SYLLABUS · TECHFOUNDATIONS

Security & AI Awareness

Fully online · self-paced · built for your phone · free

1. Course Title & Description

Security & AI Awareness is a short, self-paced course for working adults who are being told to use AI at work but were never shown how to do it safely. AI didn't just change your tools — it changed the tricks that get played on you and the power you have to work smarter. Across eleven short lessons in three movements — **Defend** (protect yourself), **Adopt** (use AI safely), and **Keep** (the habit) — you'll learn to spot AI-era scams, understand where your data goes when you paste, redact or rephrase before sending, catch hidden-instruction traps, grant AI tools the least access that does the job, respond in the first five minutes when something slips, verify AI answers before relying on them, and finish with a four-question check you can run before any AI use, for the rest of your career.

Every lesson runs in your browser, works one-handed on a phone, and follows the same arc: a short video opener, a "your call" decision before any explanation, the concept narrated over slides, and hands-on practice on your own real examples — never canned scenarios alone.

2. Course Learning Outcomes

By the end of this course (about two hours of work), you will be able to:

1. **CLO 1 · Classify risk** — Classify how AI changes a threat — an old threat supercharged, or a genuinely new AI-native risk — and update the matching habit. *(Lessons 1, 2, 7)*
2. **CLO 2 · Trace data** — Trace where a piece of sensitive data lives and every moment it leaves your control, including the moment it is pasted into an AI tool. *(Lessons 3, 5)*
3. **CLO 3 · Redact or rephrase** — Apply the redact-or-rephrase workflow before sending content to an AI tool. *(Lessons 5, 6 + cohort discussion)*
4. **CLO 4 · Review policy** — Interpret what an AI tool actually does with your input — what it sees, stores, and trains on — and what access it is asking for. *(Lessons 4, 8 + policy-skim discussion)*

5. **CLO 5 · Map your data** — Construct a one-page Data & AI-Use Map — the capstone proof of transfer. *(Lesson 11 capstone, building on Lesson 3)*
6. **CLO 6 · Evaluate & respond** — Evaluate AI output and new AI situations against your baseline: verify before relying, respond in the first five minutes when something slips, and run the four-question check on any new use. *(Lessons 9, 10, 11)*

Each outcome is taught and practiced at Bloom's Apply level or above across at least two lesson objectives; the map below shows every pairing. (CLO set revised July 2026: the earlier device-setup outcome moved to a companion course; data-tracing took its place.)

3. Alignment Map

Read each row as: where you learn it → what you'll be able to do → how you show it.

LESSON → OBJECTIVE → EVIDENCE

Lesson	Objective (what you'll be able to do)	Evidence
1. The game changed (but you're not starting over)	Given a real threat or AI capability, the learner classifies how AI changes the risk — an OLD threat supercharged, or a GENUINELY NEW AI-native risk — and names what it means for their own habits. (Bloom Evaluate/Apply: the learner sorts real examples and articulates the implication for how they work.)	in-lesson decisions, hands-on practice on your own example, quick check
2. When you can't trust your eyes (or ears)	Given an AI-era identity attack (a flawless phishing message, a cloned-voice request, or a deepfake), the learner decides the right response: verify the request through a separate trusted channel and rely on protections AI can't fake (MFA), rather than judging by how convincing it looks or sounds. (Bloom Apply+: pass = chooses out-of-band verification AND identifies MFA as the account protection.)	in-lesson decisions, hands-on practice on your own example, quick check
3. Where your data lives, and when it leaves	Given a real work scenario, the learner traces a piece of sensitive data: where it sits (at rest), when it travels (in motion), and the moment it 'leaves the building' — crosses from systems they control to systems someone else controls. (Bloom Apply+: pass = identifies the sensitive data, its at-rest/in-motion states, and correctly flags the leaving-the-building moment.)	in-lesson decisions, hands-on practice on your own example, quick check

4. What an AI actually does with what you type	Given a common belief about what an AI 'does' (it looks things up, it understands you, it keeps your input private), the learner judges whether it's accurate and states what the AI actually does — building the working mental model that the rest of the course's habits depend on. (Bloom Evaluate/Apply+: pass = correctly flags misconceptions and replaces each with what the model actually does.)	in-lesson decisions, hands-on practice on your own example, quick check
5. Prompt as egress: what you paste leaves	Given a real work prompt the learner is about to send to an AI tool, triage it: flag each element that crosses the trust boundary and state why it can't be unsent, so they decide what is safe to export before they hit enter. (Bloom Evaluate/Apply+: the learner produces a triage of a real prompt; pass = every sensitive element flagged AND each carries its egress reason.)	in-lesson decisions, hands-on practice on your own example, quick check
6. Safe prompting: use AI without leaking	Given a real mixed-sensitivity work prompt, the learner produces a redacted version that removes every element crossing the trust boundary using the right move for each (placeholder / generalize / summarize) while keeping the request answerable. (Bloom Apply+: pass = no crossing element remains AND the task is still clearly answerable AND each redaction uses an appropriate move.)	in-lesson decisions, hands-on practice on your own example, quick check
7. The AI threat model: hidden instructions	Given a real scenario where someone has AI read or act on content from a source they don't control, the learner assesses whether it's exposed to prompt injection and names the safe move. (Bloom Evaluate/Apply+: pass = correctly judges the exposure AND names a safe move that treats the outside content as untrusted.)	in-lesson decisions, hands-on practice on your own example, quick check
8. Granting AI access: the smallest keyring	Given a real AI connection request (a permissions/scopes screen) and the task it's for, the learner decides the least-privilege grant: which permissions the task actually needs and which exceed it. (Bloom Evaluate/Apply+: pass = grants only what the task needs, flags the over-asks, and justifies read-only/narrow over broad.)	in-lesson decisions, hands-on practice on your own example, quick check
9. When it goes wrong: what to do in the first five minutes	Given a scenario where someone has just leaked something into an AI, the learner produces the right immediate response: the correct containment move for what leaked, who to tell, and what to record. (Bloom Apply+: pass = names the correct kill-	in-lesson decisions, hands-on practice on your

	the-value move, escalates fast and honestly, and records what/when/where.)	own example, quick check
10. Don't trust the confidence: when to verify AI	Given an AI output and how it will be used, the learner decides the right verification move: judge the stakes, identify what needs checking, and verify it against an independent trusted source (not the AI itself). (Bloom Apply+: pass = calibrates verification to the stakes AND names an independent source to check against, rather than trusting the AI's confidence.)	in-lesson decisions, hands-on practice on your own example, quick check
11. Four questions before you use AI	Given a real AI-use situation, the learner runs a short set of responsible-use questions and reaches a decision: proceed, modify, or ask first (check the rule / escalate to a person). (Bloom Apply+: pass = applies the questions to the situation AND lands on a sound proceed/modify/ask-first decision with reasoning.)	in-lesson decisions, hands-on practice on your own example, quick check
Capstone — Data & AI-Use Map (in Lesson 11)	Construct your one-page map: sensitive data kinds × where they live × AI-tool calls, plus your four-question rule (CLO 5)	hands-on capstone artifact, self-checked against published criteria
Final knowledge check	Demonstrate the course outcomes across realistic scenarios	12-item check, 80% mastery gate, unlimited retries with per-item feedback

4. Instructor Contact

Instructor / designer: Glen Buchanan — TechFoundations
Contact: glen.m.buchanan@gmail.com · **Response time:** within one business day, usually sooner.

A note on background so you know you're in good hands: about ten years in IT and security, holding the CISSP, CISM, CISA, and CRISC. You don't need any of that — I do, so you don't have to.

5. Course Prerequisites

None technical. If you can send an email and install an app, you're ready. You do **not** need any security or IT background, and you don't need a paid AI account — Claude's free tier works for every practice task in this course.

6. Technology Requirements

What	Details	Cost
A phone, tablet, or computer	Any current one, with an up-to-date web browser	—
Internet connection	For the browser-based lessons and short videos	—
The course platform	Runs in your browser — nothing to install; progress is tracked via xAPI (completion counts, never your content)	Free
A Claude account	Claude's free tier, for the practice tasks (claude.ai)	Free
Headphones (optional)	All videos and narration are fully captioned, so audio is a nicety, not a need	—

7. Grading Policy, Grade Scale & Weights

This is a **mastery-and-completion** course — you're not competing with anyone. You earn a **Certificate of Completion** by completing all eleven lessons (including each lesson's hands-on practice) and clearing the **80% mastery gate** on the final knowledge check. Check items give instant per-item feedback and can be retried as many times as you need.

Component	Weight
Lesson completion — all 11 lessons, including each lesson's decisions and hands-on practice	60%
Final knowledge check (mastery-gated at 80%, unlimited retries)	40%

Result	Meaning

Complete with Distinction	All lessons + 90–100% on the knowledge check — certificate awarded
Complete	All lessons + 80–89% — certificate awarded
Not yet	Below 80% — per-item feedback shows exactly what to review; retry freely, no penalty

8. Late Policy

The course is **self-paced: there are no hard deadlines by default** — start and finish on your schedule, and revisit any lesson as often as you like. If you're taking it in a **facilitated cohort** with weekly checkpoints: checkpoints are soft. Work is accepted with an automatic 48-hour grace window and **no penalty**, and a longer extension is granted on request, for any reason, before or after the checkpoint. Asking for an extension never costs you anything — the goal is your learning, not the clock.

9. Schedule of Instructional Events

A welcome video, eleven short lessons, and a final knowledge check — about **two hours total**, in ~9-minute lessons you can take in any sitting or spread across two weeks (a suggested pace: one movement per week).

#	Lesson	Movement	Time
—	**Welcome & overview** (video)	—	~1 min
1	**The game changed (but you're not starting over)**	Defend — protect yourself	~8 min
2	**When you can't trust your eyes (or ears)**	Defend — protect yourself	~9 min
3	**Where your data lives, and when it leaves**	Defend — protect yourself	~9 min
4	**What an AI actually does with what you type**	Defend — protect yourself	~9 min

5	**Prompt as egress: what you paste leaves**	Defend — protect yourself	~9 min
6	**Safe prompting: use AI without leaking**	Defend — protect yourself	~9 min
7	**The AI threat model: hidden instructions**	Defend — protect yourself	~9 min
8	**Granting AI access: the smallest keyring**	Adopt — use AI safely	~9 min
9	**When it goes wrong: what to do in the first five minutes**	Adopt — use AI safely	~9 min
10	**Don't trust the confidence: when to verify AI**	Adopt — use AI safely	~9 min
11	**Four questions before you use AI**	Keep — the habit	~9 min
—	**Final knowledge check** + course feedback	Check	~10 min

10. Academic Honesty & Responsible AI Use

Do your own thinking. The work you submit — your redact-or-rephrase calls, your decisions, your safe versions — should be your own judgment about your own situation. That's the whole point; there's no answer key to copy.

Using AI in this course is allowed and encouraged — with two rules. (1) Use AI as a thought partner, not a ghost-writer: if an AI helped shape something you submit in a cohort setting, say so in a one-line note. (2) **Never paste real sensitive data** — yours or anyone else's — into an AI tool as part of an exercise. This is the exact habit the course teaches, so we hold ourselves to it here. Cite any outside source you quote.

11. Accommodations for Learners with Disabilities

Accessibility is built in, not bolted on. If you still need something, ask — you never have to disclose a diagnosis to get it.

Already in place for everyone: keyboard navigation and screen-reader-real HTML (tables are tables, not images); 4.5:1 text contrast; every video and narration clip carries accurate closed

captions; body text at a grade 8–10 reading level; self-paced with no time limits anywhere; knowledge checks retry freely; a pause control on all narration.

Available on request (no documentation required): alternative formats for any task (voice or video in place of typed work, large-print or plain-text versions), a one-on-one walk-through of any step, or a reduced-motion, low-stimulation version of any lesson. Contact the instructor (§4) — requests are handled privately, usually within one business day.

12. References

CAST. (2024). *Universal Design for Learning guidelines version 3.0*. https://udlguidelines.cast.org/

Mayer, R. E. (2021). *Multimedia learning* (3rd ed.). Cambridge University Press.

Merrill, M. D. (2002). First principles of instruction. *Educational Technology Research and Development, 50*(3), 43–59.

Online Learning Consortium / SUNY. (2024). *OSCQR course design review scorecard*. https://oscqr.suny.edu/

W3C. (2023). *Web Content Accessibility Guidelines (WCAG) 2.2*. https://www.w3.org/TR/WCAG22/

13. PDF Version

This syllabus is available as a downloadable PDF, and is linked from the course menu on every screen.