

[Your Company Name] — AI Fundamentals & Safety (LMS Module) —

Outline & Presenter Script

Version: 1.0 | Date: 2025-10-14 | Owner: [Your Company Name: Role] | Duration: ~20 minutes

Prerequisites: None | Audience: All employees, contractors, and vendors with system access

Disclaimer: This content is a framework for [Your Company Name] to customize with counsel; state/local laws and client contracts may require changes.

Learning Objectives

- Recognize approved vs. prohibited AI use and why MSP/MTSP environments require stricter controls.
- Apply prompt hygiene and a “Do-Not-Paste” list when using enterprise AI tools (e.g., Microsoft 365 Copilot Enterprise, Azure OpenAI, Amazon Bedrock, or equivalent).
- Describe responsibilities for logging, human-in-the-loop review, incident reporting, and records retention.
- Identify when higher-risk features (e.g., facial recognition/biometrics) require elevated approvals or are disallowed.
- Complete the policy acknowledgment required for AI tool access.

Module Flow (20 minutes)

1. Why this matters (2:00) — client trust, privacy contracts, audit expectations.
2. Approved tools & boundaries (3:00) — allow-listed enterprise AI; consumer endpoints disabled; “no training on client data” posture.
3. Do-Not-Paste + prompt hygiene (3:30) — secrets/keys, device logs w/ identifiers, chain-of-custody, unreleased code; minimum-necessary data.
4. Human-in-the-loop & consequential decisions (3:00) — what requires sign-off; examples (HR, legal, safety).
5. Sector overlays (3:00) — HIPAA/PHI, GLBA/financial, PCI DSS/payment, CJIS/public safety; routing FOIA/open-records via Legal.
6. Logging, incident reporting, audits (2:30) — privacy-preserving logs; 72-hour client notice after confirmation (if contractually required).
7. Quiz + acknowledgment (3:00) — must score $\geq 85\%$ and e-sign policy.

Presenter Script (bullet prompts)

- Stakes: mishandled prompts can leak PI/PHI/secrets or violate client DPAs.
- Allow-list: enterprise tenants only; disable consumer AI; map to your IdP and DLP policies.
- Do-Not-Paste: never paste credentials/keys, identifiable device logs, chain-of-custody info, unreleased code/firmware, or client source code.
- Human-in-the-loop: capture sign-offs for consequential actions; no auto-publish to clients.
- Risky features: FRT/biometrics default-deny; require executive approval + evidence + client consent + Legal sign-off if ever considered.
- Records & logging: store work only in approved repositories with retention tags; IT may log tool IDs and output hashes where permitted.
- Incident playbook: stop, capture details, notify Security immediately; do not attempt to fix silently.

LMS Build Checklist

- Create course shell: “AI Fundamentals & Safety (TAG Member)” (~20 minutes).
- Upload slides/video; attach policy PDF and quick reference (wallet card).
- Add the 12-question quiz ($\geq 85\%$ to pass; 2 attempts).
- Add e-signature acknowledgment of the policy.
- Assign to All Staff; enforce completion before AI tool access (IdP group mapping).
- Reporting: completion, score, e-sign timestamp; keep in audit folder.