WORKBOOK AND GUIDE

Advancing Outcomes and Community Well-Being: Using AI Responsibly in Behavioral Health Organizations

MCTAC Nonprofit Board and Executive Leadership Series

May 21, 2026

This tool was developed for MCTAC in partnership with Jackie Negri (Negri Management Resources LLC) and Doug Golub (Data Potato LLC).

How to Use This Workbook

This workbook is designed to accompany the presentation. Each section includes:

- **Concept summaries**: key ideas from the slides
- **Worksheets**: fill-in, checklist, or scenario to apply to your agency

The workbook covers three areas that mirror the workshop:

1. **Foundations & Frameworks**: AI risks, concerns, opportunities, definitions, best practice frameworks, regulations
2. **Policies & Vendor Oversight**: Acceptable Use Policies (AUPs), vendor risk, human-in-the-loop, and how to handle existing software, new AI modules, personal devices, and external data sharing
3. **Action Tools**: policy worksheet, vendor scorecard, scenario exercises

You will need:

- List of software your agency uses (EHR, billing, scheduling, HRIS, etc.)
- Existing trainings and policies that relate to AI at your agency
- Information about any AI features already present or being added
- Awareness of staff-owned devices and connections to state systems, RHIOs/HIEs, and partner agencies

Contents

Part 1: Foundations & Frameworks

1.1 Key AI Concepts and Definitions from the Presentation

- Large Language Model (LLM) – An AI system trained on vast text data to understand and generate human language. Examples: ChatGPT, Gemini, Claude.
- Generative AI – AI that creates new content (text, images, summaries) based on patterns learned from training data.
- Agentic AI – AI that takes actions on a user's behalf (e.g., drafting and sending a document, adjusting schedules) with limited human intervention.

1.1.1 Additional Definitions

- Open vs. Closed System – Open systems (e.g., public ChatGPT) may use your inputs for training and have no BAA; closed systems keep your data isolated and are typically HIPAA-compliant.
- Prompt Engineering – Crafting clear, structured instructions to get accurate and useful outputs from an AI.
- Grounding / Retrieval-Augmented Generation (RAG) – Connecting AI to your organization's trusted documents (e.g., policies, protocols) to reduce hallucinations and align responses with best practices.
- Human-in-the-Loop (HITL) – A requirement that a person reviews, edits, and approves any AI output before it is used or shared.
- Business Associate Agreement (BAA) – A HIPAA-required contract with any vendor that handles Protected Health Information (PHI). Without a BAA, the vendor cannot lawfully access PHI.
- Audit Logging – Tracking all AI interactions (who, what, when, what data) for oversight and compliance.
- Role-Based Access Control (RBAC) – Limiting who can use or see AI outputs based on their job role.
- Anonymization / De-identification – Removing personal identifiers (name, DOB, Medicaid IDs) from data before it enters an AI model.
- Encryption (at Rest & In Transit) – Scrambling data so it cannot be read if intercepted or stolen, both while stored and while moving.
- NIST - NIST stands for the National Institute of Standards and Technology. It is a federal agency within the U.S. Department of Commerce that promotes innovation, industrial competitiveness, and measurement science. The NIST Cybersecurity Framework and technical standards are recognized

internationally, which is why their AI Risk Management Framework (RMF) is the leading framework in this space.

Why this matters for your agency:

- AI is already inside common software (EHRs, billing systems, Microsoft 365, Google Workspace), often without your explicit knowledge or consent.
- Staff also use public AI LLMs (ChatGPT, Gemini, Claude) on personal phones and agency computers, creating data leakage risks.

1.2 Key Risks (from slides)

Risk	What it means
Bias	AI produces unfair outcomes based on race, disability, gender, etc.
Hallucinations	AI invents facts, citations, or events that never happened.
Privacy breach	Protected Health Information (PHI) entered into public LLMs can be exposed, re-identified, or used for training.
Over-reliance	Staff trust AI outputs without human review, reducing judgment and empathy.

Worksheet 1 – Identify AI-Enabled Systems in Your Environment

Complete one row per AI-enabled system, module, personal device app, or external connection.

Vendor System / Module /	Description	Privacy Level (1–4)	Open or Closed	Does your data train the model?	BAA signed?
Ex: EHR monthly summary	Summarizes progress notes	Level 2 (PHI)	Open	Temporarily	Yes
Ex: Staff ChatGPT on phone	Drafts emails, notes	Level 2 (PHI pasted)	Open	Permanently	No
Ex: RHIO data feed	Shares discharge data	Level 4 (multi-agency PHI)	Closed	No	N/A (covered entity)

Privacy Level key:

1 = Your agency only, no PII

2 = Your agency only, with PII/PHI

3 = Multi-agency data, no PII

4 = Multi-agency data, with PII/PHI

1.3 NIST AI Risk Management Framework (RMF)

A flexible, non-mandatory framework to build trustworthy AI systems.

Function	What it means for your agency
GOVERN	Establish policies, assign roles, create accountability.
MAP	Understand context: who uses AI, for what, with what data?
MEASURE	Test for bias, accuracy, safety – before and during use.
MANAGE	Respond to incidents, monitor continuously, update as needed.

Worksheet 2 – Apply NIST to Your Agency

For each function, write one specific action your agency will take in the next 90 days.

Function	Action
GOVERN	
MAP	
MEASURE	
MANAGE	

1.4 Regulatory Landscape (as of May 2026)

These regulations and frameworks do not constitute a complete list. They are selected to be representative of the types of external requirements that help shape agency controls and policies.

Regulation	Key Provision	AI Relevance & Compliance Action
HIPAA (Security Rule)	Requires annual risk analysis for AI-specific threats; proposed Security Rule updates mandate formal AI policies and vendor BAAs.	Inventory all AI tools; require HIPAA-compliant BAAs from every AI vendor handling PHI.
42 CFR Part 2	Protects Behavioral Health SUD records with heightened consent.	Obtain explicit, separate patient consent before using any AI with SUD data.
CMS (Medicare Advantage & Medicaid)	AI can support but cannot supplant clinical judgment or coverage rules; proposed guardrails ban algorithmic discrimination.	Establish human-in-the-loop controls for AI-influenced decisions; conduct regular bias reviews.
ONC HTI-1	Requires transparency for Predictive Decision Support Interventions (PDSI) in certified EHRs; mandates evaluation of FAVES (Fairness, Appropriateness, Validity, Effectiveness, Safety).	Ask your EHR vendor for PDSI source attributes and risk assessments for any predictive AI.
NYS ITS Policy	For NY State entities: AI cannot make automated final decisions impacting the public; requires human oversight, annual AUP review, and training.	Mandate documented human review for every AI output that affects a person's services or benefits.
NIST AI RMF (Not a Regulation)	Prevailing governance standard for trustworthy AI; integrates with HIPAA risk management.	Use NIST's four functions (Govern, Map, Measure, Manage) to structure your agency's AI policies and risk assessments.

1.5 Role-Based Training Requirements

Every staff member needs different AI knowledge. Use this table to plan your training, understanding that roles, required concepts, and training frequencies will differ with the needs of your agency.

Role	Must know	Training frequency
Direct support / clinical staff	What is prohibited (no PHI in public AI); how to recognize AI in software; when to flag concerns.	Annual + new hire
Supervisors	How to review AI-generated outputs; how to enforce acceptable use policy.	Annual
IT / compliance	Vendor risk assessment; audit log review; BAA management.	Semi-annual
Executives / board	AI risk governance; regulatory duties; approval of AUP.	Annual

Part 2: Policies & Vendor Oversight – Managing Complex AI Environments

2.1 The Reality: AI Is Already Everywhere

Your agency's AI landscape includes:

- **Existing software** that has added AI features (often called "silent AI" – no notification, no opt-out). Examples: EHRs adding smart summaries, billing systems adding predictive analytics, Microsoft 365 Copilot enabled by default.
- **New software** you are considering – you can ask questions before buying.
- **Modules being added** to existing software – these may introduce AI without a full re-procurement.
- **Staff personal devices** (phones, laptops) where they use public AI tools like ChatGPT, Gemini, or Claude for work tasks – often without agency knowledge.
- **External connections** to state systems (e.g., Medicaid portals, OMH/OASAS/OPWDD reporting), RHIOs/HIEs (health information exchanges), and other provider systems – these may also use AI on your data or send your data to AI models.

2.2 AI Use Cases – Safe vs. Unsafe

Not every AI use is appropriate for the services your agency provides in behavioral health. The table below shows examples of **safe** (with safeguards) and **unsafe** (prohibited) applications. The understanding, comfort, risk tolerance, and consent of your stakeholders determines which use cases are analyzed at your agency.

Area	Safe (with human review & closed systems)	Unsafe (prohibited)
Clinical documentation	Draft progress notes from de-identified bullet points	Paste raw PHI into public ChatGPT
Intake & assessment	AI-generated summary of intake form for clinician review	AI decides level of care or denies admission
Scheduling & workforce	Predict no-show risk to offer appointment reminders	Automatically discharge a person for missed visits

Billing & prior auth	Identify missing documentation before submission	Auto-submit claims without human verification
Quality reporting	Flag incomplete data elements in required quality reports	Generate outcome measures without clinical validation
Person-facing tools	Wellness chatbots that disclose AI use and offer opt-out	Unannounced AI monitoring or emotion recognition

2.3 Template – AI Acceptable-Use Policy (AUP)

Customize this template for your agency. It addresses open vs. closed systems, personal devices, and external connections.

Policy Title: Acceptable Use of Artificial Intelligence Tools
Effective Date: _____ // **Review Date:** _____ (at least annually)

1. Scope
Applies to all staff, contractors, volunteers, board members, and any external partners who share data with the agency. Covers agency-owned devices, personal devices used for work, and any third-party systems (state, RHIO, other providers) that interact with agency data.

2. Definitions

- **Open AI system:** Public AI tool (e.g., free ChatGPT, Gemini, Claude) that may use input data for training and does not sign a BAA.
- **Closed AI system:** Agency-approved, HIPAA-compliant AI tool with a signed BAA, data encryption, and no third-party data retention.

3. Approved Tools (Closed Systems Only)
Only the following AI tools are permitted for work-related tasks:

- _____
- _____

All other AI tools are prohibited unless approved in writing by [IT/Compliance/Executive].

4. Prohibited Actions

- Entering any PHI, PII, or SUD records into any open AI system (including ChatGPT, Gemini, Claude, or any public AI).

- Using any AI tool (open or closed) to make final clinical or administrative decisions without human review.
- Installing AI browser extensions or mobile apps on agency devices without approval.
- Using personal devices to access agency data through any AI tool.

5. Human-in-the-Loop
Every AI-generated output must be reviewed, edited, and approved by a designated staff member before use, sharing, or filing.

6. External Connections (State, RHIO, Other Providers)
Before sharing data with any external system that uses AI (including state reporting portals, RHIOs, or partner agencies), the agency will:

- Request written disclosure of AI use.
- Ensure a BAA or equivalent data protection agreement is in place.
- Limit data to the minimum necessary.

7. Training
All staff will complete annual training on this policy, including how to recognize AI features in existing software and how to report unauthorized AI use.

8. Enforcement
Violations may result in disciplinary action up to and including termination and may be reported to regulatory authorities (e.g., HHS OCR, OMIG).

9. Policy Review
Reviewed annually and after any material change in AI tools, regulations, or external connections.

2.4 Vendor Risk Questionnaire – For Any AI-Enabled System

Use this questionnaire for new software, major updates to existing software, and any external partner that sends or receives data.

Question	**Response**	**Follow-up needed?**
Does your system use any AI or LLM? (If no, stop here.)	☐ Yes ☐ No	
*Is the AI an **open** system (data trains the model) or **closed** (your data is isolated)?*	☐ Open ☐ Closed	If open, do not use for PHI.
Does your system require a BAA? (If PHI is involved, yes.)	☐ Signed ☐ Not signed ☐ Required but not provided	
*Is AI use **optional** (can be disabled) or **mandatory**?*	☐ Optional ☐ Mandatory	If mandatory and no BAA, do not use.
Do you have SOC-2, HITRUST, or independent audit reports?	☐ Yes ☐ No	Request copies.
Can we retain audit logs of all AI interactions?	☐ Yes ☐ No	
*Does the contract include **data liberation** (ability to export all data) and **breach notification** clauses?*	☐ Yes ☐ No	Negotiate if missing.

2.5 Human-in-the-Loop & Ethics Checklist – For Every AI Use Case

Safeguard	**In place?**	**Responsible roles**
Every AI output is reviewed by a human before use.	☐ Yes ☐ No	
The reviewer has authority to override or reject AI recommendations.	☐ Yes ☐ No	
The people you support are informed when AI is used in their care (including via external connections).	☐ Yes ☐ No	
There is a process for individuals to opt out of AI-driven features.	☐ Yes ☐ No	
A designated committee or leader reviews AI use quarterly, including outputs from external systems.	☐ Yes ☐ No	

Part 3: Action Tools

Worksheet 3 – Draft Your AUP (Fill-In)

Based on template 2.2 (AUP), fill in your agency's specifics.

Our agency will permit the following closed-system AI tools:

1. _____
2. _____

We will prohibit the following AI tools (including on personal devices):

1. _____
2. _____

The following roles are responsible for reviewing AI outputs from agency systems:

- Role: _____ – for: _____
- Role: _____ – for: _____

For external connections (state, RHIO, other providers), the person responsible for requesting AI disclosure is:
_____ (role)

Staff will be trained on AI acceptable use at least:
☐ Annually ☐ Semi-annually ☐ Quarterly

Our policy will be reviewed by: (committee/position) _____ every _____ months.

Worksheet 4 – Vendor Scorecard (For Each AI System or Module)

Rate each vendor or system on a scale of 1 (poor) to 5 (excellent).

Criteria	Score (1–5)
Signed BAA (if PHI involved)	
No open-system training on our data	
AI is optional (can be disabled)	
Bias test results available	
SOC-2 / HITRUST	
Audit logs accessible	
Exit plan (data liberation) in contract	
Total score (max 35)	

Action guide:

- 30–35: Use with confidence
- 20–29: Use with conditions (negotiate improvements)
- Below 20: Do not use; seek legal review

Worksheet 5 – Training Plan

Fill in who will be trained, by when, and on what topics. Add roles, as appropriate, for your agency's organizational and functional design.

Role group	Training topic	Delivery method (e.g., online module, live)	Completion deadline
Direct support	**Prohibited AI actions + how to report**		
Clinicians	**Human review of AI outputs**		
Supervisors	**Enforcing AUP, reviewing audits**		
IT/Compliance	**Vendor risk questionnaire, BAA**		
Board	**AI governance & regulatory duties**		

Worksheet 6 – Scenario Exercises

Discuss each scenario with your team. Write your response.

Scenario 1 – Staff uses public AI on personal phone
A clinician copies a person's psychiatric assessment into the free ChatGPT app on their personal phone to "reword it more clearly."

- Which policy was violated?
- What immediate steps should you take?
- How do you address personal device use?

Scenario 2 – EHR vendor adds AI module without notice
Your EHR vendor releases an update that includes an AI "risk prediction" feature. It is enabled by default. The vendor has not signed a BAA for this feature. The feature sends data to an external LLM.

- What risks does this create?
- Who in your agency needs to be notified?
- What is your next action (technical and legal)?

Scenario 3 – Outside provider starts using AI
An outside care management provider entity that you send discharge summaries to announces it will use an AI model to "enhance care coordination." The provider says no BAA is needed because it is already a covered entity.

- Do you have any rights to limit this use?
- What questions should you ask the RHIO?
- Should you change what data you share?

Appendix – Reference Materials

NIST AI RMF : *https://www.nist.gov/itl/ai-risk-management-framework*

- **GOVERN** → Policies, roles, accountability
- **MAP** → Context, use cases, risks
- **MEASURE** → Bias, accuracy, safety
- **MANAGE** → Incident response, monitoring

Six Non-Negotiable Questions for Any AI Vendor or External Connection

1. Is AI used? If yes, **open or closed** system?
2. Does our data train the model?
3. Is a **BAA** signed (if PHI is involved)?
4. Can we **disable** the AI if we choose?
5. Do you provide **audit logs** and **bias test results**?
6. Does the contract include **data liberation** and **breach notification**?

Next Steps for Your Agency – A 90-Day Plan

By end of week 1:
☐ Assign an AI lead (person or committee) for this work.

By end of month 1:
☐ Complete inventory of all AI in use (agency systems, personal devices, external connections).
☐ Request BAAs or AI disclosures from all vendors and external partners.

By end of month 2:
☐ Draft or update your AI acceptable-use policy (use template in this workbook).
☐ Present policy to board for approval.

By end of month 3:
☐ Train all staff on the policy and prohibited tools.
☐ Disable any AI features that lack BAAs or that expose PHI to open systems.
☐ Establish quarterly review process for AI incidents and new AI features.