); (b) Written AUP (Exhibit B) prohibits consumer AI use with PHI, with disciplinary consequences enumerated; (c) Employee training (annual, 80% pass required) covers PHI-handling boundaries in AI tools; (d) Quarterly governance committee review of new AI tool requests. Sanctioned alternatives (Microsoft Copilot E5, Otter.ai Enterprise+ with BAA) provided to reduce shadow-AI motivation.

DELIVERABLE 04 · ACCEPTABLE USE POLICY

An AUP drafted to your operational reality. Not a template.

The full AUP is provided as a separate document (DOCX, editable, 14 pages). What follows is the policy structure and an excerpt from §3 Prohibited Use, the section that generates the most policy-implementation friction.

Policy structure

- §1 · Scope and definitions — what counts as AI for purposes of this policy
- §2 · Sanctioned tools list — what employees may use, with conditions per tool
- §3 · Prohibited use — consumer LLMs with PHI, unverified Chrome extensions, etc.
- §4 · Data handling — what data may go into sanctioned AI tools, with examples
- §5 · Vendor procurement — how new AI tools are evaluated before adoption
- §6 · Reporting and escalation — how employees flag AI security concerns
- §7 · Acknowledgment and review — annual employee re-acknowledgment cycle

EXCERPT · §3 PROHIBITED USE (REPRESENTATIVE)

Employees may not enter Protected Health Information (PHI), Personally Identifiable Information (PII) beyond the employee's own, or any data classified as Confidential under §2.4 into any AI tool that does not appear on the Sanctioned Tools List (§2). 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 Family Medicine;

(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+ is sanctioned per §2.3).

Violation of this section may result in disciplinary action up to and including termination, and may trigger mandatory reporting obligations under 45 CFR 164.408 (HIPAA breach notification) if PHI is found to have been transmitted.

EXCERPT · §6 REPORTING AND ESCALATION (REPRESENTATIVE)

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. **No disciplinary action will be taken solely for the act of reporting**; the policy explicitly favors disclosure over concealment to enable timely breach assessment under HIPAA §164.402.

DELIVERABLE 04 · TRAINING PLAN

Four modules, twelve weeks, LMS-ready. Audience-segmented.

The training plan is built for distribution via your existing TalentLMS. Each module includes scripted narration, slide assets, and a scenario-based quiz. Module completion satisfies carrier rider §4.6 (AI-specific training records, 3-year retention).

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. Not "we caught you" — "here's what changed and why." Includes the carrier renewal context so employees understand the business stake. **Quiz:** 5 questions, 80% pass required for acknowledgment.

02 **The sanctioned tools and how to use them.** 25 min · role-specific · weeks 3-6

Three role tracks. **Clinical track:** EHR-embedded scribe, Otter.ai Enterprise+ (post-BAA), Microsoft Copilot E5 for clinical documentation. **Office staff track:** Microsoft Copilot E5 for admin work, Grammarly Business, Calendly AI. **Billing track:** sanctioned custom GPT (post-remediation), Excel AI features, sanctioned vendor list. Each track includes 3-4 worked examples 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 (patient names, MRNs, diagnosis codes, even paraphrased clinical scenarios that could be re-identified), what to do if you accidentally paste it (don't panic, don't conceal, report within 24 hours per §6), no-blame reporting path. Includes 6 scenario-based decision exercises with feedback.

04 **Refresh and quiz.** 10 min · annual · all employees

Annual re-acknowledgment cycle. Scenario-based quiz with 10 questions; 80% pass required for acknowledgment to count toward carrier requirement §4.6. Questions rotate annually to prevent gaming. Failure auto-routes to a 25-minute remedial module.

DISTRIBUTION CHECKLIST

(1) Upload all four modules to TalentLMS, (2) assign Module 1 to all employees with 14-day completion deadline, (3) configure role-track auto-assignment for Module 2 based on Active Directory groups, (4) configure annual re-acknowledgment trigger for Module 4. Carrier evidence package: LMS completion report + AUP acknowledgment receipts.

DELIVERABLE 05 · 90-DAY REMEDIATION ROADMAP

Every gap sequenced into a Week 1 / 30 / 60 / 90 plan. Owners, dependencies, effort estimates.

The plan is sequenced so that the carrier renewal questionnaire (May 24) lands inside Week 7 — every item required for that questionnaire is complete by Week 6. Items past Week 6 strengthen the underwriter response but are not strictly required for renewal.

WEEK 1

Block & deploy

- Enable Cisco Umbrella DLP on AI categories
  IT · 1 DAY
- Distribute AUP with required acknowledgment
  HR · 3 DAYS
- Disable ChatGPT consumer access for 4 flagged users
  IT · 1 HOUR
- Initiate Otter.ai Enterprise+ BAA repaper
  PROCUREMENT · 1 DAY

DAY 30

Sanction & train

- Microsoft Copilot E5 rollout to billing and office staff
  IT · 2 WEEKS
- Module 1 + Module 2 training delivered
  COMPLIANCE · 2 WEEKS
- Otter.ai Enterprise+ BAA executed and seats migrated
  PROCUREMENT · 2 WEEKS
- Custom GPT for billing — security review & sanction or sunset
  CIO · 1 WEEK

DAY 60

Audit & respond

- Carrier renewal questionnaire response — submit
  CIO · 3 DAYS
- Module 3 training delivered
  COMPLIANCE · 1 WEEK
- Sanctioned tool inventory published internally
  IT · 2 DAYS
- First quarterly governance committee meeting
  CIO · 2 HOURS

DAY 90

Operationalize

- Procurement workflow updated for AI vendor review
  PROCUREMENT · 2 WEEKS
- IR runbook AI scenarios tested in tabletop exercise
  IT/CIO · 1 DAY
- Module 4 refresh quiz scheduled in LMS
  COMPLIANCE · 1 DAY
- 2027 renewal calendar set with 90-day pre-flight
  CIO · 1 HOUR

IF YOUR TEAM EXECUTES

Week 1 items alone close **three of five gaps** and reduce highest-severity exposure by ~80%. Week 1 is achievable with existing internal resources and no new spend. Day 30–90 items can be executed internally OR through the AI Governance Implementation engagement — Sprint fee credits toward that engagement if you decide to proceed.

DELIVERABLE 05 · PER-ACTION ACCEPTANCE CRITERIA

How you know each Week 1 action is actually complete, not just attempted.

Each Week 1 action has a verifiable completion test. The test is what an auditor or underwriter would accept as evidence. "We worked on it" is not a completion test; "Configuration X is set to value Y and screenshot Z exists" is.

WEEK 1 · 01 **Enable Cisco Umbrella DLP on AI categories.** Complete when: (a) Umbrella admin console shows 18 AI-category domains in "Block" or "Monitor" state; (b) screenshot dated within 7 days of activation retained in compliance vault; (c) test query from an internal device to `chat.openai.com` is blocked and logged; (d) blocked-traffic log retained for 90 days.

WEEK 1 · 02 **Distribute AUP with required acknowledgment.** Complete when: (a) TalentLMS shows AUP module assigned to 118 employees; (b) acknowledgment rate ≥ 90% within 14 days; (c) remaining 10% (on leave, etc.) have documented return-from-leave training plan; (d) signed acknowledgment receipts exported and retained in HRIS.

WEEK 1 · 03 **Disable ChatGPT consumer access for 4 flagged users.** Complete when: (a) the 4 named users' Okta access logs show no successful sessions to `chat.openai.com` after the disable date; (b) each user received the standard "consumer-AI-disabled" notification email; (c) replacement (Microsoft Copilot E5) access confirmed via Okta entitlement audit.

WEEK 1 · 04 **Initiate Otter.ai Enterprise+ BAA repaper.** Complete when: (a) procurement has emailed Otter.ai legal requesting Enterprise+ upgrade with BAA; (b) Otter.ai has acknowledged receipt and provided estimated BAA execution date; (c) procurement has communicated estimated migration date to clinical leadership.

WHY ACCEPTANCE CRITERIA MATTER

Carrier underwriters and HIPAA auditors don't grade on effort. They grade on verifiable evidence. Acceptance criteria turn the roadmap from a to-do list into an evidence-production plan. **Side benefit:** if your team executes against these criteria and retains the artifacts, the next assessment in 2027 is materially cheaper because the evidence trail already exists.

DELIVERABLE 06 · EXECUTIVE READOUT · DECISIONS REQUIRED

Five things leadership should leave the 60-minute readout with.

The Day-14 executive readout is structured as a decision session, not a presentation. Each item below is a decision leadership owns — Shadow AI Labs presents the analysis and recommendation, leadership decides. Decisions captured in the meeting minutes become the basis for Implementation engagement scope (if elected).

01 **Microsoft Copilot E5 rollout sequencing.**

The license already exists. The question is whether to roll out to all eligible users in Week 2 or stage by department. Staging gives training time per cohort; bulk rollout closes the consumer-ChatGPT replacement gap faster. **Strategist recommends staging by department** starting with billing, since billing is the largest contributor to consumer-ChatGPT exposure.

02 **Otter.ai upgrade authority.**

Enterprise+ tier upgrade is ~$140/seat/year incremental. Seven seats currently use Otter on clinical/intake calls. Annual cost difference: ~$980. Material? No. Worth executing this week? Yes — vendor lead time for BAA execution is 5–7 business days, and the carrier questionnaire arrives in week 7. **Strategist recommends proceeding immediately.**

03 **Custom GPT for billing.**

The custom GPT built by your billing manager is the most interesting finding — useful work, used daily, built outside the formal sanction process. **Strategist recommends formal sanction** (security review of prompt/data flow, BAA-covered API key, documented owner) rather than sunset. Cost: ~6 hours of CIO time. Benefit: preserves a productivity gain employees value, prevents the "shadow rebuild" if you sunset without offering a sanctioned replacement.

04 **Carrier conversation positioning.**

When the carrier questionnaire arrives, leadership should be prepared to attach this report as the §4.4 documentation. Whether to also attach the AUP, Inventory, Matrix, and Appendix B as separate exhibits is tactical — your broker can guide. **Strategist default: attach all of them.** The cost is low and the underwriter signal is high.

05 **AI Governance Implementation engagement.**

The Sprint identifies the work. The Implementation engagement executes Week 30–90 of the roadmap with Shadow AI Labs leading vendor coordination, BAA repaper, DLP rollout, IR runbook update, and governance committee setup. Indicative scope: 8 weeks, $24K. Sprint fee credits toward it. Decision can wait through Week 4; Implementation begins Week 5 if elected.

DELIVERABLE 06 · PRE-READOUT PREPARATION & FOLLOW-UP TRACKING

What to read before the meeting. What to track after.

The 60-minute readout is most effective when leadership arrives with positions formed on the five decisions on the previous page. The pre-readout reading list below is structured so each attendee can prepare in 20-25 minutes.

Pre-readout reading list · by attendee role

CEO / Practice Administrator

Pages 2-3 (Executive Summary), Page 14 (Decisions Required), Page 22 (Engagement options). ~15 min.

CIO / IT Director

Full report. Particular attention to Pages 6-9 (Matrix + Gap Analysis), Pages 12-13 (Roadmap + Acceptance), Appendix A. ~45 min.

CFO / Finance

Pages 2-3 (Executive Summary), Page 14 decisions 02 + 05 (Otter cost, Implementation cost), Page 22 (pricing). ~10 min.

Compliance / HR

Pages 10-11 (AUP + Training), Page 13 (Week-1 acceptance), Page 9 (Carrier response drafts). ~20 min.

Follow-up tracking · what gets recorded post-meeting

DECISION LOG TEMPLATE (PROVIDED AS DOCX)

Each of the five decisions captured with: (a) decision text (e.g., "Approve staged Copilot E5 rollout starting with billing"); (b) decision owner (named); (c) target completion date; (d) any conditions or open questions; (e) escalation path if condition not met by date. Decision log circulates within 24 hours of the meeting.

30-day check-in

A 30-minute follow-up call is included in the Sprint engagement (no additional cost). Scheduled for Day 30 ± 3 days. Agenda: status against the five decisions, Week 1 completion check, any blockers, decision on the Implementation engagement. If Implementation has been elected, the 30-day call becomes the Implementation kickoff.

WHAT THIS IS NOT

This Sprint engagement does **not** include ongoing monthly check-ins, ad-hoc questions after Day 30, vendor coordination on your behalf, or remediation execution. Those are deliberately bundled into the AI Governance Implementation engagement so the Sprint stays productized at fixed scope and fixed cost.

METHODOLOGY · HOW THIS ENGAGEMENT RAN

Two weeks. Fixed scope. Six deliverables. One executive readout.

Every AI Risk Sprint runs the same scope on the same calendar. No discovery calls to "figure out what you need." No hourly billing. No surprise change orders. The structure is the productization — clients buy a known quantity at a known price, and we deliver against a known checklist.

DAYS 1–4

Discovery

Day 1 · kickoff call, scope confirmation, access provisioning
90 MIN

Days 1–4 · browser telemetry pull, employee survey, procurement audit
PARALLEL

DAYS 3–9

Classification

Days 3–6 · risk classification matrix, NIST AI RMF mapping
SEQUENTIAL

Days 5–9 · cyber insurance rider gap analysis (carrier policy review)
SEQUENTIAL

DAYS 6–13

Synthesis

Days 6–10 · AUP draft and training plan
DRAFTING

Days 9–13 · 90-day remediation roadmap with sequenced owners
DRAFTING

DAY 14

Delivery

60-minute executive readout (Zoom, leadership only)
60 MIN

Final PDF + DOCX deliverables emailed same day
SAME DAY

Data handling and confidentiality

All client data accessed during the Sprint is processed under the SAL standard engagement MSA, which includes a Mutual NDA and a Business Associate Agreement for HIPAA-covered clients. Browser telemetry, survey responses, and procurement records are accessed via read-only credentials and are not exfiltrated to SAL infrastructure beyond what is required to produce this report. All working files are retained for 12 months in encrypted storage (SAL Vault, AES-256, US-region) and destroyed after retention.

Tools and references

This Sprint used: Cisco Umbrella egress data, anonymized internal survey via SurveyMonkey, Okta SSO log exports, Microsoft 365 admin console reports, Northgate's procurement system (read-only). Framework references: NIST AI Risk Management Framework (NIST AI 100-1, January 2023, plus Generative AI Profile NIST AI 600-1, July 2024), ISO/IEC 42001:2023, HIPAA Security Rule (45 CFR Part 164 Subpart C), and the Travelers Cyber Insurance AI Security Rider (Form CY-AISR-26-001, effective January 2026).

APPENDIX A · FULL AI TOOL INVENTORY (1 OF 3)

All 47 AI tools surfaced, with classification, business unit, and contract status.

#	TOOL	CATEGORY	BU	RISK	SANC.	BAA
01	**ChatGPT (consumer/Plus)**	Consumer LLM	Billing (3)	CRIT	NO	NO
02	**Claude.ai (consumer)**	Consumer LLM	Physician (1)	CRIT	NO	NO
03	**Bardeen.ai (Chrome ext)**	Browser AI	Admin (2)	CRIT	NO	NO
04	**AnyText.ai (Chrome ext)**	Browser AI	Admin (1)	CRIT	NO	NO
05	**Otter.ai Enterprise**	Transcription AI	Physician (7)	HIGH	YES	NO
06	**Custom GPT (billing)**	Custom LLM	Billing (1)	HIGH	NO	NO
07	**Gamma.app**	AI presentations	Marketing (6)	HIGH	YES	N/A
08	**Notion AI**	Workspace AI	Admin (12)	HIGH	YES	NO
09	**Loom AI**	Video AI	Leadership (4)	HIGH	YES	N/A
10	**EHR-embedded AI scribe**	Clinical AI	Physician (18)	LOW	YES	YES
11	**Dragon Medical One**	Clinical AI	Physician (14)	LOW	YES	YES
12	**eClinicalWorks AI features**	Clinical AI	Admin (22)	LOW	YES	YES
13	**NextGen Mobile AI**	Clinical AI	Physician (8)	LOW	YES	YES
14	**Microsoft Copilot E5**	Sanctioned LLM	All eligible	MED	YES	YES
15	**Outlook AI features**	Office AI	All (118)	LOW	YES	YES
16	**Teams AI meeting summaries**	Office AI	Leadership (12)	LOW	YES	YES

Continued on page 18. Risk: per Deliverable 02 rubric. Sanc: sanctioned per AUP §2. BAA: Business Associate Agreement on file (N/A if vendor processes no PHI).

APPENDIX A · FULL AI TOOL INVENTORY (2 OF 3)

Tools 17 through 32 · Office stack and employee productivity.

#	TOOL	CATEGORY	BU	RISK	SANC.	BAA
17	**Word AI rewrite**	Office AI	All (118)	LOW	YES	YES
18	**Excel AI formula gen**	Office AI	Finance (8)	LOW	YES	YES
19	**PowerPoint Designer AI**	Office AI	Marketing (6)	LOW	YES	YES
20	**Microsoft Loop AI**	Office AI	Admin (14)	LOW	YES	YES
21	**OneNote AI**	Office AI	Physician (9)	LOW	YES	YES
22	**Grammarly Business**	Writing AI	Admin (24)	LOW	YES	YES
23	**Calendly AI scheduling**	Scheduling AI	Admin (16)	LOW	YES	N/A
24	**SaneBox AI email triage**	Email AI	Physician (6)	MED	YES	YES
25	**Superhuman AI**	Email AI	Leadership (3)	MED	YES	YES
26	**Reclaim AI**	Scheduling AI	Physician (4)	LOW	YES	N/A
27	**Motion AI**	Task AI	Admin (5)	LOW	YES	N/A
28	**Krisp AI noise cancel**	Audio AI	Physician (11)	LOW	YES	N/A
29	**Zoom AI Companion**	Meeting AI	All (118)	MED	YES	YES
30	**HubSpot AI**	Marketing AI	Marketing (3)	LOW	YES	N/A
31	**Mailchimp AI subject lines**	Marketing AI	Marketing (2)	LOW	YES	N/A
32	**Canva AI Design**	Design AI	Marketing (5)	LOW	YES	N/A

Continued on page 19.

APPENDIX A · FULL AI TOOL INVENTORY (3 OF 3)

Tools 33 through 47 · Communication, development, and niche tools.

#	TOOL	CATEGORY	BU	RISK	SANC.	BAA
33	**Descript AI editing**	Media AI	Marketing (1)	LOW	YES	N/A
34	**ChatGPT Team (sanctioned)**	Sanctioned LLM	Leadership (8)	MED	YES	YES
35	**Perplexity Pro**	Search AI	Leadership (3)	LOW	YES	N/A
36	**Slack AI**	Comms AI	Admin (18)	LOW	YES	YES
37	**Fathom Sales**	Sales AI	Sales (2)	HIGH	YES	N/A
38	**Otter.ai Personal (leadership)**	Transcription AI	Leadership (2)	LOW	YES	N/A
39	**Read.ai**	Meeting AI	Leadership (3)	LOW	YES	N/A
40	**Tidio AI chat (website)**	Customer AI	Front desk (3)	LOW	YES	N/A
41	**Notion AI (engineering)**	Workspace AI	Engineering (3)	LOW	YES	N/A
42	**GitHub Copilot**	Code AI	Engineering (2)	LOW	YES	N/A
43	**Cursor AI**	Code AI	Engineering (1)	LOW	YES	N/A
44	**Postman AI**	Dev AI	Engineering (2)	LOW	YES	N/A
45	**Adobe Firefly**	Design AI	Marketing (2)	LOW	YES	N/A
46	**Beautiful.ai**	AI presentations	Marketing (1)	LOW	YES	N/A
47	**Tome AI**	AI presentations	Marketing (1)	LOW	YES	N/A

INVENTORY SUMMARY

47 total tools. **4 critical exposures** (consumer LLMs and Chrome extensions with PHI access, no contract). **5 high-severity contractual gaps** (vendors with BAA-required workflows but no BAA or wrong tier). **38 sanctioned-eligible** (either already governed or trivially closeable with policy enumeration). Inventory is the basis for AUP §2 sanctioned tools list and §3 prohibited use enumeration.

APPENDIX B · CARRIER RIDER FULL MAPPING (1 OF 2)

Each requirement broken into sub-controls with NIST AI RMF cross-reference and evidence.

Requirements 01–05 below. Each card expands the carrier's high-level requirement into the underlying sub-controls an underwriter expects to see evidence for, mapped to the corresponding NIST AI RMF control families.

§4.1 · Written AI Acceptable Use Policy

NIST AI RMF GV-1.1, GV-1.3

Sub-controls: (a) written policy exists · (b) distributed to all employees · (c) acknowledgment captured · (d) annual review cycle defined · (e) version-controlled

Evidence at Northgate: Pre-Sprint, only employee handbook reference to "approved tools" — no AI-specific policy. **Remediation:** Deliverable 04 provides 14-page AUP DOCX. Week-1 action: distribute via TalentLMS with required acknowledgment.

§4.2 · Inventory of AI tools in use

NIST AI RMF MP-3.1, MS-1.1

Sub-controls: (a) enumeration of all AI tools · (b) classification by data sensitivity · (c) business owner per tool · (d) annual review · (e) procurement gate for new tools

Evidence at Northgate: No inventory existed. **Remediation:** Appendix A (47 tools enumerated with classification, BU, sanction status, BAA status). §5 of AUP adds procurement gate.

§4.3 · Risk classification methodology

NIST AI RMF MS-2.1 through MS-2.7

Sub-controls: (a) documented methodology · (b) NIST or ISO equivalent · (c) applied to all inventory tools · (d) scoring rationale retained · (e) review on tool changes

Evidence at Northgate: No methodology in place. **Remediation:** Deliverables 02a (matrix) and 02b (per-tool scoring rubric) provide NIST AI RMF–aligned methodology. Scoring rationale documented per tool in working files.

§4.4 · Annual AI risk assessment

NIST AI RMF MS-1.1, MS-1.3

Sub-controls: (a) assessment performed · (b) documented · (c) leadership reviewed · (d) retained for 3 years · (e) cadence ≥ annual

Evidence at Northgate: This Sprint engagement satisfies 2026. **Open item:** 2027 recurrence — calendar entry recommended for Q1 2027. Sprint engagement retained in client compliance vault for 3 years.

§4.5 · Vendor security review for AI vendors

NIST AI RMF GV-6.1, MS-3.1

Sub-controls: (a) DPA or BAA on file · (b) SOC 2 or equivalent · (c) data residency documented · (d) model-training opt-out · (e) breach notification SLA

Evidence at Northgate: Microsoft DPA covers majority. **Gaps:** Otter.ai BAA missing (Enterprise+ upgrade required). ChatGPT consumer not contracted (blocked Week 1). 3 Chrome extensions not contracted (blocked Week 1). **Remediation:** Otter upgrade in flight; consumer LLM and Chrome ext access blocked via Umbrella.

APPENDIX B · CARRIER RIDER FULL MAPPING (2 OF 2)

Requirements 06 through 10.

§4.6 · AI-specific training records

NIST AI RMF GV-4.1, GV-4.2

Sub-controls: (a) training content covers AI specifically · (b) annual cadence · (c) acknowledgment captured · (d) 3-year retention · (e) role-segmented

Evidence at Northgate: Existing HIPAA training already references AI use in clinical workflow. **Enhancement (Deliverable 04 Training Plan):** 4-module AI-specific training adds depth, role-segmented (clinical / office / billing tracks). LMS retains records natively.

§4.7 · DLP coverage of AI domains

NIST AI RMF MS-2.6, MG-2.2

Sub-controls: (a) AI domains enumerated · (b) egress blocked or monitored · (c) bypass requires approval · (d) logs retained · (e) periodic review

Evidence at Northgate: Cisco Umbrella has 6 AI domains in default category but enforcement is off. **Remediation (Week 1):** Enable enforcement on existing 6 + add 12 more from Appendix A inventory = 18 enumerated AI domains under DLP. Logs retain 90 days; quarterly review cycle added to governance committee charter.

§4.8 · Incident response runbook · AI scenarios

NIST AI RMF MG-2.1, MG-2.3

Sub-controls: (a) AI scenarios documented · (b) escalation path defined · (c) tabletop tested annually · (d) lessons-learned integrated · (e) external party notification template

Evidence at Northgate: Existing IR runbook (HIPAA scope) does not include AI scenarios. **Remediation:** AI runbook delivered Day 12 of Sprint (4 scenarios: consumer-AI-with-PHI, vendor-breach, prompt-injection on customer-facing AI, model-output-error in clinical context). **Day 90:** tabletop exercise scheduled.

§4.9 · Governance with AI oversight authority

NIST AI RMF GV-2.1, GV-3.1

Sub-controls: (a) named accountable officer · (b) stop-authority documented · (c) cross-functional review · (d) quarterly minimum cadence · (e) board reporting

Evidence at Northgate: CIO has documented stop-authority via existing IT Steering Committee charter (last updated Q2 2025). **Enhancement:** Add AI-specific standing agenda item; first quarterly meeting scheduled Day 60.

§4.10 · BAA coverage for AI processing of PHI

NIST AI RMF MS-2.7 + HIPAA 45 CFR 164.504(e)

Sub-controls: (a) BAA on file for every PHI-processing AI vendor · (b) BAA scope covers AI workflow · (c) sub-processor disclosure · (d) breach notification SLA · (e) annual BAA review

Evidence at Northgate: Microsoft Copilot E5 BAA covers approved use including AI features. EHR vendor BAA covers embedded AI scribe. **Closing items:** Otter Enterprise+ BAA (Week 2). Consumer LLMs blocked, no BAA needed (removed from PHI workflow).

ENGAGING SHADOW AI LABS

This is what a sample looks like. Yours would carry your name, your tools, your carrier, your gaps.

This report is a sample. Northgate Family Medicine is a composite client constructed from patterns we see across SMB healthcare engagements. The methodology, deliverables, regulatory citations, carrier rider language, and remediation patterns are all real and representative. The names and specific numbers are anonymized.

If your organization is heading into a 2026 cyber insurance renewal, an audit, a customer-due-diligence questionnaire, or any other AI governance forcing function — a Sprint of your own produces a report like this one, tailored to your operational reality, in two weeks at fixed cost.

THREE WAYS TO ENGAGE — FROM HERE

AI Risk Sprint

$5,500 · 2 weeks · fixed scope

The engagement that produces a report like this one. Productized, fixed-price, fixed-deliverables. The front door for most engagements. Sprint fee credits toward Implementation if you decide to proceed.

AI Governance Implementation

$15K–$35K · 6–10 weeks

Closes the gaps the Sprint identifies. DPA repaper, vendor BAA execution, DLP configuration, IR runbook update, control rollout. Begins where the Sprint Roadmap leaves off.

Fractional AI Security Officer

$4.5K–$8.5K/mo · 6 month minimum

Named AI security officer on retainer for organizations that need ongoing carrier liaison, vendor reviews, employee training, and quarterly board readouts. Activates once governance is in place.

Start here

Free assessment: **shadowailabs.com/assess**

20 questions, 6 minutes, a risk score and a personalized teaser report. The right move if you're not sure whether the forcing function justifies the Sprint yet.

DIRECT CONTACT

Peter Kwidzinski, founder · peter@shadowailabs.com · shadowailabs.com
Open to discovery calls without obligation. The Sprint is a fit for organizations with $1–25M revenue and 10–200 employees who are navigating a specific forcing function — cyber insurance renewal, audit, customer DD, or regulatory deadline. If we're not a fit, we'll say so in the discovery call.